

avr-libc
2.0.0

Generated by Doxygen 1.6.3

Wed May 11 11:27:14 2022

Contents

1	AVR Libc	1
1.1	Introduction	1
1.2	General information about this library	1
1.3	Supported Devices	2
1.4	avr-libc License	12
2	Toolchain Overview	13
2.1	Introduction	13
2.2	FSF and GNU	13
2.3	GCC	14
2.4	GNU Binutils	14
2.5	avr-libc	16
2.6	Building Software	16
2.7	AVRDUDE	16
2.8	GDB / Insight / DDD	16
2.9	AVaRICE	17
2.10	SimulAVR	17
2.11	Utilities	17
2.12	Toolchain Distributions (Distros)	17
2.13	Open Source	17
3	Memory Areas and Using malloc()	18
3.1	Introduction	18
3.2	Internal vs. external RAM	19
3.3	Tunables for malloc()	19
3.4	Implementation details	21
4	Memory Sections	22
4.1	The .text Section	23
4.2	The .data Section	23
4.3	The .bss Section	23
4.4	The .eeprom Section	23
4.5	The .noinit Section	23

4.6	The .initN Sections	24
4.7	The .finiN Sections	25
4.8	The .note.gnu.avr.deviceinfo Section	26
4.9	Using Sections in Assembler Code	27
4.10	Using Sections in C Code	27
5	Data in Program Space	28
5.1	Introduction	28
5.2	A Note On const	28
5.3	Storing and Retrieving Data in the Program Space	29
5.4	Storing and Retrieving Strings in the Program Space	30
5.5	Caveats	32
6	avr-libc and assembler programs	32
6.1	Introduction	32
6.2	Invoking the compiler	33
6.3	Example program	33
6.4	Pseudo-ops and operators	37
7	Inline Assembler Cookbook	38
7.1	GCC asm Statement	39
7.2	Assembler Code	40
7.3	Input and Output Operands	41
7.4	Clobbers	45
7.5	Assembler Macros	47
7.6	C Stub Functions	48
7.7	C Names Used in Assembler Code	49
7.8	Links	49
8	How to Build a Library	50
8.1	Introduction	50
8.2	How the Linker Works	50
8.3	How to Design a Library	50
8.4	Creating a Library	51
8.5	Using a Library	52

9	Benchmarks	52
9.1	A few of libc functions.	53
9.2	Math functions.	54
10	Porting From IAR to AVR GCC	55
10.1	Introduction	55
10.2	Registers	56
10.3	Interrupt Service Routines (ISRs)	56
10.4	Intrinsic Routines	57
10.5	Flash Variables	57
10.6	Non-Returning main()	58
10.7	Locking Registers	59
11	Frequently Asked Questions	59
11.1	FAQ Index	59
11.2	My program doesn't recognize a variable updated within an interrupt routine	61
11.3	I get "undefined reference to..." for functions like "sin()"	61
11.4	How to permanently bind a variable to a register?	62
11.5	How to modify MCUCR or WDTCR early?	62
11.6	What is all this _BV() stuff about?	63
11.7	Can I use C++ on the AVR?	63
11.8	Shouldn't I initialize all my variables?	64
11.9	Why do some 16-bit timer registers sometimes get trashed?	65
11.10	How do I use a #define'd constant in an asm statement?	65
11.11	Why does the PC randomly jump around when single-stepping through my program in avr-gdb?	66
11.12	How do I trace an assembler file in avr-gdb?	67
11.13	How do I pass an IO port as a parameter to a function?	68
11.14	What registers are used by the C compiler?	70
11.15	How do I put an array of strings completely in ROM?	71
11.16	How to use external RAM?	73
11.17	Which -O flag to use?	74
11.18	How do I relocate code to a fixed address?	74
11.19	My UART is generating nonsense! My ATmega128 keeps crashing! Port F is completely broken!	75

11.20	Why do all my "foo...bar" strings eat up the SRAM?	75
11.21	Why does the compiler compile an 8-bit operation that uses bitwise operators into a 16-bit operation in assembly?	76
11.22	How to detect RAM memory and variable overlap problems?	77
11.23	Is it really impossible to program the ATtinyXX in C?	77
11.24	What is this "clock skew detected" message?	77
11.25	Why are (many) interrupt flags cleared by writing a logical 1?	78
11.26	Why have "programmed" fuses the bit value 0?	79
11.27	Which AVR-specific assembler operators are available?	79
11.28	Why are interrupts re-enabled in the middle of writing the stack pointer?	79
11.29	Why are there five different linker scripts?	80
11.30	How to add a raw binary image to linker output?	80
11.31	How do I perform a software reset of the AVR?	81
11.32	I am using floating point math. Why is the compiled code so big? Why does my code not work?	82
11.33	What pitfalls exist when writing reentrant code?	82
11.34	Why are some addresses of the EEPROM corrupted (usually address zero)?	85
11.35	Why is my baud rate wrong?	86
11.36	On a device with more than 128 KiB of flash, how to make function pointers work?	86
11.37	Why is assigning ports in a "chain" a bad idea?	86
12	Building and Installing the GNU Tool Chain	87
12.1	Building and Installing under Linux, FreeBSD, and Others	87
12.2	Required Tools	88
12.3	Optional Tools	88
12.4	GNU Binutils for the AVR target	89
12.5	GCC for the AVR target	90
12.6	AVR LibC	90
12.7	AVRDUDE	91
12.8	GDB for the AVR target	91
12.9	SimulAVR	92
12.10	AVaRICE	92
12.11	Building and Installing under Windows	93

12.12 Tools Required for Building the Toolchain for Windows	93
12.13 Building the Toolchain for Windows	96
13 Using the GNU tools	101
13.1 Options for the C compiler avr-gcc	101
13.1.1 Machine-specific options for the AVR	101
13.1.2 Selected general compiler options	110
13.2 Options for the assembler avr-as	111
13.2.1 Machine-specific assembler options	111
13.2.2 Examples for assembler options passed through the C compiler	112
13.3 Controlling the linker avr-ld	113
13.3.1 Selected linker options	113
13.3.2 Passing linker options from the C compiler	114
14 Compiler optimization	115
14.1 Problems with reordering code	115
15 Using the avrdude program	117
16 Release Numbering and Methodology	119
16.1 Release Version Numbering Scheme	119
16.2 Releasing AVR Libc	119
16.2.1 Creating an SVN branch	119
16.2.2 Making a release	120
17 Acknowledgments	122
18 Todo List	123
19 Deprecated List	123
20 Module Index	124
20.1 Modules	124
21 Data Structure Index	126
21.1 Data Structures	126
22 File Index	126

22.1 File List	126
23 Module Documentation	128
23.1 <alloca.h>: Allocate space in the stack	128
23.1.1 Detailed Description	128
23.1.2 Function Documentation	128
23.2 <assert.h>: Diagnostics	129
23.2.1 Detailed Description	129
23.2.2 Define Documentation	129
23.3 <ctype.h>: Character Operations	130
23.3.1 Detailed Description	130
23.3.2 Function Documentation	131
23.4 <errno.h>: System Errors	132
23.4.1 Detailed Description	133
23.4.2 Define Documentation	133
23.4.3 Variable Documentation	133
23.5 <inttypes.h>: Integer Type conversions	133
23.5.1 Detailed Description	136
23.5.2 Define Documentation	137
23.5.3 Typedef Documentation	147
23.6 <math.h>: Mathematics	147
23.6.1 Detailed Description	149
23.6.2 Define Documentation	149
23.6.3 Function Documentation	154
23.7 <setjmp.h>: Non-local goto	160
23.7.1 Detailed Description	160
23.7.2 Function Documentation	161
23.8 <stdint.h>: Standard Integer Types	162
23.8.1 Detailed Description	165
23.8.2 Define Documentation	165
23.8.3 Typedef Documentation	171
23.9 <stdio.h>: Standard IO facilities	174
23.9.1 Detailed Description	175
23.9.2 Define Documentation	178

23.9.3 Typedef Documentation	181
23.9.4 Function Documentation	181
23.10<stdlib.h>: General utilities	192
23.10.1 Detailed Description	193
23.10.2 Define Documentation	193
23.10.3 Typedef Documentation	194
23.10.4 Function Documentation	194
23.10.5 Variable Documentation	202
23.11<string.h>: Strings	203
23.11.1 Detailed Description	204
23.11.2 Define Documentation	204
23.11.3 Function Documentation	204
23.12<time.h>: Time	216
23.12.1 Detailed Description	217
23.12.2 Define Documentation	218
23.12.3 Typedef Documentation	219
23.12.4 Enumeration Type Documentation	219
23.12.5 Function Documentation	220
23.13<avr/boot.h>: Bootloader Support Utilities	225
23.13.1 Detailed Description	225
23.13.2 Define Documentation	227
23.14<avr/cpufunc.h>: Special AVR CPU functions	231
23.14.1 Detailed Description	232
23.14.2 Define Documentation	232
23.14.3 Function Documentation	232
23.15<avr/eeprom.h>: EEPROM handling	232
23.15.1 Detailed Description	233
23.15.2 Define Documentation	234
23.15.3 Function Documentation	235
23.16<avr/fuse.h>: Fuse Support	237
23.17<avr/interrupt.h>: Interrupts	240
23.17.1 Detailed Description	240
23.17.2 Define Documentation	257
23.18<avr/io.h>: AVR device-specific IO definitions	260

23.18.1 Detailed Description	260
23.18.2 Define Documentation	261
23.19<avr/lock.h>: Lockbit Support	261
23.20<avr/pgmspace.h>: Program Space Utilities	264
23.20.1 Detailed Description	266
23.20.2 Define Documentation	266
23.20.3 Typedef Documentation	270
23.20.4 Function Documentation	274
23.21<avr/power.h>: Power Reduction Management	288
23.21.1 Detailed Description	288
23.21.2 Function Documentation	291
23.22Additional notes from <avr/sfr_defs.h>	292
23.23<avr/sfr_defs.h>: Special function registers	293
23.23.1 Detailed Description	293
23.23.2 Define Documentation	294
23.24<avr/signature.h>: Signature Support	295
23.25<avr/sleep.h>: Power Management and Sleep Modes	296
23.25.1 Detailed Description	296
23.25.2 Function Documentation	297
23.26<avr/version.h>: avr-libc version macros	298
23.26.1 Detailed Description	298
23.26.2 Define Documentation	299
23.27<avr/wdt.h>: Watchdog timer handling	299
23.27.1 Detailed Description	300
23.27.2 Define Documentation	301
23.27.3 Function Documentation	303
23.28<util/atomic.h> Atomically and Non-Atomically Executed Code Blocks	303
23.28.1 Detailed Description	303
23.28.2 Define Documentation	305
23.29<util/crc16.h>: CRC Computations	306
23.29.1 Detailed Description	306
23.29.2 Function Documentation	307
23.30<util/delay.h>: Convenience functions for busy-wait delay loops . . .	310
23.30.1 Detailed Description	310

23.30.2 Define Documentation	311
23.30.3 Function Documentation	311
23.31 <util/delay_basic.h>: Basic busy-wait delay loops	313
23.31.1 Detailed Description	313
23.31.2 Function Documentation	313
23.32 <util/parity.h>: Parity bit generation	314
23.32.1 Detailed Description	314
23.32.2 Define Documentation	314
23.33 <util/setbaud.h>: Helper macros for baud rate calculations	314
23.33.1 Detailed Description	315
23.33.2 Define Documentation	316
23.34 <util/twi.h>: TWI bit mask definitions	317
23.34.1 Detailed Description	318
23.34.2 Define Documentation	318
23.35 <compat/deprecated.h>: Deprecated items	321
23.35.1 Detailed Description	322
23.35.2 Define Documentation	322
23.35.3 Function Documentation	324
23.36 <compat/ina90.h>: Compatibility with IAR EWB 3.x	324
23.37 Demo projects	325
23.37.1 Detailed Description	325
23.38 Combining C and assembly source files	326
23.38.1 Hardware setup	326
23.38.2 A code walkthrough	327
23.38.3 The source code	329
23.39 A simple project	329
23.39.1 The Project	329
23.39.2 The Source Code	331
23.39.3 Compiling and Linking	333
23.39.4 Examining the Object File	333
23.39.5 Linker Map Files	338
23.39.6 Generating Intel Hex Files	340
23.39.7 Letting Make Build the Project	341
23.39.8 Reference to the source code	343

23.40	A more sophisticated project	343
23.40.1	Hardware setup	344
23.40.2	Functional overview	347
23.40.3	A code walkthrough	347
23.40.4	The source code	350
23.41	Using the standard IO facilities	350
23.41.1	Hardware setup	350
23.41.2	Functional overview	352
23.41.3	A code walkthrough	352
23.41.4	The source code	357
23.42	Example using the two-wire interface (TWI)	357
23.42.1	Introduction into TWI	358
23.42.2	The TWI example project	358
23.42.3	The Source Code	358
24	Data Structure Documentation	362
24.1	div_t Struct Reference	362
24.1.1	Detailed Description	362
24.1.2	Field Documentation	362
24.2	ldiv_t Struct Reference	363
24.2.1	Detailed Description	363
24.2.2	Field Documentation	363
24.3	tm Struct Reference	363
24.3.1	Detailed Description	363
24.3.2	Field Documentation	364
24.4	week_date Struct Reference	365
24.4.1	Detailed Description	365
24.4.2	Field Documentation	365
25	File Documentation	365
25.1	assert.h File Reference	365
25.1.1	Detailed Description	366
25.2	atoi.S File Reference	366
25.2.1	Detailed Description	366
25.3	atol.S File Reference	366

25.3.1 Detailed Description	366
25.4 atomic.h File Reference	366
25.4.1 Detailed Description	366
25.5 boot.h File Reference	366
25.5.1 Detailed Description	367
25.6 cpufunc.h File Reference	367
25.6.1 Detailed Description	367
25.7 crc16.h File Reference	367
25.7.1 Detailed Description	367
25.8 ctype.h File Reference	367
25.8.1 Detailed Description	368
25.9 delay.h File Reference	368
25.9.1 Detailed Description	368
25.10delay_basic.h File Reference	368
25.10.1 Detailed Description	368
25.11errno.h File Reference	368
25.11.1 Detailed Description	369
25.12fdevopen.c File Reference	369
25.12.1 Detailed Description	369
25.13fuse.h File Reference	369
25.13.1 Detailed Description	369
25.14interrupt.h File Reference	369
25.14.1 Detailed Description	370
25.15inttypes.h File Reference	370
25.15.1 Detailed Description	372
25.16io.h File Reference	372
25.16.1 Detailed Description	372
25.17lock.h File Reference	372
25.17.1 Detailed Description	372
25.18math.h File Reference	372
25.18.1 Detailed Description	375
25.19parity.h File Reference	375
25.19.1 Detailed Description	375
25.20pgmspace.h File Reference	375

25.20.1 Detailed Description	377
25.21 power.h File Reference	377
25.21.1 Detailed Description	377
25.21.2 Define Documentation	377
25.22 setbaud.h File Reference	378
25.22.1 Detailed Description	378
25.23 setjmp.h File Reference	378
25.23.1 Detailed Description	378
25.24 signature.h File Reference	378
25.24.1 Detailed Description	378
25.25 sleep.h File Reference	378
25.25.1 Detailed Description	378
25.26 stdint.h File Reference	378
25.26.1 Detailed Description	381
25.27 stdio.h File Reference	381
25.27.1 Detailed Description	383
25.28 stdlib.h File Reference	383
25.28.1 Detailed Description	384
25.29 string.h File Reference	384
25.29.1 Detailed Description	385
25.30 time.h File Reference	385
25.30.1 Detailed Description	387
25.31 twi.h File Reference	387
25.31.1 Detailed Description	388
25.32 wdt.h File Reference	388
25.32.1 Detailed Description	388

1 AVR Libc

1.1 Introduction

The latest version of this document is always available from <http://savannah.nongnu.org/projects/avr-libc/>

The AVR Libc package provides a subset of the standard C library for *Atmel AVR 8-bit RISC microcontrollers*. In addition, the library provides the basic

startup code needed by most applications.

There is a wealth of information in this document which goes beyond simply describing the interfaces and routines provided by the library. We hope that this document provides enough information to get a new AVR developer up to speed quickly using the freely available development tools: binutils, gcc avr-libc and many others.

If you find yourself stuck on a problem which this document doesn't quite address, you may wish to post a message to the avr-gcc mailing list. Most of the developers of the AVR binutils and gcc ports in addition to the developers of avr-libc subscribe to the list, so you will usually be able to get your problem resolved. You can subscribe to the list at <http://lists.nongnu.org/mailman/listinfo/avr-gcc-list>. Before posting to the list, you might want to try reading the [Frequently Asked Questions](#) chapter of this document.

Note

If you think you've found a bug, or have a suggestion for an improvement, either in this documentation or in the library itself, please use the bug tracker at <https://savannah.nongnu.org/bugs/?group=avr-libc> to ensure the issue won't be forgotten.

1.2 General information about this library

In general, it has been the goal to stick as best as possible to established standards while implementing this library. Commonly, this refers to the C library as described by the ANSI X3.159-1989 and ISO/IEC 9899:1990 ("ANSI-C") standard, as well as parts of their successor ISO/IEC 9899:1999 ("C99"). Some additions have been inspired by other standards like IEEE Std 1003.1-1988 ("POSIX.1"), while other extensions are purely AVR-specific (like the entire program-space string interface).

Unless otherwise noted, functions of this library are *not* guaranteed to be reentrant. In particular, any functions that store local state are known to be non-reentrant, as well as functions that manipulate IO registers like the EEPROM access routines. If these functions are used within both standard and interrupt contexts undefined behaviour will result. See the FAQ for a more detailed discussion.

1.3 Supported Devices

The following is a list of AVR devices currently supported by the library. Note that actual support for some newer devices depends on the ability of the compiler/assembler to support these devices at library compile-time.

megaAVR Devices:

- atmega103
- atmega128

- atmega128a
- atmega1280
- atmega1281
- atmega1284
- atmega1284p
- atmega16
- atmega161
- atmega162
- atmega163
- atmega164a
- atmega164p
- atmega164pa
- atmega165
- atmega165a
- atmega165p
- atmega165pa
- atmega168
- atmega168a
- atmega168p
- atmega168pa
- atmega168pb
- atmega16a
- atmega2560
- atmega2561
- atmega32
- atmega32a
- atmega323
- atmega324a
- atmega324p
- atmega324pa

- atmega325
- atmega325a
- atmega325p
- atmega325pa
- atmega3250
- atmega3250a
- atmega3250p
- atmega3250pa
- atmega328
- atmega328p
- atmega48
- atmega48a
- atmega48pa
- atmega48pb
- atmega48p
- atmega64
- atmega64a
- atmega640
- atmega644
- atmega644a
- atmega644p
- atmega644pa
- atmega645
- atmega645a
- atmega645p
- atmega6450
- atmega6450a
- atmega6450p
- atmega8
- atmega8a

- atmega88
- atmega88a
- atmega88p
- atmega88pa
- atmega88pb
- atmega8515
- atmega8535

tinyAVR Devices:

- attiny4
- attiny5
- attiny10
- attiny11 [\[1\]](#)
- attiny12 [\[1\]](#)
- attiny13
- attiny13a
- attiny15 [\[1\]](#)
- attiny20
- attiny22
- attiny24
- attiny24a
- attiny25
- attiny26
- attiny261
- attiny261a
- attiny28 [\[1\]](#)
- attiny2313
- attiny2313a
- attiny40
- attiny4313

- attiny43u
- attiny44
- attiny44a
- attiny44l
- attiny45
- attiny46l
- attiny46la
- attiny48
- attiny828
- attiny84
- attiny84a
- attiny84l
- attiny85
- attiny86l
- attiny86la
- attiny87
- attiny88
- attiny1634

Automotive AVR Devices:

- atmega16m1
- atmega32c1
- atmega32m1
- atmega64c1
- atmega64m1
- attiny167
- ata5505
- ata5272
- ata5702m322
- ata5782

- ata5790
- ata5790n
- ata5791
- ata5795
- ata5831
- ata6612c
- ata6613c
- ata6614q
- ata6616c
- ata6617c
- ata664251
- ata8210
- ata8510

CAN AVR Devices:

- at90can32
- at90can64
- at90can128

LCD AVR Devices:

- atmega169
- atmega169a
- atmega169p
- atmega169pa
- atmega329
- atmega329a
- atmega329p
- atmega329pa
- atmega3290
- atmega3290a

- atmega3290p
- atmega3290pa
- atmega649
- atmega649a
- atmega6490
- atmega6490a
- atmega6490p
- atmega649p

Lighting AVR Devices:

- at90pwm1
- at90pwm2
- at90pwm2b
- at90pwm216
- at90pwm3
- at90pwm3b
- at90pwm316
- at90pwm161
- at90pwm81

Smart Battery AVR Devices:

- atmega8hva
- atmega16hva
- atmega16hva2
- atmega16hvb
- atmega16hvbrevb
- atmega32hvb
- atmega32hvbrevb
- atmega64hve

- atmega64hve2
- atmega406

USB AVR Devices:

- at90usb82
- at90usb162
- at90usb646
- at90usb647
- at90usb1286
- at90usb1287
- atmega8u2
- atmega16u2
- atmega16u4
- atmega32u2
- atmega32u4
- atmega32u6

XMEGA Devices:

- atxmega8e5
- atxmega16a4
- atxmega16a4u
- atxmega16c4
- atxmega16d4
- atxmega16e5
- atxmega32a4
- atxmega32a4u
- atxmega32c3
- atxmega32c4
- atxmega32d3

- atxmega32d4
- atxmega32e5
- atxmega64a1
- atxmega64a1u
- atxmega64a3
- atxmega64a3u
- atxmega64a4u
- atxmega64b1
- atxmega64b3
- atxmega64c3
- atxmega64d3
- atxmega64d4
- atxmega128a1
- atxmega128a1u
- atxmega128a3
- atxmega128a3u
- atxmega128a4u
- atxmega128b1
- atxmega128b3
- atxmega128c3
- atxmega128d3
- atxmega128d4
- atxmega192a3
- atxmega192a3u
- atxmega192c3
- atxmega192d3
- atxmega256a3
- atxmega256a3u
- atxmega256a3b
- atxmega256a3bu

- atxmega256c3
- atxmega256d3
- atxmega384c3
- atxmega384d3

Wireless AVR devices:

- atmega644rfr2
- atmega64rfr2
- atmega128rfa1
- atmega1284rfr2
- atmega128rfr2
- atmega2564rfr2
- atmega256rfr2

Miscellaneous Devices:

- at94K [\[2\]](#)
- at76c711 [\[3\]](#)
- at43usb320
- at43usb355
- at86rf401
- at90scr100
- ata6285
- ata6286
- ata6289
- m3000 [\[4\]](#)

Classic AVR Devices:

- at90s1200 [\[1\]](#)
- at90s2313

- at90s2323
- at90s2333
- at90s2343
- at90s4414
- at90s4433
- at90s4434
- at90s8515
- at90c8534
- at90s8535

Note

[1] Assembly only. There is no direct support for these devices to be programmed in C since they do not have a RAM based stack. Still, it could be possible to program them in C, see the [FAQ](#) for an option.

Note

[2] The at94K devices are a combination of FPGA and AVR microcontroller. [TRoth-2002/11/12: Not sure of the level of support for these. More information would be welcomed.]

Note

[3] The at76c711 is a USB to fast serial interface bridge chip using an AVR core.

Note

[4] The m3000 is a motor controller AVR ASIC from Intelligent Motion Systems (IMS) / Schneider Electric.

1.4 avr-libc License

avr-libc can be freely used and redistributed, provided the following license conditions are met.

```
Portions of avr-libc are Copyright (c) 1999-2016
Werner Boellmann,
Dean Camera,
Pieter Conradie,
Brian Dean,
Keith Gudger,
Wouter van Gulik,
Bjoern Haase,
```


Steinar Haugen,
Peter Jansen,
Reinhard Jessich,
Magnus Johansson,
Harald Kipp,
Carlos Lamas,
Cliff Lawson,
Artur Lipowski,
Marek Michalkiewicz,
Todd C. Miller,
Rich Neswold,
Colin O'Flynn,
Bob Paddock,
Andrey Pashchenko,
Reiner Patommel,
Florin-Viorel Petrov,
Alexander Popov,
Michael Rickman,
Theodore A. Roth,
Juergen Schilling,
Philip Soeberg,
Anatoly Sokolov,
Nils Kristian Strom,
Michael Stumpf,
Stefan Swanepoel,
Helmut Wallner,
Eric B. Weddington,
Joerg Wunsch,
Dmitry Xmelkov,
Atmel Corporation,
egnite Software GmbH,
The Regents of the University of California.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the copyright holders nor the names of contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

2 Toolchain Overview

2.1 Introduction

Welcome to the open source software development toolset for the Atmel AVR!

There is not a single tool that provides everything needed to develop software for the AVR. It takes many tools working together. Collectively, the group of tools are called a toolset, or commonly a toolchain, as the tools are chained together to produce the final executable application for the AVR microcontroller.

The following sections provide an overview of all of these tools. You may be used to cross-compilers that provide everything with a GUI front-end, and not know what goes on "underneath the hood". You may be coming from a desktop or server computer background and not used to embedded systems. Or you may be just learning about the most common software development toolchain available on Unix and Linux systems. Hopefully the following overview will be helpful in putting everything in perspective.

2.2 FSF and GNU

According to its website, "the Free Software Foundation (FSF), established in 1985, is dedicated to promoting computer users' rights to use, study, copy, modify, and redistribute computer programs. The FSF promotes the development and use of free software, particularly the GNU operating system, used widely in its GNU/Linux variant." The FSF remains the primary sponsor of the GNU project.

The GNU Project was launched in 1984 to develop a complete Unix-like operating system which is free software: the GNU system. GNU is a recursive acronym for »GNU's Not Unix«; it is pronounced guh-noo, approximately like canoe.

One of the main projects of the GNU system is the GNU Compiler Collection, or GCC, and its sister project, GNU Binutils. These two open source projects provide a foundation for a software development toolchain. Note that these projects were designed to originally run on Unix-like systems.

2.3 GCC

GCC stands for GNU Compiler Collection. GCC is highly flexible compiler system. It has different compiler front-ends for different languages. It has many back-ends that generate assembly code for many different processors and host operating systems. All share a common "middle-end", containing the generic parts of the compiler, including a lot of optimizations.

In GCC, a *host* system is the system (processor/OS) that the compiler runs on. A *target* system is the system that the compiler compiles code for. And, a *build* system is the system that the compiler is built (from source code) on. If a compiler has the same system for *host* and for *target*, it is known as a *native* compiler. If a compiler has different systems for *host* and *target*, it is known as a cross-compiler. (And if all three, *build*, *host*, and *target* systems are different, it is known as a Canadian cross compiler, but we won't discuss that here.) When GCC is built to execute on a *host*

system such as FreeBSD, Linux, or Windows, and it is built to generate code for the AVR microcontroller *target*, then it is a cross compiler, and this version of GCC is commonly known as "AVR GCC". In documentation, or discussion, AVR GCC is used when referring to GCC targeting specifically the AVR, or something that is AVR specific about GCC. The term "GCC" is usually used to refer to something generic about GCC, or about GCC as a whole.

GCC is different from most other compilers. GCC focuses on translating a high-level language to the target assembly only. AVR GCC has three available compilers for the AVR: C language, C++, and Ada. The compiler itself does not assemble or link the final code.

GCC is also known as a "driver" program, in that it knows about, and drives other programs seamlessly to create the final output. The assembler, and the linker are part of another open source project called GNU Binutils. GCC knows how to drive the GNU assembler (gas) to assemble the output of the compiler. GCC knows how to drive the GNU linker (ld) to link all of the object modules into a final executable.

The two projects, GCC and Binutils, are very much interrelated and many of the same volunteers work on both open source projects.

When GCC is built for the AVR target, the actual program names are prefixed with "avr-". So the actual executable name for AVR GCC is: avr-gcc. The name "avr-gcc" is used in documentation and discussion when referring to the program itself and not just the whole AVR GCC system.

See the GCC Web Site and GCC User Manual for more information about GCC.

2.4 GNU Binutils

The name GNU Binutils stands for "Binary Utilities". It contains the GNU assembler (gas), and the GNU linker (ld), but also contains many other utilities that work with binary files that are created as part of the software development toolchain.

Again, when these tools are built for the AVR target, the actual program names are prefixed with "avr-". For example, the assembler program name, for a native assembler is "as" (even though in documentation the GNU assembler is commonly referred to as "gas"). But when built for an AVR target, it becomes "avr-as". Below is a list of the programs that are included in Binutils:

avr-as

The Assembler.

avr-ld

The Linker.

avr-ar

Create, modify, and extract from libraries (archives).

avr-ranlib

Generate index to library (archive) contents.

avr-objcopy

Copy and translate object files to different formats.

avr-objdump

Display information from object files including disassembly.

avr-size

List section sizes and total size.

avr-nm

List symbols from object files.

avr-strings

List printable strings from files.

avr-strip

Discard symbols from files.

avr-readelf

Display the contents of ELF format files.

avr-addr2line

Convert addresses to file and line.

avr-c++filt

Filter to demangle encoded C++ symbols.

2.5 avr-libc

GCC and Binutils provides a lot of the tools to develop software, but there is one critical component that they do not provide: a Standard C Library.

There are different open source projects that provide a Standard C Library depending upon your system time, whether for a native compiler (GNU Libc), for some other embedded system (newlib), or for some versions of Linux (uCLibc). The open source AVR toolchain has its own Standard C Library project: avr-libc.

AVR-Libc provides many of the same functions found in a regular Standard C Library and many additional library functions that is specific to an AVR. Some of the Standard C Library functions that are commonly used on a PC environment have limitations or additional issues that a user needs to be aware of when used on an embedded system.

AVR-Libc also contains the most documentation about the whole AVR toolchain.

2.6 Building Software

Even though GCC, Binutils, and avr-libc are the core projects that are used to build software for the AVR, there is another piece of software that ties it all together: Make. GNU Make is a program that makes things, and mainly software. Make interprets and executes a Makefile that is written for a project. A Makefile contains dependency rules, showing which output files are dependent upon which input files, and instructions on how to build output files from input files.

Some distributions of the toolchains, and other AVR tools such as MFile, contain a Makefile template written for the AVR toolchain and AVR applications that you can copy and modify for your application.

See the GNU Make User Manual for more information.

2.7 AVRDUDE

After creating your software, you'll want to program your device. You can do this by using the program AVRDUDE which can interface with various hardware devices to program your processor.

AVRDUDE is a very flexible package. All the information about AVR processors and various hardware programmers is stored in a text database. This database can be modified by any user to add new hardware or to add an AVR processor if it is not already listed.

2.8 GDB / Insight / DDD

The GNU Debugger (GDB) is a command-line debugger that can be used with the rest of the AVR toolchain. Insight is GDB plus a GUI written in Tcl/Tk. Both GDB and Insight are configured for the AVR and the main executables are prefixed with the target name: avr-gdb, and avr-insight. There is also a "text mode" GUI for GDB: avr-gdbtui. DDD (Data Display Debugger) is another popular GUI front end to GDB, available on Unix and Linux systems.

2.9 AVaRICE

AVaRICE is a back-end program to AVR GDB and interfaces to the Atmel JTAG In-Circuit Emulator (ICE), to provide emulation capabilities.

2.10 SimulAVR

SimulAVR is an AVR simulator used as a back-end with AVR GDB.

2.11 Utilities

There are also other optional utilities available that may be useful to add to your toolset.

`SRecord` is a collection of powerful tools for manipulating EPROM load files. It reads and writes numerous EPROM file formats, and can perform many different manipulations.

`MFile` is a simple Makefile generator is meant as an aid to quickly customize a Makefile to use for your AVR application.

2.12 Toolchain Distributions (Distros)

All of the various open source projects that comprise the entire toolchain are normally distributed as source code. It is left up to the user to build the tool application from its source code. This can be a very daunting task to any potential user of these tools.

Luckily there are people who help out in this area. Volunteers take the time to build the application from source code on particular host platforms and sometimes packaging the tools for convenient installation by the end user. These packages contain the binary executables of the tools, pre-made and ready to use. These packages are known as "distributions" of the AVR toolchain, or by a more shortened name, "distros".

AVR toolchain distros are available on FreeBSD, Windows, Mac OS X, and certain flavors of Linux.

2.13 Open Source

All of these tools, from the original source code in the multitude of projects, to the various distros, are put together by many, many volunteers. All of these projects could always use more help from other people who are willing to volunteer some of their time. There are many different ways to help, for people with varying skill levels, abilities, and available time.

You can help to answer questions in mailing lists such as the `avr-gcc-list`, or on forums at the AVR Freaks website. This helps many people new to the open source AVR tools.

If you think you found a bug in any of the tools, it is always a big help to submit a good bug report to the proper project. A good bug report always helps other volunteers to analyze the problem and to get it fixed for future versions of the software.

You can also help to fix bugs in various software projects, or to add desirable new features.

Volunteers are always welcome! :-)

3 Memory Areas and Using `malloc()`

3.1 Introduction

Many of the devices that are possible targets of `avr-libc` have a minimal amount of RAM. The smallest parts supported by the C environment come with 128 bytes of RAM. This needs to be shared between initialized and uninitialized variables ([sections](#)

.data and .bss), the dynamic memory allocator, and the stack that is used for calling subroutines and storing local (automatic) variables.

Also, unlike larger architectures, there is no hardware-supported memory management which could help in separating the mentioned RAM regions from being overwritten by each other.

The standard RAM layout is to place .data variables first, from the beginning of the internal RAM, followed by .bss. The stack is started from the top of internal RAM, growing downwards. The so-called "heap" available for the dynamic memory allocator will be placed beyond the end of .bss. Thus, there's no risk that dynamic memory will ever collide with the RAM variables (unless there were bugs in the implementation of the allocator). There is still a risk that the heap and stack could collide if there are large requirements for either dynamic memory or stack space. The former can even happen if the allocations aren't all that large but dynamic memory allocations get fragmented over time such that new requests don't quite fit into the "holes" of previously freed regions. Large stack space requirements can arise in a C function containing large and/or numerous local variables or when recursively calling function.

Note

The pictures shown in this document represent typical situations where the RAM locations refer to an ATmega128. The memory addresses used are not displayed in a linear scale.

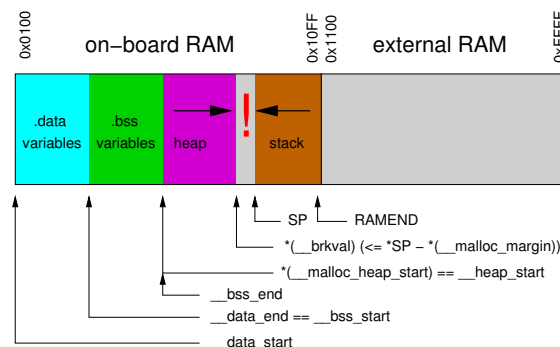


Figure 1: RAM map of a device with internal RAM

On a simple device like a microcontroller it is a challenge to implement a dynamic memory allocator that is simple enough so the code size requirements will remain low, yet powerful enough to avoid unnecessary memory fragmentation and to get it all done with reasonably few CPU cycles. Microcontrollers are often low on space and also run at much lower speeds than the typical PC these days.

The memory allocator implemented in avr-libc tries to cope with all of these constraints, and offers some tuning options that can be used if there are more resources available than in the default configuration.

3.2 Internal vs. external RAM

Obviously, the constraints are much harder to satisfy in the default configuration where only internal RAM is available. Extreme care must be taken to avoid a stack-heap collision, both by making sure functions aren't nesting too deeply, and don't require too much stack space for local variables, as well as by being cautious with allocating too much dynamic memory.

If external RAM is available, it is strongly recommended to move the heap into the external RAM, regardless of whether or not the variables from the `.data` and `.bss` sections are also going to be located there. The stack should always be kept in internal RAM. Some devices even require this, and in general, internal RAM can be accessed faster since no extra wait states are required. When using dynamic memory allocation and stack and heap are separated in distinct memory areas, this is the safest way to avoid a stack-heap collision.

3.3 Tunables for `malloc()`

There are a number of variables that can be tuned to adapt the behavior of `malloc()` to the expected requirements and constraints of the application. Any changes to these tunables should be made before the very first call to `malloc()`. Note that some library functions might also use dynamic memory (notably those from the `<stdio.h>`: [Standard IO facilities](#)), so make sure the changes will be done early enough in the startup sequence.

The variables `__malloc_heap_start` and `__malloc_heap_end` can be used to restrict the `malloc()` function to a certain memory region. These variables are statically initialized to point to `__heap_start` and `__heap_end`, respectively, where `__heap_start` is filled in by the linker to point just beyond `.bss`, and `__heap_end` is set to 0 which makes `malloc()` assume the heap is below the stack.

If the heap is going to be moved to external RAM, `__malloc_heap_end` *must* be adjusted accordingly. This can either be done at run-time, by writing directly to this variable, or it can be done automatically at link-time, by adjusting the value of the symbol `__heap_end`.

The following example shows a linker command to relocate the entire `.data` and `.bss` segments, and the heap to location 0x1100 in external RAM. The heap will extend up to address 0xffff.

```
avr-gcc ... -Wl,--section-start,.data=0x801100,--defsym=__heap_end=0x80ffff ...
```

Note

See [explanation](#) for offset 0x800000. See the chapter about [using gcc](#) for the `-Wl` options.

The `ld` (linker) user manual states that using `-Tdata=<x>` is equivalent to using `--section-start,.data=<x>`. However, you have to use `--section-start` as above because the GCC frontend also sets the `-Tdata` option for all MCU types where the SRAM doesn't start at 0x800060. Thus, the linker is being faced with two `-Tdata` options. Starting with binutils 2.16, the linker changed the preference, and picks the "wrong" option in this situation.




Figure 2: Internal RAM: stack only, external RAM: variables and heap

If dynamic memory should be placed in external RAM, while keeping the variables in internal RAM, something like the following could be used. Note that for demonstration purposes, the assignment of the various regions has not been made adjacent in this example, so there are "holes" below and above the heap in external RAM that remain completely inaccessible by regular variables or dynamic memory allocations (shown in light bisque color in the picture below).

```
avr-gcc ... -Wl,--defsym=__heap_start=0x802000,--defsym=__heap_end=0x803fff ...
```




Figure 3: Internal RAM: variables and stack, external RAM: heap

If `__malloc_heap_end` is 0, the allocator attempts to detect the bottom of stack in order to prevent a stack-heap collision when extending the actual size of the heap to gain more space for dynamic memory. It will not try to go beyond the current stack limit, decreased by `__malloc_margin` bytes. Thus, all possible stack frames of interrupt routines that could interrupt the current function, plus all further nested function calls must not require more stack space, or they will risk colliding with the data segment.

The default value of `__malloc_margin` is set to 32.

3.4 Implementation details

Dynamic memory allocation requests will be returned with a two-byte header prepended that records the size of the allocation. This is later used by `free()`. The returned address points just beyond that header. Thus, if the application accidentally writes before the returned memory region, the internal consistency of the memory allocator is compromised.

The implementation maintains a simple freelist that accounts for memory blocks that have been returned in previous calls to `free()`. Note that all of this memory is considered to be successfully added to the heap already, so no further checks against stack-heap collisions are done when recycling memory from the freelist.

The freelist itself is not maintained as a separate data structure, but rather by modifying the contents of the freed memory to contain pointers chaining the pieces together. That way, no additional memory is required to maintain this list except for a variable that keeps track of the lowest memory segment available for reallocation. Since both, a chain pointer and the size of the chunk need to be recorded in each chunk, the minimum chunk size on the freelist is four bytes.

When allocating memory, first the freelist is walked to see if it could satisfy the request. If there's a chunk available on the freelist that will fit the request exactly, it will be taken, disconnected from the freelist, and returned to the caller. If no exact match could be found, the closest match that would just satisfy the request will be used. The chunk will normally be split up into one to be returned to the caller, and another (smaller) one that will remain on the freelist. In case this chunk was only up to two bytes larger than the request, the request will simply be altered internally to also account for these additional bytes since no separate freelist entry could be split off in that case.

If nothing could be found on the freelist, heap extension is attempted. This is where `__malloc_margin` will be considered if the heap is operating below the stack, or where `__malloc_heap_end` will be verified otherwise.

If the remaining memory is insufficient to satisfy the request, `NULL` will eventually be returned to the caller.

When calling `free()`, a new freelist entry will be prepared. An attempt is then made to aggregate the new entry with possible adjacent entries, yielding a single larger entry available for further allocations. That way, the potential for heap fragmentation is hopefully reduced. When deallocating the topmost chunk of memory, the size of the heap is reduced.

A call to `realloc()` first determines whether the operation is about to grow or shrink the current allocation. When shrinking, the case is easy: the existing chunk is split, and the tail of the region that is no longer to be used is passed to the standard `free()` function for insertion into the freelist. Checks are first made whether the tail chunk is large enough to hold a chunk of its own at all, otherwise `realloc()` will simply do nothing, and return the original region.

When growing the region, it is first checked whether the existing allocation can be extended in-place. If so, this is done, and the original pointer is returned without copying any data contents. As a side-effect, this check will also record the size of the largest chunk on the freelist.

If the region cannot be extended in-place, but the old chunk is at the top of heap, and the above freelist walk did not reveal a large enough chunk on the freelist to satisfy the new request, an attempt is made to quickly extend this topmost chunk (and thus the heap), so no need arises to copy over the existing data. If there's no more space available in the heap (same check is done as in [malloc\(\)](#)), the entire request will fail.

Otherwise, [malloc\(\)](#) will be called with the new request size, the existing data will be copied over, and [free\(\)](#) will be called on the old region.

4 Memory Sections

Remarks

Need to list all the sections which are available to the avr.

Weak Bindings

FIXME: need to discuss the .weak directive.

The following describes the various sections available.

4.1 The .text Section

The .text section contains the actual machine instructions which make up your program. This section is further subdivided by the .initN and .finiN sections discussed below.

Note

The `avr-size` program (part of `binutils`), coming from a Unix background, doesn't account for the .data initialization space added to the .text section, so in order to know how much flash the final program will consume, one needs to add the values for both, .text and .data (but not .bss), while the amount of pre-allocated SRAM is the sum of .data and .bss.

4.2 The .data Section

This section contains static data which was defined in your code. Things like the following would end up in .data:

```
char err_str[] = "Your program has died a horrible death!";

struct point pt = { 1, 1 };
```

It is possible to tell the linker the SRAM address of the beginning of the .data section. This is accomplished by adding `-Wl,-Tdata,addr` to the `avr-gcc` command used to link your program. Note that `addr` must be offset by adding 0x800000 to the real SRAM address so that the linker knows that the address is in the SRAM memory space. Thus, if you want the .data section to start at 0x1100, pass 0x801100 at the address to the linker. [offset [explained](#)]

Note

When using `malloc()` in the application (which could even happen inside library calls), [additional adjustments](#) are required.

4.3 The .bss Section

Uninitialized global or static variables end up in the .bss section.

4.4 The .eeprom Section

This is where eeprom variables are stored.

4.5 The .noinit Section

This section is a part of the .bss section. What makes the .noinit section special is that variables which are defined as such:

```
int foo __attribute__((section(".noinit")));
```

will not be initialized to zero during startup as would normal .bss data.

Only uninitialized variables can be placed in the .noinit section. Thus, the following code will cause `avr-gcc` to issue an error:

```
int bar __attribute__((section(".noinit"))) = 0xaa;
```

It is possible to tell the linker explicitly where to place the .noinit section by adding `-Wl,--section-start=.noinit=0x802000` to the `avr-gcc` command line at the linking stage. For example, suppose you wish to place the .noinit section at SRAM address 0x2000:

```
$ avr-gcc ... -Wl,--section-start=.noinit=0x802000 ...
```

Note

Because of the Harvard architecture of the AVR devices, you must manually add 0x800000 to the address you pass to the linker as the start of the section. Otherwise, the linker thinks you want to put the .noinit section into the .text section instead of .data/.bss and will complain.

Alternatively, you can write your own linker script to automate this. [FIXME: need an example or ref to dox for writing linker scripts.]

4.6 The .initN Sections

These sections are used to define the startup code from reset up through the start of `main()`. These all are subparts of the [.text section](#).

The purpose of these sections is to allow for more specific placement of code within your program.

Note

Sometimes, it is convenient to think of the `.initN` and `.finiN` sections as functions, but in reality they are just symbolic names which tell the linker where to stick a chunk of code which is *not* a function. Notice that the examples for [asm](#) and [C](#) can not be called as functions and should not be jumped into.

The `.initN` sections are executed in order from 0 to 9.

.init0:

Weakly bound to `__init()`. If user defines `__init()`, it will be jumped into immediately after a reset.

.init1:

Unused. User definable.

.init2:

In C programs, weakly bound to initialize the stack, and to clear `__zero_reg__` (r1).

.init3:

Unused. User definable.

.init4:

For devices with > 64 KB of ROM, `.init4` defines the code which takes care of copying the contents of `.data` from the flash to SRAM. For all other devices, this code as well as the code to zero out the `.bss` section is loaded from `libgcc.a`.

.init5:

Unused. User definable.

.init6:

Unused for C programs, but used for constructors in C++ programs.

.init7:

Unused. User definable.

.init8:

Unused. User definable.

.init9:

Jumps into main().

4.7 The .finiN Sections

These sections are used to define the exit code executed after return from main() or a call to [exit\(\)](#). These all are subparts of the [.text section](#).

The **.finiN** sections are executed in descending order from 9 to 0.

.fini9:

Unused. User definable. This is effectively where `_exit()` starts.

.fini8:

Unused. User definable.

.fini7:

Unused. User definable.

.fini6:

Unused for C programs, but used for destructors in C++ programs.

.fini5:

Unused. User definable.

.fini4:

Unused. User definable.

.fini3:

Unused. User definable.

.fini2:

Unused. User definable.

.fini1:

Unused. User definable.

.fini0:

Goes into an infinite loop after program termination and completion of any `_exit()` code (execution of code in the `.fini9` -> `.fini1` sections).

4.8 The .note.gnu.avr.deviceinfo Section

This section contains device specific information picked up from the device header file and compiler builtin macros. The layout conforms to the standard ELF note section layout (http://docs.oracle.com/cd/E23824_01/html/819-0690/chapter6-18048.html).

The section contents are laid out as below.

```
#define __NOTE_NAME_LEN 4
struct __note_gnu_avr_deviceinfo
{
    struct
    {
        uint32_t namesz;          /* = __NOTE_NAME_LEN */
        uint32_t descsize;        /* = size of avr_desc */
        uint32_t type;            /* = 1 - no other AVR note types exist */
        char note_name[__NOTE_NAME_LEN]; /* = "AVR\0" */
    }
    note_header;
    struct
    {
        uint32_t flash_start;
        uint32_t flash_size;
        uint32_t sram_start;
        uint32_t sram_size;
        uint32_t eeprom_start;
        uint32_t eeprom_size;
        uint32_t offset_table_size;
        uint32_t offset_table[1]; /* Offset table containing byte offsets into
                                   string table that immediately follows it.
                                   index 0: Device name byte offset
                                   */
        char str_table [2 +
            strlen(__AVR_DEVICE_NAME__)] /* Standard ELF string table.
                                           index 0 : NULL
                                           index 1 : Device name
                                           index 2 : NULL
                                           */
    }
    avr_desc;
};
```

4.9 Using Sections in Assembler Code

Example:

```
#include <avr/io.h>

.section .init1,"ax",@progbits
ldi    r0, 0xff
out    _SFR_IO_ADDR(PORTB), r0
out    _SFR_IO_ADDR(DDRB), r0
```

Note

The **, "ax", @progbits** tells the assembler that the section is allocatable ("a"), executable ("x") and contains data ("@progbits"). For more detailed information

on the `.section` directive, see the gas user manual.

4.10 Using Sections in C Code

Example:

```
#include <avr/io.h>

void my_init_portb (void) __attribute__ ((naked)) \
    __attribute__ ((section (".init3")))
    __attribute__ ((used));

void
my_init_portb (void)
{
    PORTB = 0xff;
    DDRB = 0xff;
}
```

Note

Section `.init3` is used in this example, as this ensures the internal `__zero_reg__` has already been set up. The code generated by the compiler might blindly rely on `__zero_reg__` being really 0. `__attribute__ ((used))` tells the compiler that code must be generated for this function even if it appears that the function is not referenced - this is necessary to prevent compiler optimizations (like LTO) from eliminating the function.

5 Data in Program Space

5.1 Introduction

So you have some constant data and you're running out of room to store it? Many AVRs have limited amount of RAM in which to store data, but may have more Flash space available. The AVR is a Harvard architecture processor, where Flash is used for the program, RAM is used for data, and they each have separate address spaces. It is a challenge to get constant data to be stored in the Program Space, and to retrieve that data to use it in the AVR application.

The problem is exacerbated by the fact that the C Language was not designed for Harvard architectures, it was designed for Von Neumann architectures where code and data exist in the same address space. This means that any compiler for a Harvard architecture processor, like the AVR, has to use other means to operate with separate address spaces.

Some compilers use non-standard C language keywords, or they extend the standard syntax in ways that are non-standard. The AVR toolset takes a different approach.

GCC has a special keyword, `__attribute__` that is used to attach different attributes to things such as function declarations, variables, and types. This keyword is followed by an attribute specification in double parentheses. In AVR GCC, there is a

special attribute called `progmem`. This attribute is use on data declarations, and tells the compiler to place the data in the Program Memory (Flash).

AVR-Libc provides a simple macro `PROGMEM` that is defined as the attribute syntax of GCC with the `progmem` attribute. This macro was created as a convenience to the end user, as we will see below. The `PROGMEM` macro is defined in the `<avr/pgmspace.h>` system header file.

It is difficult to modify GCC to create new extensions to the C language syntax, so instead, `avr-libc` has created macros to retrieve the data from the Program Space. These macros are also found in the `<avr/pgmspace.h>` system header file.

5.2 A Note On const

Many users bring up the idea of using C's keyword `const` as a means of declaring data to be in Program Space. Doing this would be an abuse of the intended meaning of the `const` keyword.

`const` is used to tell the compiler that the data is to be "read-only". It is used to help make it easier for the compiler to make certain transformations, or to help the compiler check for incorrect usage of those variables.

For example, the `const` keyword is commonly used in many functions as a modifier on the parameter type. This tells the compiler that the function will only use the parameter as read-only and will not modify the contents of the parameter variable.

`const` was intended for uses such as this, not as a means to identify where the data should be stored. If it were used as a means to define data storage, then it loses its correct meaning (changes its semantics) in other situations such as in the function parameter example.

5.3 Storing and Retrieving Data in the Program Space

Let's say you have some global data:

```
unsigned char mydata[11][10] =
{
    {0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07, 0x08, 0x09},
    {0x0A, 0x0B, 0x0C, 0x0D, 0x0E, 0x0F, 0x10, 0x11, 0x12, 0x13},
    {0x14, 0x15, 0x16, 0x17, 0x18, 0x19, 0x1A, 0x1B, 0x1C, 0x1D},
    {0x1E, 0x1F, 0x20, 0x21, 0x22, 0x23, 0x24, 0x25, 0x26, 0x27},
    {0x28, 0x29, 0x2A, 0x2B, 0x2C, 0x2D, 0x2E, 0x2F, 0x30, 0x31},
    {0x32, 0x33, 0x34, 0x35, 0x36, 0x37, 0x38, 0x39, 0x3A, 0x3B},
    {0x3C, 0x3D, 0x3E, 0x3F, 0x40, 0x41, 0x42, 0x43, 0x44, 0x45},
    {0x46, 0x47, 0x48, 0x49, 0x4A, 0x4B, 0x4C, 0x4D, 0x4E, 0x4F},
    {0x50, 0x51, 0x52, 0x53, 0x54, 0x55, 0x56, 0x57, 0x58, 0x59},
    {0x5A, 0x5B, 0x5C, 0x5D, 0x5E, 0x5F, 0x60, 0x61, 0x62, 0x63},
    {0x64, 0x65, 0x66, 0x67, 0x68, 0x69, 0x6A, 0x6B, 0x6C, 0x6D}
};
```

and later in your code you access this data in a function and store a single byte into a variable like so:

```
byte = mydata[i][j];
```

Now you want to store your data in Program Memory. Use the `PROGMEM` macro found in `<avr/pgmspace.h>` and put it after the declaration of the variable, but before the initializer, like so:

```
#include <avr/pgmspace.h>
.
.
.
const unsigned char mydata[11][10] PROGMEM =
{
    {0x00,0x01,0x02,0x03,0x04,0x05,0x06,0x07,0x08,0x09},
    {0x0A,0x0B,0x0C,0x0D,0x0E,0x0F,0x10,0x11,0x12,0x13},
    {0x14,0x15,0x16,0x17,0x18,0x19,0x1A,0x1B,0x1C,0x1D},
    {0x1E,0x1F,0x20,0x21,0x22,0x23,0x24,0x25,0x26,0x27},
    {0x28,0x29,0x2A,0x2B,0x2C,0x2D,0x2E,0x2F,0x30,0x31},
    {0x32,0x33,0x34,0x35,0x36,0x37,0x38,0x39,0x3A,0x3B},
    {0x3C,0x3D,0x3E,0x3F,0x40,0x41,0x42,0x43,0x44,0x45},
    {0x46,0x47,0x48,0x49,0x4A,0x4B,0x4C,0x4D,0x4E,0x4F},
    {0x50,0x51,0x52,0x53,0x54,0x55,0x56,0x57,0x58,0x59},
    {0x5A,0x5B,0x5C,0x5D,0x5E,0x5F,0x60,0x61,0x62,0x63},
    {0x64,0x65,0x66,0x67,0x68,0x69,0x6A,0x6B,0x6C,0x6D}
};
```

That's it! Now your data is in the Program Space. You can compile, link, and check the map file to verify that `mydata` is placed in the correct section.

Now that your data resides in the Program Space, your code to access (read) the data will no longer work. The code that gets generated will retrieve the data that is located at the address of the `mydata` array, plus offsets indexed by the `i` and `j` variables. However, the final address that is calculated where to retrieve the data points to the Data Space! Not the Program Space where the data is actually located. It is likely that you will be retrieving some garbage. The problem is that AVR GCC does not intrinsically know that the data resides in the Program Space.

The solution is fairly simple. The "rule of thumb" for accessing data stored in the Program Space is to access the data as you normally would (as if the variable is stored in Data Space), like so:

```
byte = mydata[i][j];
```

then take the address of the data:

```
byte = &(mydata[i][j]);
```

then use the appropriate `pgm_read_*` macro, and the address of your data becomes the parameter to that macro:

```
byte = pgm_read_byte(&(mydata[i][j]));
```

The `pgm_read_*` macros take an address that points to the Program Space, and retrieves the data that is stored at that address. This is why you take the address of the offset into the array. This address becomes the parameter to the macro so it can generate the correct code to retrieve the data from the Program Space. There are different `pgm_read_*` macros to read different sizes of data at the address given.

5.4 Storing and Retrieving Strings in the Program Space

Now that you can successfully store and retrieve simple data from Program Space you want to store and retrieve strings from Program Space. And specifically you want to store an array of strings to Program Space. So you start off with your array, like so:

```
char *string_table[] =
{
    "String 1",
    "String 2",
    "String 3",
    "String 4",
    "String 5"
};
```

and then you add your `PROGMEM` macro to the end of the declaration:

```
char *string_table[] PROGMEM =
{
    "String 1",
    "String 2",
    "String 3",
    "String 4",
    "String 5"
};
```

Right? WRONG!

Unfortunately, with GCC attributes, they affect only the declaration that they are attached to. So in this case, we successfully put the `string_table` variable, the array itself, in the Program Space. This DOES NOT put the actual strings themselves into Program Space. At this point, the strings are still in the Data Space, which is probably not what you want.

In order to put the strings in Program Space, you have to have explicit declarations for each string, and put each string in Program Space:

```
const char string_1[] PROGMEM = "String 1";
const char string_2[] PROGMEM = "String 2";
const char string_3[] PROGMEM = "String 3";
const char string_4[] PROGMEM = "String 4";
const char string_5[] PROGMEM = "String 5";
```

Then use the new symbols in your table, like so:

```
PGM_P const string_table[] PROGMEM =
{
    string_1,
    string_2,
    string_3,
    string_4,
    string_5
};
```

Now this has the effect of putting `string_table` in Program Space, where `string_table` is an array of pointers to characters (strings), where each pointer is a pointer to the Program Space, where each string is also stored.

The `PGM_P` type above is also a macro that defined as a pointer to a character in the Program Space.

Retrieving the strings are a different matter. You probably don't want to pull the string out of Program Space, byte by byte, using the `pgm_read_byte()` macro. There are other functions declared in the `<avr/pgmspace.h>` header file that work with strings that are stored in the Program Space.

For example if you want to copy the string from Program Space to a buffer in RAM (like an automatic variable inside a function, that is allocated on the stack), you can do this:

```
void foo(void)
{
    char buffer[10];

    for (unsigned char i = 0; i < 5; i++)
    {
        strcpy_P(buffer, (PGM_P)pgm_read_word(&(string_table[i])));

        // Display buffer on LCD.
    }
    return;
}
```

Here, the `string_table` array is stored in Program Space, so we access it normally, as if were stored in Data Space, then take the address of the location we want to access, and use the address as a parameter to `pgm_read_word`. We use the `pgm_read_word` macro to read the string pointer out of the `string_table` array. Remember that a pointer is 16-bits, or word size. The `pgm_read_word` macro will return a 16-bit unsigned integer. We then have to typecast it as a true pointer to program memory, `PGM_P`. This pointer is an address in Program Space pointing to the string that we want to copy. This pointer is then used as a parameter to the function `strcpy_P`. The function `strcpy_P` is just like the regular `strcpy` function, except that it copies a string from Program Space (the second parameter) to a buffer in the Data Space (the first parameter).

There are many string functions available that work with strings located in Program Space. All of these special string functions have a suffix of `_P` in the function name, and are declared in the `<avr/pgmspace.h>` header file.

5.5 Caveats

The macros and functions used to retrieve data from the Program Space have to generate some extra code in order to actually load the data from the Program Space. This incurs some extra overhead in terms of code space (extra opcodes) and execution time. Usually, both the space and time overhead is minimal compared to the space savings of putting data in Program Space. But you should be aware of this so you can minimize the number of calls within a single function that gets the same piece of data from Program Space. It is always instructive to look at the resulting disassembly from the compiler.

6 avr-libc and assembler programs

6.1 Introduction

There might be several reasons to write code for AVR microcontrollers using plain assembler source code. Among them are:

- Code for devices that do not have RAM and are thus not supported by the C compiler.
- Code for very time-critical applications.
- Special tweaks that cannot be done in C.

Usually, all but the first could probably be done easily using the [inline assembler](#) facility of the compiler.

Although avr-libc is primarily targeted to support programming AVR microcontrollers using the C (and C++) language, there's limited support for direct assembler usage as well. The benefits of it are:

- Use of the C preprocessor and thus the ability to use the same symbolic constants that are available to C programs, as well as a flexible macro concept that can use any valid C identifier as a macro (whereas the assembler's macro concept is basically targeted to use a macro in place of an assembler instruction).
- Use of the runtime framework like automatically assigning interrupt vectors. For devices that have RAM, [initializing the RAM variables](#) can also be utilized.

6.2 Invoking the compiler

For the purpose described in this document, the assembler and linker are usually not invoked manually, but rather using the C compiler frontend (`avr-gcc`) that in turn will call the assembler and linker as required.

This approach has the following advantages:

- There is basically only one program to be called directly, `avr-gcc`, regardless of the actual source language used.
- The invocation of the C preprocessor will be automatic, and will include the appropriate options to locate required include files in the filesystem.
- The invocation of the linker will be automatic, and will include the appropriate options to locate additional libraries as well as the application start-up code (`crtXXX.o`) and linker script.

Note that the invocation of the C preprocessor will be automatic when the filename provided for the assembler file ends in `.S` (the capital letter "s"). This would even apply to operating systems that use case-insensitive filesystems since the actual decision is

made based on the case of the filename suffix given on the command-line, not based on the actual filename from the file system.

As an alternative to using `.S`, the suffix `.sx` is recognized for this purpose (starting with GCC 4.3.0). This is primarily meant to be compatible with other compiler environments that have been providing this variant before in order to cope with operating systems where filenames are case-insensitive (and, with some versions of `make` that could not distinguish between `.s` and `.S` on such systems).

Alternatively, the language can explicitly be specified using the `-x assembler-with-cpp` option.

6.3 Example program

The following annotated example features a simple 100 kHz square wave generator using an AT90S1200 clocked with a 10.7 MHz crystal. Pin PD6 will be used for the square wave output.

```
#include <avr/io.h>          ; Note [1]

work    =    16              ; Note [2]
tmp     =    17

inttmp  =    19

intsav  =    0

SQUARE  =    PD6            ; Note [3]

                        ; Note [4]:
tmconst= 10700000 / 200000 ; 100 kHz => 200000 edges/s
fuzz=    8                  ; # clocks in ISR until TCNT0 is set

.section .text

.global main              ; Note [5]
main:
    rcall    ioinit
1:
    rjmp     1b            ; Note [6]

.global TIMER0_OVF_vect   ; Note [7]
TIMER0_OVF_vect:
    ldi inttmp, 256 - tmconst + fuzz
    out _SFR_IO_ADDR(TCNT0), inttmp ; Note [8]

    in  intsav, _SFR_IO_ADDR(SREG) ; Note [9]

    sbic     _SFR_IO_ADDR(PORTD), SQUARE
    rjmp     1f
    sbi _SFR_IO_ADDR(PORTD), SQUARE
    rjmp     2f
1: cbi _SFR_IO_ADDR(PORTD), SQUARE
2:

    out _SFR_IO_ADDR(SREG), intsav
    reti
```

```
ioinit:
    sbi _SFR_IO_ADDR(DDRD), SQUARE

    ldi work, _BV(TOIE0)
    out _SFR_IO_ADDR(TIMSK), work

    ldi work, _BV(CS00)      ; tmr0: CK/1
    out _SFR_IO_ADDR(TCCR0), work

    ldi work, 256 - tmconst
    out _SFR_IO_ADDR(TCNT0), work

    sei

    ret

    .global __vector_default      ; Note [10]
__vector_default:
    reti

    .end
```

Note [1]

As in C programs, this includes the central processor-specific file containing the IO port definitions for the device. Note that not all include files can be included into assembler sources.

Note [2]

Assignment of registers to symbolic names used locally. Another option would be to use a C preprocessor macro instead:

```
#define work 16
```

Note [3]

Our bit number for the square wave output. Note that the right-hand side consists of a CPP macro which will be substituted by its value (6 in this case) before actually being passed to the assembler.

Note [4]

The assembler uses integer operations in the host-defined integer size (32 bits or longer) when evaluating expressions. This is in contrast to the C compiler that uses the C type `int` by default in order to calculate constant integer expressions.

In order to get a 100 kHz output, we need to toggle the PD6 line 200000 times per second. Since we use timer 0 without any prescaling options in order to get the desired frequency and accuracy, we already run into serious timing considerations: while accepting and processing the timer overflow interrupt, the timer already continues to count. When pre-loading the `TCCNT0` register, we therefore have to account for the number of clock cycles required for interrupt acknowledge and for the instructions to reload `TCCNT0` (4 clock cycles for interrupt acknowledge, 2 cycles for the jump from the interrupt vector, 2 cycles for the 2 instructions that reload `TCCNT0`). This is what the constant `fuzz` is for.

Note [5]

External functions need to be declared to be `.global`. `main` is the application entry point that will be jumped to from the initialization routine in `crt0.o`.

Note [6]

The main loop is just a single jump back to itself. Square wave generation itself is completely handled by the timer 0 overflow interrupt service. A `sleep` instruction (using idle mode) could be used as well, but probably would not conserve much energy anyway since the interrupt service is executed quite frequently.

Note [7]

Interrupt functions can get the [usual names](#) that are also available to C programs. The linker will then put them into the appropriate interrupt vector slots. Note that they must be declared `.global` in order to be acceptable for this purpose. This will only work if `<avr/io.h>` has been included. Note that the assembler or linker have no chance to check the correct spelling of an interrupt function, so it should be double-checked. (When analyzing the resulting object file using `avr-objdump` or `avr-nm`, a name like `__vector_N` should appear, with `N` being a small integer number.)

Note [8]

As explained in the section about [special function registers](#), the actual IO port address should be obtained using the macro `_SFR_IO_ADDR`. (The AT90S1200 does not have RAM thus the memory-mapped approach to access the IO registers is not available. It would be slower than using `in / out` instructions anyway.)

Since the operation to reload `TCCNT0` is time-critical, it is even performed before saving `SREG`. Obviously, this requires that the instructions involved would not change any of the flag bits in `SREG`.

Note [9]

Interrupt routines must not clobber the global CPU state. Thus, it is usually necessary to save at least the state of the flag bits in `SREG`. (Note that this serves as an example here only since actually, all the following instructions would not modify `SREG` either, but that's not commonly the case.)

Also, it must be made sure that registers used inside the interrupt routine do not conflict with those used outside. In the case of a RAM-less device like the AT90S1200, this can only be done by agreeing on a set of registers to be used exclusively inside the interrupt routine; there would not be any other chance to "save" a register anywhere.

If the interrupt routine is to be linked together with C modules, care must be taken to follow the [register usage guidelines](#) imposed by the C compiler. Also, any register modified inside the interrupt service needs to be saved, usually on the stack.

Note [10]

As explained in [Interrupts](#), a global "catch-all" interrupt handler that gets all unassigned interrupt vectors can be installed using the name `__vector_default`. This must be `.global`, and obviously, should end in a `reti` instruction. (By default, a jump to location 0 would be implied instead.)

6.4 Pseudo-ops and operators

The available pseudo-ops in the assembler are described in the GNU assembler (`gas`) manual. The manual can be found online as part of the current `binutils` release under <http://sources.redhat.com/binutils/>.

As `gas` comes from a Unix origin, its pseudo-op and overall assembler syntax is slightly different than the one being used by other assemblers. Numeric constants follow the C notation (prefix `0x` for hexadecimal constants), expressions use a C-like syntax.

Some common pseudo-ops include:

- `.byte` allocates single byte constants
- `.ascii` allocates a non-terminated string of characters
- `.asciz` allocates a `\0`-terminated string of characters (C string)
- `.data` switches to the `.data` section (initialized RAM variables)
- `.text` switches to the `.text` section (code and ROM constants)
- `.set` declares a symbol as a constant expression (identical to `.equ`)

- `.global` (or `.globl`) declares a public symbol that is visible to the linker (e. g. function entry point, global variable)
- `.extern` declares a symbol to be externally defined; this is effectively a comment only, as `gas` treats all undefined symbols it encounters as globally undefined anyway

Note that `.org` is available in `gas` as well, but is a fairly pointless pseudo-op in an assembler environment that uses relocatable object files, as it is the linker that determines the final position of some object in ROM or RAM.

Along with the architecture-independent standard operators, there are some AVR-specific operators available which are unfortunately not yet described in the official documentation. The most notable operators are:

- `lo8` Takes the least significant 8 bits of a 16-bit integer
- `hi8` Takes the most significant 8 bits of a 16-bit integer
- `pm` Takes a program-memory (ROM) address, and converts it into a RAM address. This implies a division by 2 as the AVR handles ROM addresses as 16-bit words (e.g. in an `IJMP` or `ICALL` instruction), and can also handle relocatable symbols on the right-hand side.

Example:

```
ldi r24, lo8(pm(somefunc))
ldi r25, hi8(pm(somefunc))
call something
```

This passes the address of function `somefunc` as the first parameter to function `something`.

7 Inline Assembler Cookbook

AVR-GCC

Inline Assembler Cookbook

About this Document

The GNU C compiler for Atmel AVR RISC processors offers, to embed assembly language code into C programs. This cool feature may be used for manually optimizing time critical parts of the software or to use specific processor instruction, which are not available in the C language.

Because of a lack of documentation, especially for the AVR version of the compiler, it may take some time to figure out the implementation details by studying the compiler and assembler source code. There are also a few sample programs available in the net. Hopefully this document will help to increase their number.

It's assumed, that you are familiar with writing AVR assembler programs, because this is not an AVR assembler programming tutorial. It's not a C language tutorial either.

Note that this document does not cover file written completely in assembler language, refer to [avr-libc and assembler programs](#) for this.

Copyright (C) 2001-2002 by egnite Software GmbH

Permission is granted to copy and distribute verbatim copies of this manual provided that the copyright notice and this permission notice are preserved on all copies. Permission is granted to copy and distribute modified versions of this manual provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

This document describes version 3.3 of the compiler. There may be some parts, which hadn't been completely understood by the author himself and not all samples had been tested so far. Because the author is German and not familiar with the English language, there are definitely some typos and syntax errors in the text. As a programmer the author knows, that a wrong documentation sometimes might be worse than none. Anyway, he decided to offer his little knowledge to the public, in the hope to get enough response to improve this document. Feel free to contact the author via e-mail. For the latest release check <http://www.ethernut.de/>.

Herne, 17th of May 2002 Harald Kipp harald.kipp-at-egnite.de

Note

As of 26th of July 2002, this document has been merged into the documentation for avr-libc. The latest version is now available at <http://savannah.nongnu.org/projects/avr-libc/>.

7.1 GCC asm Statement

Let's start with a simple example of reading a value from port D:

```
asm("in %0, %1" : "=r" (value) : "I" (_SFR_IO_ADDR(PORTD)) );
```

Each `asm` statement is divided by colons into (up to) four parts:

1. The assembler instructions, defined as a single string constant:

```
"in %0, %1"
```

2. A list of output operands, separated by commas. Our example uses just one:

```
"=r" (value)
```

3. A comma separated list of input operands. Again our example uses one operand only:

```
"I" (_SFR_IO_ADDR(PORTD))
```

4. Clobbered registers, left empty in our example.

You can write assembler instructions in much the same way as you would write assembler programs. However, registers and constants are used in a different way if they refer to expressions of your C program. The connection between registers and C operands is specified in the second and third part of the `asm` instruction, the list of input and output operands, respectively. The general form is

```
asm(code : output operand list : input operand list [: clobber list]);
```

In the code section, operands are referenced by a percent sign followed by a single digit. 0 refers to the first 1 to the second operand and so forth. From the above example:

0 refers to `"=r" (value)` and

1 refers to `"I" (_SFR_IO_ADDR(PORTD))`.

This may still look a little odd now, but the syntax of an operand list will be explained soon. Let us first examine the part of a compiler listing which may have been generated from our example:

```
        lds r24,value
/* #APP */
        in r24, 12
/* #NOAPP */
        sts value,r24
```

The comments have been added by the compiler to inform the assembler that the included code was not generated by the compilation of C statements, but by inline assembler statements. The compiler selected register `r24` for storage of the value read from `PORTD`. The compiler could have selected any other register, though. It may not explicitly load or store the value and it may even decide not to include your assembler code at all. All these decisions are part of the compiler's optimization strategy. For example, if you never use the variable `value` in the remaining part of the C program, the compiler will most likely remove your code unless you switched off optimization. To avoid this, you can add the `volatile` attribute to the `asm` statement:

```
asm volatile("in %0, %1" : "=r" (value) : "I" (_SFR_IO_ADDR(PORTD)));
```

Alternatively, operands can be given names. The name is prepended in brackets to the constraints in the operand list, and references to the named operand use the bracketed name instead of a number after the `%` sign. Thus, the above example could also be written as

```
asm("in %[retval], %[port]" :
    [retval] "=r" (value) :
    [port] "I" (_SFR_IO_ADDR(PORTD)) );
```

The last part of the `asm` instruction, the clobber list, is mainly used to tell the compiler about modifications done by the assembler code. This part may be omitted, all other parts are required, but may be left empty. If your assembler routine won't use any input or output operand, two colons must still follow the assembler code string. A good example is a simple statement to disable interrupts:

```
asm volatile("cli");
```

7.2 Assembler Code

You can use the same assembler instruction mnemonics as you'd use with any other AVR assembler. And you can write as many assembler statements into one code string as you like and your flash memory is able to hold.

Note

The available assembler directives vary from one assembler to another.

To make it more readable, you should put each statement on a separate line:

```
asm volatile("nop\n\t"
             "nop\n\t"
             "nop\n\t"
             "nop\n\t"
             :);
```

The linefeed and tab characters will make the assembler listing generated by the compiler more readable. It may look a bit odd for the first time, but that's the way the compiler creates its own assembler code.

You may also make use of some special registers.

Symbol	Register
<code>__SREG__</code>	Status register at address 0x3F
<code>__SP_H__</code>	Stack pointer high byte at address 0x3E
<code>__SP_L__</code>	Stack pointer low byte at address 0x3D
<code>__tmp_reg__</code>	Register r0, used for temporary storage
<code>__zero_reg__</code>	Register r1, always zero

Register r0 may be freely used by your assembler code and need not be restored at the end of your code. It's a good idea to use `__tmp_reg__` and `__zero_reg__` instead of r0 or r1, just in case a new compiler version changes the register usage definitions.

7.3 Input and Output Operands

Each input and output operand is described by a constraint string followed by a C expression in parentheses. AVR-GCC 3.3 knows the following constraint characters:

Note

The most up-to-date and detailed information on constraints for the avr can be found in the gcc manual.

The x register is r27:r26, the y register is r29:r28, and the z register is r31:r30

Constraint	Used for	Range
a	Simple upper registers	r16 to r23
b	Base pointer registers pairs	y, z
d	Upper register	r16 to r31
e	Pointer register pairs	x, y, z
q	Stack pointer register	SPH:SPL
r	Any register	r0 to r31
t	Temporary register	r0
w	Special upper register pairs	r24, r26, r28, r30
x	Pointer register pair X	x (r27:r26)
y	Pointer register pair Y	y (r29:r28)
z	Pointer register pair Z	z (r31:r30)
G	Floating point constant	0.0
I	6-bit positive integer constant	0 to 63
J	6-bit negative integer constant	-63 to 0
K	Integer constant	2
L	Integer constant	0
l	Lower registers	r0 to r15
M	8-bit integer constant	0 to 255
N	Integer constant	-1
O	Integer constant	8, 16, 24
P	Integer constant	1
Q	(GCC >= 4.2.x) A memory address based on Y or Z pointer with displacement.	
R	(GCC >= 4.3.x) Integer constant.	-6 to 5

The selection of the proper constraint depends on the range of the constants or registers, which must be acceptable to the AVR instruction they are used with. The C compiler doesn't check any line of your assembler code. But it is able to check the constraint against your C expression. However, if you specify the wrong constraints, then the compiler may silently pass wrong code to the assembler. And, of course, the assembler will fail with some cryptic output or internal errors. For example, if you specify the constraint "r" and you are using this register with an "ori" instruction in your assembler code, then the compiler may select any register. This will fail, if the compiler chooses r2 to r15. (It will never choose r0 or r1, because these are used for special purposes.) That's why the correct constraint in that case is "d". On the other hand, if you use the constraint "M", the compiler will make sure that you don't pass anything else but an 8-bit value. Later on we will see how to pass multibyte expression results to the assembler code.

The following table shows all AVR assembler mnemonics which require operands, and the related constraints. Because of the improper constraint definitions in version 3.3, they aren't strict enough. There is, for example, no constraint, which restricts integer constants to the range 0 to 7 for bit set and bit clear operations.

Mnemonic	Constraints		Mnemonic	Constraints
adc	r,r		add	r,r
adiw	w,I		and	r,r
andi	d,M		asr	r
bclr	I		bld	r,I
brbc	I,label		brbs	I,label
bset	I		bst	r,I
cbi	I,I		cbr	d,I
com	r		cp	r,r
cpc	r,r		cpi	d,M
cpse	r,r		dec	r
elpm	t,z		eor	r,r
in	r,I		inc	r
ld	r,e		ldd	r,b
ldi	d,M		lds	r,label
lpm	t,z		lsl	r
lsr	r		mov	r,r
movw	r,r		mul	r,r
neg	r		or	r,r
ori	d,M		out	I,r
pop	r		push	r
rol	r		ror	r
sbc	r,r		sbc	d,M
sbi	I,I		sbic	I,I
sbiw	w,I		sbr	d,M
sbr	r,I		sbrs	r,I
ser	d		st	e,r
std	b,r		sts	label,r
sub	r,r		subi	d,M
swap	r			

Constraint characters may be prepended by a single constraint modifier. Constraints without a modifier specify read-only operands. Modifiers are:

Modifier	Specifies
=	Write-only operand, usually used for all output operands.
+	Read-write operand
&	Register should be used for output only

Output operands must be write-only and the C expression result must be an lvalue, which means that the operands must be valid on the left side of assignments. Note, that the compiler will not check if the operands are of reasonable type for the kind of operation used in the assembler instructions.

Input operands are, you guessed it, read-only. But what if you need the same operand for input and output? As stated above, read-write operands are not supported in inline assembler code. But there is another solution. For input operators it is possible to use a single digit in the constraint string. Using digit *n* tells the compiler to use the same register as for the *n*-th operand, starting with zero. Here is an example:

```
asm volatile("swap %0" : "=r" (value) : "0" (value));
```

This statement will swap the nibbles of an 8-bit variable named `value`. Constraint `"0"` tells the compiler, to use the same input register as for the first operand. Note however, that this doesn't automatically imply the reverse case. The compiler may choose the same registers for input and output, even if not told to do so. This is not a problem in most cases, but may be fatal if the output operator is modified by the assembler code before the input operator is used. In the situation where your code depends on different registers used for input and output operands, you must add the `&` constraint modifier to your output operand. The following example demonstrates this problem:

```
asm volatile("in %0,%1"      "\n\t"
             "out %1, %2"    "\n\t"
             : "=&r" (input)
             : "I" (_SFR_IO_ADDR(port)), "r" (output)
             );
```

In this example an input value is read from a port and then an output value is written to the same port. If the compiler would have chosen the same register for input and output, then the output value would have been destroyed on the first assembler instruction. Fortunately, this example uses the `&` constraint modifier to instruct the compiler not to select any register for the output value, which is used for any of the input operands. Back to swapping. Here is the code to swap high and low byte of a 16-bit value:

```
asm volatile("mov __tmp_reg__, %A0" "\n\t"
             "mov %A0, %B0"         "\n\t"
             "mov %B0, __tmp_reg__" "\n\t"
             : "=r" (value)
             : "0" (value)
             );
```

First you will notice the usage of register `__tmp_reg__`, which we listed among other special registers in the [Assembler Code](#) section. You can use this register without saving its contents. Completely new are those letters A and B in `%A0` and `%B0`. In fact they refer to two different 8-bit registers, both containing a part of value.

Another example to swap bytes of a 32-bit value:

```
asm volatile("mov __tmp_reg__, %A0" "\n\t"
             "mov %A0, %D0"         "\n\t"
             "mov %D0, __tmp_reg__" "\n\t"
             "mov __tmp_reg__, %B0" "\n\t"
             "mov %B0, %C0"         "\n\t"
             "mov %C0, __tmp_reg__" "\n\t"
             : "=r" (value)
             : "0" (value)
             );
```

Instead of listing the same operand as both, input and output operand, it can also be declared as a read-write operand. This must be applied to an output operand, and the respective input operand list remains empty:

```
asm volatile("mov __tmp_reg__, %A0" "\n\t"
             "mov %A0, %D0"         "\n\t"
             "mov %D0, __tmp_reg__" "\n\t"
```



```

"mov __tmp_reg__, %B0" "\n\t"
"mov %B0, %C0"          "\n\t"
"mov %C0, __tmp_reg__"  "\n\t"
: "+r" (value));

```

If operands do not fit into a single register, the compiler will automatically assign enough registers to hold the entire operand. In the assembler code you use %A0 to refer to the lowest byte of the first operand, %A1 to the lowest byte of the second operand and so on. The next byte of the first operand will be %B0, the next byte %C0 and so on.

This also implies, that it is often necessary to cast the type of an input operand to the desired size.

A final problem may arise while using pointer register pairs. If you define an input operand

```
"e" (ptr)
```

and the compiler selects register Z (r30:r31), then

%A0 refers to r30 and

%B0 refers to r31.

But both versions will fail during the assembly stage of the compiler, if you explicitly need Z, like in

```
ld r24,Z
```

If you write

```
ld r24, %a0
```

with a lower case a following the percent sign, then the compiler will create the proper assembler line.

7.4 Clobbers

As stated previously, the last part of the `asm` statement, the list of clobbers, may be omitted, including the colon separator. However, if you are using registers, which had not been passed as operands, you need to inform the compiler about this. The following example will do an atomic increment. It increments an 8-bit value pointed to by a pointer variable in one go, without being interrupted by an interrupt routine or another thread in a multithreaded environment. Note, that we must use a pointer, because the incremented value needs to be stored before interrupts are enabled.

```

asm volatile(
    "cli"                "\n\t"
    "ld r24, %a0"         "\n\t"
    "inc r24"            "\n\t"
    "st %a0, r24"         "\n\t"
    "sei"                "\n\t"

```

```

:
: "e" (ptr)
: "r24"
);

```

The compiler might produce the following code:

```

cli
ld r24, Z
inc r24
st Z, r24
sei

```

One easy solution to avoid clobbering register `r24` is, to make use of the special temporary register `__tmp_reg__` defined by the compiler.

```

asm volatile(
    "cli"                "\n\t"
    "ld __tmp_reg__, %a0" "\n\t"
    "inc __tmp_reg__"     "\n\t"
    "st %a0, __tmp_reg__" "\n\t"
    "sei"                "\n\t"
    :
    : "e" (ptr)
);

```

The compiler is prepared to reload this register next time it uses it. Another problem with the above code is, that it should not be called in code sections, where interrupts are disabled and should be kept disabled, because it will enable interrupts at the end. We may store the current status, but then we need another register. Again we can solve this without clobbering a fixed, but let the compiler select it. This could be done with the help of a local C variable.

```

{
    uint8_t s;
    asm volatile(
        "in %0, __SREG__"    "\n\t"
        "cli"                "\n\t"
        "ld __tmp_reg__, %a1" "\n\t"
        "inc __tmp_reg__"     "\n\t"
        "st %a1, __tmp_reg__" "\n\t"
        "out __SREG__, %0"    "\n\t"
        : "=&r" (s)
        : "e" (ptr)
    );
}

```

Now every thing seems correct, but it isn't really. The assembler code modifies the variable, that `ptr` points to. The compiler will not recognize this and may keep its value in any of the other registers. Not only does the compiler work with the wrong value, but the assembler code does too. The C program may have modified the value too, but the compiler didn't update the memory location for optimization reasons. The worst thing you can do in this case is:

```

{
    uint8_t s;
    asm volatile(
        "in %0, __SREG__"           "\n\t"
        "cli"                       "\n\t"
        "ld __tmp_reg__, %a1"       "\n\t"
        "inc __tmp_reg__"           "\n\t"
        "st %a1, __tmp_reg__"       "\n\t"
        "out __SREG__, %0"          "\n\t"
        : "=&r" (s)
        : "e" (ptr)
        : "memory"
    );
}

```

The special clobber "memory" informs the compiler that the assembler code may modify any memory location. It forces the compiler to update all variables for which the contents are currently held in a register before executing the assembler code. And of course, everything has to be reloaded again after this code.

In most situations, a much better solution would be to declare the pointer destination itself volatile:

```
volatile uint8_t *ptr;
```

This way, the compiler expects the value pointed to by `ptr` to be changed and will load it whenever used and store it whenever modified.

Situations in which you need clobbers are very rare. In most cases there will be better ways. Clobbered registers will force the compiler to store their values before and reload them after your assembler code. Avoiding clobbers gives the compiler more freedom while optimizing your code.

7.5 Assembler Macros

In order to reuse your assembler language parts, it is useful to define them as macros and put them into include files. AVR Libc comes with a bunch of them, which could be found in the directory `avr/include`. Using such include files may produce compiler warnings, if they are used in modules, which are compiled in strict ANSI mode. To avoid that, you can write `__asm__` instead of `asm` and `__volatile__` instead of `volatile`. These are equivalent aliases.

Another problem with reused macros arises if you are using labels. In such cases you may make use of the special pattern `=`, which is replaced by a unique number on each `asm` statement. The following code had been taken from `avr/include/iomacros.h`:

```

#define loop_until_bit_is_clear(port,bit) \
    __asm__ __volatile__ ( \
        "L_%=: " "sbic %0, %1" "\n\t" \
        "rjmp L_%" \
        : /* no outputs */ \
        : "I" (_SFR_IO_ADDR(port)), \
        "I" (bit) \
    )

```

When used for the first time, `L_` may be translated to `L_1404`, the next usage might create `L_1405` or whatever. In any case, the labels became unique too.

Another option is to use Unix-assembler style numeric labels. They are explained in [How do I trace an assembler file in avr-gdb?](#). The above example would then look like:

```
#define loop_until_bit_is_clear(port,bit)
    __asm__ __volatile__ (
        "1: " "sbic %0, %1" "\n\t"
            "rjmp 1b"
        : /* no outputs */
        : "I" (_SFR_IO_ADDR(port)),
            "I" (bit)
    )
```

7.6 C Stub Functions

Macro definitions will include the same assembler code whenever they are referenced. This may not be acceptable for larger routines. In this case you may define a C stub function, containing nothing other than your assembler code.

```
void delay(uint8_t ms)
{
    uint16_t cnt;
    asm volatile (
        "\n"
        "L_d11%=: " "\n\t"
        "mov %A0, %A2" "\n\t"
        "mov %B0, %B2" "\n"
        "L_d12%=: " "\n\t"
        "sbiw %A0, 1" "\n\t"
        "brne L_d12%=" "\n\t"
        "dec %1" "\n\t"
        "brne L_d11%=" "\n\t"
        : "=&w" (cnt)
        : "r" (ms), "r" (delay_count)
    );
}
```

The purpose of this function is to delay the program execution by a specified number of milliseconds using a counting loop. The global 16 bit variable `delay_count` must contain the CPU clock frequency in Hertz divided by 4000 and must have been set before calling this routine for the first time. As described in the [clobber](#) section, the routine uses a local variable to hold a temporary value.

Another use for a local variable is a return value. The following function returns a 16 bit value read from two successive port addresses.

```
uint16_t inw(uint8_t port)
{
    uint16_t result;
    asm volatile (
        "in %A0,%1" "\n\t"
        "in %B0, (%1) + 1"
        : "=r" (result)
        : "I" (_SFR_IO_ADDR(port))
    )
```

```
    );  
    return result;  
}
```

Note

`inw()` is supplied by `avr-libc`.

7.7 C Names Used in Assembler Code

By default AVR-GCC uses the same symbolic names of functions or variables in C and assembler code. You can specify a different name for the assembler code by using a special form of the `asm` statement:

```
unsigned long value asm("clock") = 3686400;
```

This statement instructs the compiler to use the symbol name `clock` rather than `value`. This makes sense only for external or static variables, because local variables do not have symbolic names in the assembler code. However, local variables may be held in registers.

With AVR-GCC you can specify the use of a specific register:

```
void Count(void)  
{  
    register unsigned char counter asm("r3");  
  
    ... some code...  
    asm volatile("clr r3");  
    ... more code...  
}
```

The assembler instruction, `"clr r3"`, will clear the variable `counter`. AVR-GCC will not completely reserve the specified register. If the optimizer recognizes that the variable will not be referenced any longer, the register may be re-used. But the compiler is not able to check whether this register usage conflicts with any predefined register. If you reserve too many registers in this way, the compiler may even run out of registers during code generation.

In order to change the name of a function, you need a prototype declaration, because the compiler will not accept the `asm` keyword in the function definition:

```
extern long Calc(void) asm ("CALCULATE");
```

Calling the function `Calc()` will create assembler instructions to call the function `CALCULATE`.

7.8 Links

For a more thorough discussion of inline assembly usage, see the gcc user manual. The latest version of the gcc manual is always available here: <http://gcc.gnu.org/onlinedocs/>

8 How to Build a Library

8.1 Introduction

So you keep reusing the same functions that you created over and over? Tired of cut and paste going from one project to the next? Would you like to reduce your maintenance overhead? Then you're ready to create your own library! Code reuse is a very laudable goal. With some upfront investment, you can save time and energy on future projects by having ready-to-go libraries. This chapter describes some background information, design considerations, and practical knowledge that you will need to create and use your own libraries.

8.2 How the Linker Works

The compiler compiles a single high-level language file (C language, for example) into a single object module file. The linker (`ld`) can only work with object modules to link them together. Object modules are the smallest unit that the linker works with.

Typically, on the linker command line, you will specify a set of object modules (that has been previously compiled) and then a list of libraries, including the Standard C Library. The linker takes the set of object modules that you specify on the command line and links them together. Afterwards there will probably be a set of "undefined references". A reference is essentially a function call. An undefined reference is a function call, with no defined function to match the call.

The linker will then go through the libraries, in order, to match the undefined references with function definitions that are found in the libraries. If it finds the function that matches the call, the linker will then link in the object module in which the function is located. This part is important: the linker links in THE ENTIRE OBJECT MODULE in which the function is located. Remember, the linker knows nothing about the functions internal to an object module, other than symbol names (such as function names). The smallest unit the linker works with is object modules.

When there are no more undefined references, the linker has linked everything and is done and outputs the final application.

8.3 How to Design a Library

How the linker behaves is very important in designing a library. Ideally, you want to design a library where only the functions that are called are the only functions to be linked into the final application. This helps keep the code size to a minimum. In order to do this, with the way the linker works, is to only write one function per code module. This will compile to one function per object module. This is usually a very different way of doing things than writing an application!

There are always exceptions to the rule. There are generally two cases where you would want to have more than one function per object module.

The first is when you have very complementary functions that it doesn't make much sense to split them up. For example, `malloc()` and `free()`. If someone is going to use

`malloc()`, they will very likely be using `free()` (or at least should be using `free()`). In this case, it makes more sense to aggregate those two functions in the same object module.

The second case is when you want to have an Interrupt Service Routine (ISR) in your library that you want to link in. The problem in this case is that the linker looks for unresolved references and tries to resolve them with code in libraries. A reference is the same as a function call. But with ISRs, there is no function call to initiate the ISR. The ISR is placed in the Interrupt Vector Table (IVT), hence no call, no reference, and no linking in of the ISR. In order to do this, you have to trick the linker in a way. Aggregate the ISR, with another function in the same object module, but have the other function be something that is required for the user to call in order to use the ISR, like perhaps an initialization function for the subsystem, or perhaps a function that enables the ISR in the first place.

8.4 Creating a Library

The librarian program is called `ar` (for "archiver") and is found in the GNU Binutils project. This program will have been built for the AVR target and will therefore be named `avr-ar`.

The job of the librarian program is simple: aggregate a list of object modules into a single library (archive) and create an index for the linker to use. The name that you create for the library filename must follow a specific pattern: `libname.a`. The *name* part is the unique part of the filename that you create. It makes it easier if the *name* part relates to what the library is about. This *name* part must be prefixed by "lib", and it must have a file extension of `.a`, for "archive". The reason for the special form of the filename is for how the library gets used by the toolchain, as we will see later on.

Note

The filename is case-sensitive. Use a lowercase "lib" prefix, and a lowercase ".a" as the file extension.

The command line is fairly simple:

```
avr-ar rcs <library name> <list of object modules>
```

The `r` command switch tells the program to insert the object modules into the archive with replacement. The `c` command line switch tells the program to create the archive. And the `s` command line switch tells the program to write an object-file index into the archive, or update an existing one. This last switch is very important as it helps the linker to find what it needs to do its job.

Note

The command line switches are case sensitive! There are uppercase switches that have completely different actions.

MFile and the WinAVR distribution contain a Makefile Template that includes the necessary command lines to build a library. You will have to manually modify the template to switch it over to build a library instead of an application.

See the GNU Binutils manual for more information on the `ar` program.

8.5 Using a Library

To use a library, use the `-l` switch on your linker command line. The string immediately following the `-l` is the unique part of the library filename that the linker will link in. For example, if you use:

```
-lm
```

this will expand to the library filename:

```
libm.a
```

which happens to be the math library included in `avr-libc`.

If you use this on your linker command line:

```
-lprintf_flt
```

then the linker will look for a library called:

```
libprintf_flt.a
```

This is why naming your library is so important when you create it!

The linker will search libraries in the order that they appear on the command line. Whichever function is found first that matches the undefined reference, it will be linked in.

There are also command line switches that tell GCC which directory to look in (`-L`) for the libraries that are specified to be linked in with `-l`.

See the GNU Binutils manual for more information on the GNU linker (`ld`) program.

9 Benchmarks

The results below can only give a rough estimate of the resources necessary for using certain library functions. There is a number of factors which can both increase or reduce the effort required:

- Expenses for preparation of operands and their stack are not considered.
- In the table, the size includes all additional functions (for example, function to multiply two integers) but they are only linked from the library.
- Expenses of time of performance of some functions essentially depend on parameters of a call, for example, `qsort()` is recursive, and `sprintf()` receives parameters in a stack.
- Different versions of the compiler can give a significant difference in code size and execution time. For example, the `dtostre()` function, compiled with `avr-gcc 3.4.6`, requires 930 bytes. After transition to `avr-gcc 4.2.3`, the size become 1088 bytes.

9.1 A few of libc functions.

Avr-gcc version is 4.7.1

The size of function is given in view of all picked up functions. By default Avr-libc is compiled with `-mcall-prologues` option. In brackets the size without taking into account modules of a prologue and an epilogue is resulted. Both of the size can coincide, if function does not cause a prologue/epilogue.

Function	Units	Avr2	Avr25	Avr4
atoi ("12345")	Flash bytes	82 (82)	78 (78)	74 (74)
	Stack bytes	2	2	2
	MCU clocks	155	149	149
atol ("12345")	Flash bytes	122 (122)	118 (118)	118 (118)
	Stack bytes	2	2	2
	MCU clocks	221	219	219
dtoa (1.2345, s, 6, 0)	Flash bytes	1116 (1004)	1048 (938)	1048 (938)
	Stack bytes	17	17	17
	MCU clocks	1247	1105	1105
dtostrf (1.2345, 15, 6, s)	Flash bytes	1616 (1616)	1508 (1508)	1508 (1508)
	Stack bytes	38	38	38
	MCU clocks	1634	1462	1462
itoa (12345, s, 10)	Flash bytes	110 (110)	102 (102)	102 (102)
	Stack bytes	2	2	2
	MCU clocks	879	875	875
ltoa (12345L, s, 10)	Flash bytes	134 (134)	126 (126)	126 (126)
	Stack bytes	2	2	2
	MCU clocks	1597	1593	1593
malloc (1)	Flash bytes	768 (712)	714 (660)	714 (660)
	Stack bytes	6	6	6
	MCU clocks	215	201	201
realloc ((void *)0, 1)	Flash bytes	1284 (1172)	1174 (1064)	1174 (1064)
	Stack bytes	18	18	18
	MCU clocks	305	286	286
qsort (s, sizeof(s), 1, cmp)	Flash bytes	1252 (1140)	1022 (912)	1028 (918)
	Stack bytes	42	42	42
	MCU clocks	21996	19905	17541
sprintf_min (s, "%d", 12345)	Flash bytes	1224 (1112)	1092 (982)	1088 (978)
	Stack bytes	53	53	53
	MCU clocks	1841	1694	1689
sprintf (s, "%d", 12345)	Flash bytes	1614 (1502)	1476 (1366)	1454 (1344)
	Stack bytes	58	58	58
	MCU clocks	1647	1552	1547
sprintf_flt (s, "%e", 1.2345)	Flash bytes	3228 (3116)	2990 (2880)	2968 (2858)
	Stack bytes	67	67	67
	MCU clocks	2573	2311	2311
sscanf_min ("12345", "%d", &i)	Flash bytes	1532 (1420)	1328 (1218)	1328 (1218)
	Stack bytes	55	55	55
	MCU clocks	1607	1446	1446
sscanf ("12345", "%d", &i)	Flash bytes	2008 (1896)	1748 (1638)	1748 (1638)
	Stack bytes	55	55	55
	MCU clocks	1610	1449	1449
sscanf ("point,color", "%[a-z]", s)	Flash bytes	2008 (1896)	1748 (1638)	1748 (1638)
	Stack bytes	86	86	86
	MCU clocks	3067	2806	2806
sscanf_flt ("1.2345", "%e", &x)	Flash bytes	3464 (3352)	3086 (2976)	3070 (2960)
	Stack bytes	71	71	71
	MCU clocks	2497	2281	2078
strtod ("1.2345", &p)	Flash bytes	1632 (1520)	1536 (1426)	1480 (1480)
	Stack bytes	20	20	21
	MCU clocks	1235	1177	1124
strtoul ("12345", &p, 0)	Flash bytes	918 (806)	834 (724)	792 (792)
	Stack bytes	22	22	28
	MCU clocks	956	891	794

9.2 Math functions.

The table contains the number of MCU clocks to calculate a function with a given argument(s). The main reason of a big difference between Avr2 and Avr4 is a hardware

multiplication.

Function	Avr2	Avr4
__addsf3 (1.234, 5.678)	113	108
__mulsf3 (1.234, 5.678)	375	138
__divsf3 (1.234, 5.678)	466	465
acos (0.54321)	4411	2455
asin (0.54321)	4517	2556
atan (0.54321)	4710	2271
atan2 (1.234, 5.678)	5270	2857
cbrt (1.2345)	2684	2555
ceil (1.2345)	177	177
cos (1.2345)	3387	1671
cosh (1.2345)	4922	2979
exp (1.2345)	4708	2765
fdim (5.678, 1.234)	111	111
floor (1.2345)	180	180
fmax (1.234, 5.678)	39	37
fmin (1.234, 5.678)	35	35
fmod (5.678, 1.234)	131	131
frexp (1.2345, 0)	42	41
hypot (1.234, 5.678)	1341	866
ldexp (1.2345, 6)	42	42
log (1.2345)	4142	2134
log10 (1.2345)	4498	2260
modf (1.2345, 0)	433	429
pow (1.234, 5.678)	9293	5047
round (1.2345)	150	150
sin (1.2345)	3353	1653
sinh (1.2345)	4946	3003
sqrt (1.2345)	494	492
tan (1.2345)	4381	2426
tanh (1.2345)	5126	3173
trunc (1.2345)	178	178

10 Porting From IAR to AVR GCC

10.1 Introduction

C language was designed to be a portable language. There two main types of porting activities: porting an application to a different platform (OS and/or processor), and porting to a different compiler. Porting to a different compiler can be exacerbated when the application is an embedded system. For example, the C language Standard, strangely, does not specify a standard for declaring and defining Interrupt Service Routines (ISRs). Different compilers have different ways of defining registers, some of which use non-standard language constructs.

This chapter describes some methods and pointers on porting an AVR application built with the IAR compiler to the GNU toolchain (AVR GCC). Note that this may not be

an exhaustive list.

10.2 Registers

IO header files contain identifiers for all the register names and bit names for a particular processor. IAR has individual header files for each processor and they must be included when registers are being used in the code. For example:

```
#include <iom169.h>
```

Note

IAR does not always use the same register names or bit names that are used in the AVR datasheet.

AVR GCC also has individual IO header files for each processor. However, the actual processor type is specified as a command line flag to the compiler. (Using the `-mmcu=processor` flag.) This is usually done in the Makefile. This allows you to specify only a single header file for any processor type:

```
#include <avr/io.h>
```

Note

The forward slash in the `<avr/io.h>` file name that is used to separate subdirectories can be used on Windows distributions of the toolchain and is the recommended method of including this file.

The compiler knows the processor type and through the single header file above, it can pull in and include the correct individual IO header file. This has the advantage that you only have to specify one generic header file, and you can easily port your application to another processor type without having to change every file to include the new IO header file.

The AVR toolchain tries to adhere to the exact names of the registers and names of the bits found in the AVR datasheet. There may be some discrepancies between the register names found in the IAR IO header files and the AVR GCC IO header files.

10.3 Interrupt Service Routines (ISRs)

As mentioned above, the C language Standard, strangely, does not specify a standard way of declaring and defining an ISR. Hence, every compiler seems to have their own special way of doing so.

IAR declares an ISR like so:

```
#pragma vector=TIMER0_OVF_vect
__interrupt void MotorPWMBottom()
{
    // code
}
```

In AVR GCC, you declare an ISR like so:

```
ISR(PCINT1_vect)
{
    //code
}
```

AVR GCC uses the `ISR` macro to define an ISR. This macro requires the header file:

```
#include <avr/interrupt.h>
```

The names of the various interrupt vectors are found in the individual processor IO header files that you must include with `<avr/io.h>`.

Note

The names of the interrupt vectors in AVR GCC has been changed to match the names of the vectors in IAR. This significantly helps in porting applications from IAR to AVR GCC.

10.4 Intrinsic Routines

IAR has a number of intrinsic routine such as

```
__enable_interrupts() __disable_interrupts() __watchdog__
reset()
```

These intrinsic functions compile to specific AVR opcodes (SEI, CLI, WDR).

There are equivalent macros that are used in AVR GCC, however they are not located in a single include file.

AVR GCC has `sei()` for `__enable_interrupts()`, and `cli()` for `__disable_interrupts()`. Both of these macros are located in `<avr/interrupt.h>`.

AVR GCC has the macro `wdt_reset()` in place of `__watchdog_reset()`. However, there is a whole Watchdog Timer API available in AVR GCC that can be found in `<avr/wdt.h>`.

10.5 Flash Variables

The C language was not designed for Harvard architecture processors with separate memory spaces. This means that there are various non-standard ways to define a variable whose data resides in the Program Memory (Flash).

IAR uses a non-standard keyword to declare a variable in Program Memory:

```
__flash int mydata[] = ....
```

AVR GCC uses Variable Attributes to achieve the same effect:

```
int mydata[] __attribute__((progmem))
```

Note

See the GCC User Manual for more information about Variable Attributes.

avr-libc provides a convenience macro for the Variable Attribute:

```
#include <avr/pgmspace.h>
.
.
.
int mydata[] PROGMEM = ....
```

Note

The `PROGMEM` macro expands to the Variable Attribute of `progmem`. This macro requires that you include `<avr/pgmspace.h>`. This is the canonical method for defining a variable in Program Space.

To read back flash data, use the `pgm_read_*`() macros defined in `<avr/pgmspace.h>`. All Program Memory handling macros are defined there.

There is also a way to create a method to define variables in Program Memory that is common between the two compilers (IAR and AVR GCC). Create a header file that has these definitions:

```
#if defined(__ICCAVR__) // IAR C Compiler
#define FLASH_DECLARE(x) __flash x
#endif
#if defined(__GNUC__) // GNU Compiler
#define FLASH_DECLARE(x) x __attribute__((__progmem__))
#endif
```

This code snippet checks for the IAR compiler or for the GCC compiler and defines a macro `FLASH_DECLARE(x)` that will declare a variable in Program Memory using the appropriate method based on the compiler that is being used. Then you would use it like so:

```
FLASH_DECLARE(int mydata[] = ...);
```

10.6 Non-Returning main()

To declare `main()` to be a non-returning function in IAR, it is done like this:

```
__C_task void main(void)
{
    // code
}
```

To do the equivalent in AVR GCC, do this:

```
void main(void) __attribute__((noreturn));

void main(void)
{
    //...
}
```

Note

See the GCC User Manual for more information on Function Attributes.

In AVR GCC, a prototype for `main()` is required so you can declare the function attribute to specify that the `main()` function is of type "noreturn". Then, define `main()` as normal. Note that the return type for `main()` is now `void`.

10.7 Locking Registers

The IAR compiler allows a user to lock general registers from r15 and down by using compiler options and this keyword syntax:

```
__regvar __no_init volatile unsigned int filteredTimeSinceCommutation @14;
```

This line locks r14 for use only when explicitly referenced in your code through the var name "filteredTimeSinceCommutation". This means that the compiler cannot dispose of it at its own will.

To do this in AVR GCC, do this:

```
register unsigned char counter asm("r3");
```

Typically, it should be possible to use r2 through r15 that way.

Note

Do not reserve r0 or r1 as these are used internally by the compiler for a temporary register and for a zero value.

Locking registers is not recommended in AVR GCC as it removes this register from the control of the compiler, which may make code generation worse. Use at your own risk.

11 Frequently Asked Questions

11.1 FAQ Index

1. [My program doesn't recognize a variable updated within an interrupt routine](#)
2. [I get "undefined reference to..." for functions like "sin\(\)"](#)
3. [How to permanently bind a variable to a register?](#)

4. [How to modify MCUCR or WDTCR early?](#)
5. [What is all this _BV\(\) stuff about?](#)
6. [Can I use C++ on the AVR?](#)
7. [Shouldn't I initialize all my variables?](#)
8. [Why do some 16-bit timer registers sometimes get trashed?](#)
9. [How do I use a #define'd constant in an asm statement?](#)
10. [Why does the PC randomly jump around when single-stepping through my program in avr-gdb?](#)
11. [How do I trace an assembler file in avr-gdb?](#)
12. [How do I pass an IO port as a parameter to a function?](#)
13. [What registers are used by the C compiler?](#)
14. [How do I put an array of strings completely in ROM?](#)
15. [How to use external RAM?](#)
16. [Which -O flag to use?](#)
17. [How do I relocate code to a fixed address?](#)
18. [My UART is generating nonsense! My ATmega128 keeps crashing! Port F is completely broken!](#)
19. [Why do all my "foo...bar" strings eat up the SRAM?](#)
20. [Why does the compiler compile an 8-bit operation that uses bitwise operators into a 16-bit operation in assembly?](#)
21. [How to detect RAM memory and variable overlap problems?](#)
22. [Is it really impossible to program the ATtinyXX in C?](#)
23. [What is this "clock skew detected" message?](#)
24. [Why are \(many\) interrupt flags cleared by writing a logical 1?](#)
25. [Why have "programmed" fuses the bit value 0?](#)
26. [Which AVR-specific assembler operators are available?](#)
27. [Why are interrupts re-enabled in the middle of writing the stack pointer?](#)
28. [Why are there five different linker scripts?](#)
29. [How to add a raw binary image to linker output?](#)
30. [How do I perform a software reset of the AVR?](#)
31. [I am using floating point math. Why is the compiled code so big? Why does my code not work?](#)

- 32. [What pitfalls exist when writing reentrant code?](#)
- 33. [Why are some addresses of the EEPROM corrupted \(usually address zero\)?](#)
- 34. [Why is my baud rate wrong?](#)
- 35. [On a device with more than 128 KiB of flash, how to make function pointers work?](#)
- 36. [Why is assigning ports in a "chain" a bad idea?](#)

11.2 My program doesn't recognize a variable updated within an interrupt routine

When using the optimizer, in a loop like the following one:

```
uint8_t flag;
...
ISR(SOME_vect) {
    flag = 1;
}
...

while (flag == 0) {
    ...
}
```

the compiler will typically access `flag` only once, and optimize further accesses completely away, since its code path analysis shows that nothing inside the loop could change the value of `flag` anyway. To tell the compiler that this variable could be changed outside the scope of its code path analysis (e. g. from within an interrupt routine), the variable needs to be declared like:

```
volatile uint8_t flag;
```

Back to [FAQ Index](#).

11.3 I get "undefined reference to..." for functions like "sin()"

In order to access the mathematical functions that are declared in `<math.h>`, the linker needs to be told to also link the mathematical library, `libm.a`.

Typically, system libraries like `libm.a` are given to the final C compiler command line that performs the linking step by adding a flag `-lm` at the end. (That is, the initial *lib* and the filename suffix from the library are written immediately after a *-l* flag. So for a `libfoo.a` library, `-lfoo` needs to be provided.) This will make the linker search the library in a path known to the system.

An alternative would be to specify the full path to the `libm.a` file at the same place on the command line, i. e. *after* all the object files (`*.o`). However, since this requires knowledge of where the build system will exactly find those library files, this is deprecated for system libraries.

Back to [FAQ Index](#).

11.4 How to permanently bind a variable to a register?

This can be done with

```
register unsigned char counter asm("r3");
```

Typically, it should be safe to use r2 through r7 that way.

Registers r8 through r15 can be used for argument passing by the compiler in case many or long arguments are being passed to callees. If this is not the case throughout the entire application, these registers could be used for register variables as well.

Extreme care should be taken that the entire application is compiled with a consistent set of register-allocated variables, including possibly used library functions.

See [C Names Used in Assembler Code](#) for more details.

Back to [FAQ Index](#).

11.5 How to modify MCUCR or WDTCR early?

The method of early initialization (MCUCR, WDTCR or anything else) is different (and more flexible) in the current version. Basically, write a small assembler file which looks like this:

```
;; begin xram.S

#include <avr/io.h>

        .section .init1,"ax",@progbits

        ldi r16,_BV(SRE) | _BV(SRW)
        out _SFR_IO_ADDR(MCUCR),r16

;; end xram.S
```

Assemble it, link the resulting `xram.o` with other files in your program, and this piece of code will be inserted in initialization code, which is run right after reset. See the linker script for comments about the new `.initN` sections (which one to use, etc.).

The advantage of this method is that you can insert any initialization code you want (just remember that this is very early startup -- no stack and no `__zero_reg__` yet), and no program memory space is wasted if this feature is not used.

There should be no need to modify linker scripts anymore, except for some very special cases. It is best to leave `__stack` at its default value (end of internal SRAM -- faster, and required on some devices like ATmega161 because of errata), and add `-Wl,-Tdata,0x801100` to start the data section above the stack.

For more information on using sections, see [Memory Sections](#). There is also an example for [Using Sections in C Code](#). Note that in C code, any such function would preferably be placed into section `.init3` as the code in `.init2` ensures the internal register `__zero_reg__` is already cleared.

Back to [FAQ Index](#).

11.6 What is all this `_BV()` stuff about?

When performing low-level output work, which is a very central point in microcontroller programming, it is quite common that a particular bit needs to be set or cleared in some IO register. While the device documentation provides mnemonic names for the various bits in the IO registers, and the [AVR device-specific IO definitions](#) reflect these names in definitions for numerical constants, a way is needed to convert a bit number (usually within a byte register) into a byte value that can be assigned directly to the register. However, sometimes the direct bit numbers are needed as well (e. g. in an `SBI()` instruction), so the definitions cannot usefully be made as byte values in the first place.

So in order to access a particular bit number as a byte value, use the `_BV()` macro. Of course, the implementation of this macro is just the usual bit shift (which is done by the compiler anyway, thus doesn't impose any run-time penalty), so the following applies:

```
_BV(3) => 1 << 3 => 0x08
```

However, using the macro often makes the program better readable.

"BV" stands for "bit value", in case someone might ask you. :-)

Example: clock timer 2 with full IO clock (`CS2x = 0b001`), toggle OC2 output on compare match (`COM2x = 0b01`), and clear timer on compare match (`CTC2 = 1`). Make OC2 (`PD7`) an output.

```
TCCR2 = _BV(COM20) | _BV(CTC2) | _BV(CS20);  
DDRD = _BV(PD7);
```

Back to [FAQ Index](#).

11.7 Can I use C++ on the AVR?

Basically yes, C++ is supported (assuming your compiler has been configured and compiled to support it, of course). Source files ending in `.cc`, `.cpp` or `.C` will automatically cause the compiler frontend to invoke the C++ compiler. Alternatively, the C++ compiler could be explicitly called by the name `avr-c++`.

However, there's currently no support for `libstdc++`, the standard support library needed for a complete C++ implementation. This imposes a number of restrictions on the C++ programs that can be compiled. Among them are:

- Obviously, none of the C++ related standard functions, classes, and template classes are available.
- The operators `new` and `delete` are not implemented, attempting to use them will cause the linker to complain about undefined external references. (This could perhaps be fixed.)

- Some of the supplied include files are not C++ safe, i. e. they need to be wrapped into

```
extern "C" { . . . }
```

(This could certainly be fixed, too.)

- Exceptions are not supported. Since exceptions are enabled by default in the C++ frontend, they explicitly need to be turned off using `-fno-exceptions` in the compiler options. Failing this, the linker will complain about an undefined external reference to `__gxx_personality_sj0`.

Constructors and destructors *are* supported though, including global ones.

When programming C++ in space- and runtime-sensitive environments like microcontrollers, extra care should be taken to avoid unwanted side effects of the C++ calling conventions like implied copy constructors that could be called upon function invocation etc. These things could easily add up into a considerable amount of time and program memory wasted. Thus, casual inspection of the generated assembler code (using the `-S` compiler option) seems to be warranted.

Back to [FAQ Index](#).

11.8 Shouldn't I initialize all my variables?

Global and static variables are guaranteed to be initialized to 0 by the C standard. `avr-gcc` does this by placing the appropriate code into section `.init4` (see [The .initN Sections](#)). With respect to the standard, this sentence is somewhat simplified (because the standard allows for machines where the actual bit pattern used differs from all bits being 0), but for the AVR target, in general, all integer-type variables are set to 0, all pointers to a NULL pointer, and all floating-point variables to 0.0.

As long as these variables are not initialized (i. e. they don't have an equal sign and an initialization expression to the right within the definition of the variable), they go into the `.bss` section of the file. This section simply records the size of the variable, but otherwise doesn't consume space, neither within the object file nor within flash memory. (Of course, being a variable, it will consume space in the target's SRAM.)

In contrast, global and static variables that have an initializer go into the `.data` section of the file. This will cause them to consume space in the object file (in order to record the initializing value), *and* in the flash ROM of the target device. The latter is needed since the flash ROM is the only way that the compiler can tell the target device the value this variable is going to be initialized to.

Now if some programmer "wants to make doubly sure" their variables really get a 0 at program startup, and adds an initializer just containing 0 on the right-hand side, they waste space. While this waste of space applies to virtually any platform C is implemented on, it's usually not noticeable on larger machines like PCs, while the waste of flash ROM storage can be very painful on a small microcontroller like the AVR.

So in general, variables should only be explicitly initialized if the initial value is non-zero.

Note

Recent versions of GCC are now smart enough to detect this situation, and revert variables that are explicitly initialized to 0 to the .bss section. Still, other compilers might not do that optimization, and as the C standard guarantees the initialization, it is safe to rely on it.

Back to [FAQ Index](#).

11.9 Why do some 16-bit timer registers sometimes get trashed?

Some of the timer-related 16-bit IO registers use a temporary register (called TEMP in the Atmel datasheet) to guarantee an atomic access to the register despite the fact that two separate 8-bit IO transfers are required to actually move the data. Typically, this includes access to the current timer/counter value register (TCNT*n*), the input capture register (ICR*n*), and write access to the output compare registers (OCR*nM*). Refer to the actual datasheet for each device's set of registers that involves the TEMP register.

When accessing one of the registers that use TEMP from the main application, and possibly any other one from within an interrupt routine, care must be taken that no access from within an interrupt context could clobber the TEMP register data of an in-progress transaction that has just started elsewhere.

To protect interrupt routines against other interrupt routines, it's usually best to use the [ISR\(\)](#) macro when declaring the interrupt function, and to ensure that interrupts are still disabled when accessing those 16-bit timer registers.

Within the main program, access to those registers could be encapsulated in calls to the [cli\(\)](#) and [sei\(\)](#) macros. If the status of the global interrupt flag before accessing one of those registers is uncertain, something like the following example code can be used.

```
uint16_t
read_timer1(void)
{
    uint8_t sreg;
    uint16_t val;

    sreg = SREG;
    cli();
    val = TCNT1;
    SREG = sreg;

    return val;
}
```

Back to [FAQ Index](#).

11.10 How do I use a #define'd constant in an asm statement?

So you tried this:

```
asm volatile("sbi 0x18,0x07;");
```

11.11 Why does the PC randomly jump around when single-stepping through my program in avr-gdb?

66

Which works. When you do the same thing but replace the address of the port by its macro name, like this:

```
asm volatile("sbi PORTB,0x07;");
```

you get a compilation error: "Error: constant value required".

PORTB is a precompiler definition included in the processor specific file included in [avr/io.h](#). As you may know, the precompiler will not touch strings and PORTB, instead of 0x18, gets passed to the assembler. One way to avoid this problem is:

```
asm volatile("sbi %0, 0x07" : "I" (_SFR_IO_ADDR(PORTB)));
```

Note

For C programs, rather use the standard C bit operators instead, so the above would be expressed as `PORTB |= (1 << 7)`. The optimizer will take care to transform this into a single SBI instruction, assuming the operands allow for this.

Back to [FAQ Index](#).

11.11 Why does the PC randomly jump around when single-stepping through my program in avr-gdb?

When compiling a program with both optimization (`-O`) and debug information (`-g`) which is fortunately possible in `avr-gcc`, the code watched in the debugger is optimized code. While it is not guaranteed, very often this code runs with the exact same optimizations as it would run without the `-g` switch.

This can have unwanted side effects. Since the compiler is free to reorder code execution as long as the semantics do not change, code is often rearranged in order to make it possible to use a single branch instruction for conditional operations. Branch instructions can only cover a short range for the target PC (-63 through +64 words from the current PC). If a branch instruction cannot be used directly, the compiler needs to work around it by combining a skip instruction together with a relative jump (`rjmp`) instruction, which will need one additional word of ROM.

Another side effect of optimization is that variable usage is restricted to the area of code where it is actually used. So if a variable was placed in a register at the beginning of some function, this same register can be re-used later on if the compiler notices that the first variable is no longer used inside that function, even though the variable is still in lexical scope. When trying to examine the variable in `avr-gdb`, the displayed result will then look garbled.

So in order to avoid these side effects, optimization can be turned off while debugging. However, some of these optimizations might also have the side effect of uncovering bugs that would otherwise not be obvious, so it must be noted that turning off optimization can easily change the bug pattern. In most cases, you are better off leaving optimizations enabled while debugging.

Back to [FAQ Index](#).

11.12 How do I trace an assembler file in avr-gdb?

When using the `-g` compiler option, `avr-gcc` only generates line number and other debug information for C (and C++) files that pass the compiler. Functions that don't have line number information will be completely skipped by a single `step` command in `gdb`. This includes functions linked from a standard library, but by default also functions defined in an assembler source file, since the `-g` compiler switch does not apply to the assembler.

So in order to debug an assembler input file (possibly one that has to be passed through the C preprocessor), it's the assembler that needs to be told to include line-number information into the output file. (Other debug information like data types and variable allocation cannot be generated, since unlike a compiler, the assembler basically doesn't know about this.) This is done using the (GNU) assembler option `--gstabs`.

Example:

```
$ avr-as -mmcu=atmega128 --gstabs -o foo.o foo.s
```

When the assembler is not called directly but through the C compiler frontend (either implicitly by passing a source file ending in `.S`, or explicitly using `-x assembler-with-cpp`), the compiler frontend needs to be told to pass the `--gstabs` option down to the assembler. This is done using `-Wa,--gstabs`. Please take care to *only* pass this option when compiling an assembler input file. Otherwise, the assembler code that results from the C compilation stage will also get line number information, which confuses the debugger.

Note

You can also use `-Wa,--gstabs` since the compiler will add the extra `'-'` for you.

Example:

```
$ EXTRA_OPTS="-Wall -mmcu=atmega128 -x assembler-with-cpp"
$ avr-gcc -Wa,--gstabs ${EXTRA_OPTS} -c -o foo.o foo.S
```

Also note that the debugger might get confused when entering a piece of code that has a non-local label before, since it then takes this label as the name of a new function that appears to have been entered. Thus, the best practice to avoid this confusion is to only use non-local labels when declaring a new function, and restrict anything else to local labels. Local labels consist just of a number only. References to these labels consist of the number, followed by the letter **b** for a backward reference, or **f** for a forward reference. These local labels may be re-used within the source file, references will pick the closest label with the same number and given direction.

Example:

```
myfunc: push    r16
        push    r17
        push    r18
```

```

    push    YL
    push    YH
    ...
    eor r16, r16    ; start loop
    ldi YL, lo8(sometable)
    ldi YH, hi8(sometable)
    rjmp    2f      ; jump to loop test at end
1: ld r17, Y+      ; loop continues here
    ...
    breq    1f      ; return from myfunc prematurely
    ...
    inc r16
2: cmp r16, r18
    brlo    1b      ; jump back to top of loop

1: pop YH
   pop YL
   pop r18
   pop r17
   pop r16
   ret

```

Back to [FAQ Index](#).

11.13 How do I pass an IO port as a parameter to a function?

Consider this example code:

```

#include <inttypes.h>
#include <avr/io.h>

void
set_bits_func_wrong (volatile uint8_t port, uint8_t mask)
{
    port |= mask;
}

void
set_bits_func_correct (volatile uint8_t *port, uint8_t mask)
{
    *port |= mask;
}

#define set_bits_macro(port,mask) ((port) |= (mask))

int main (void)
{
    set_bits_func_wrong (PORTB, 0xaa);
    set_bits_func_correct (&PORTB, 0x55);
    set_bits_macro (PORTB, 0xf0);

    return (0);
}

```

The first function will generate object code which is not even close to what is intended. The major problem arises when the function is called. When the compiler sees this call, it will actually pass the value of the PORTB register (using an IN instruction), instead

of passing the address of PORTB (e.g. memory mapped io addr of 0x38, io port 0x18 for the mega128). This is seen clearly when looking at the disassembly of the call:

```

    set_bits_func_wrong (PORTB, 0xaa);
10a:  6a ea          ldi     r22, 0xAA          ; 170
10c:  88 b3          in      r24, 0x18        ; 24
10e:  0e 94 65 00    call    0xca

```

So, the function, once called, only sees the value of the port register and knows nothing about which port it came from. At this point, whatever object code is generated for the function by the compiler is irrelevant. The interested reader can examine the full disassembly to see that the function's body is completely fubar.

The second function shows how to pass (by reference) the memory mapped address of the io port to the function so that you can read and write to it in the function. Here's the object code generated for the function call:

```

    set_bits_func_correct (&PORTB, 0x55);
112:  65 e5          ldi     r22, 0x55          ; 85
114:  88 e3          ldi     r24, 0x38          ; 56
116:  90 e0          ldi     r25, 0x00          ; 0
118:  0e 94 7c 00    call    0xf8

```

You can clearly see that 0x0038 is correctly passed for the address of the io port. Looking at the disassembled object code for the body of the function, we can see that the function is indeed performing the operation we intended:

```

void
set_bits_func_correct (volatile uint8_t *port, uint8_t mask)
{
    f8:  fc 01          movw    r30, r24
        *port |= mask;
    fa:  80 81          ld      r24, Z
    fc:  86 2b          or      r24, r22
    fe:  80 83          st      Z, r24
}
100:  08 95          ret

```

Notice that we are accessing the io port via the LD and ST instructions.

The `port` parameter must be volatile to avoid a compiler warning.

Note

Because of the nature of the IN and OUT assembly instructions, they can not be used inside the function when passing the port in this way. Readers interested in the details should consult the *Instruction Set* datasheet.

Finally we come to the macro version of the operation. In this contrived example, the macro is the most efficient method with respect to both execution speed and code size:

```

    set_bits_macro (PORTB, 0xf0);
11c:  88 b3          in      r24, 0x18        ; 24
11e:  80 6f          ori     r24, 0xF0          ; 240
120:  88 bb          out     0x18, r24          ; 24

```

Of course, in a real application, you might be doing a lot more in your function which uses a passed by reference io port address and thus the use of a function over a macro could save you some code space, but still at a cost of execution speed.

Care should be taken when such an indirect port access is going to one of the 16-bit IO registers where the order of write access is critical (like some timer registers). All versions of avr-gcc up to 3.3 will generate instructions that use the wrong access order in this situation (since with normal memory operands where the order doesn't matter, this sometimes yields shorter code).

See <http://mail.nongnu.org/archive/html/avr-libc-dev/2003-01/msg00044.html> for a possible workaround.

avr-gcc versions after 3.3 have been fixed in a way where this optimization will be disabled if the respective pointer variable is declared to be `volatile`, so the correct behaviour for 16-bit IO ports can be forced that way.

Back to [FAQ Index](#).

11.14 What registers are used by the C compiler?

- **Data types:**

`char` is 8 bits, `int` is 16 bits, `long` is 32 bits, `long long` is 64 bits, `float` and `double` are 32 bits (this is the only supported floating point format), pointers are 16 bits (function pointers are word addresses, to allow addressing up to 128K program memory space). There is a `-mint8` option (see [Options for the C compiler avr-gcc](#)) to make `int` 8 bits, but that is not supported by avr-libc and violates C standards (`int` *must* be at least 16 bits). It may be removed in a future release.

- **Call-used registers (r18-r27, r30-r31):**

May be allocated by gcc for local data. You *may* use them freely in assembler subroutines. Calling C subroutines can clobber any of them - the caller is responsible for saving and restoring.

- **Call-saved registers (r2-r17, r28-r29):**

May be allocated by gcc for local data. Calling C subroutines leaves them unchanged. Assembler subroutines are responsible for saving and restoring these registers, if changed. `r29:r28` (Y pointer) is used as a frame pointer (points to local data on stack) if necessary. The requirement for the callee to save/preserve the contents of these registers even applies in situations where the compiler assigns them for argument passing.

- **Fixed registers (r0, r1):**

Never allocated by gcc for local data, but often used for fixed purposes:

`r0` - temporary register, can be clobbered by any C code (except interrupt handlers which save it), *may* be used to remember something for a while within one piece of assembler code

`r1` - assumed to be always zero in any C code, *may* be used to remember something for a while within one piece of assembler code, but *must* then be cleared after use (`clr r1`). This includes any use of the `[f]mul[s[u]]` instructions, which return their result in `r1:r0`. Interrupt handlers save and clear `r1` on entry, and restore `r1` on exit (in case it was non-zero).

- **Function call conventions:**

Arguments - allocated left to right, `r25` to `r8`. All arguments are aligned to start in even-numbered registers (odd-sized arguments, including `char`, have one free register above them). This allows making better use of the `movw` instruction on the enhanced core.

If too many, those that don't fit are passed on the stack.

Return values: 8-bit in `r24` (not `r25`!), 16-bit in `r25:r24`, up to 32 bits in `r22-r25`, up to 64 bits in `r18-r25`. 8-bit return values are zero/sign-extended to 16 bits by the called function (unsigned `char` is more efficient than signed `char` - just `clr r25`). Arguments to functions with variable argument lists (`printf` etc.) are all passed on stack, and `char` is extended to `int`.

Warning

There was no such alignment before 2000-07-01, including the old patches for `gcc-2.95.2`. Check your old assembler subroutines, and adjust them accordingly.

Back to [FAQ Index](#).

11.15 How do I put an array of strings completely in ROM?

There are times when you may need an array of strings which will never be modified. In this case, you don't want to waste ram storing the constant strings. The most obvious (and incorrect) thing to do is this:

```
#include <avr/pgmspace.h>

PGM_P array[2] PROGMEM = {
    "Foo",
    "Bar"
};

int main (void)
{
    char buf[32];
    strcpy_P (buf, array[1]);
    return 0;
}
```

The result is not what you want though. What you end up with is the array stored in ROM, while the individual strings end up in RAM (in the `.data` section).

To work around this, you need to do something like this:

```
#include <avr/pgmspace.h>

const char foo[] PROGMEM = "Foo";
const char bar[] PROGMEM = "Bar";

PGM_P array[2] PROGMEM = {
    foo,
    bar
};

int main (void)
{
    char buf[32];
    PGM_P p;
    int i;

    memcpy_P(&p, &array[i], sizeof(PGM_P));
    strcpy_P(buf, p);
    return 0;
}
```

Looking at the disassembly of the resulting object file we see that array is in flash as such:

```
00000026 <array>:
 26:  2e 00          .word  0x002e  ; ????
 28:  2a 00          .word  0x002a  ; ????

0000002a <bar>:
 2a:  42 61 72 00          Bar.

0000002e <foo>:
 2e:  46 6f 6f 00          Foo.
```

foo is at addr 0x002e.

bar is at addr 0x002a.

array is at addr 0x0026.

Then in main we see this:

```
    memcpy_P(&p, &array[i], sizeof(PGM_P));
70:  66 0f          add     r22, r22
72:  77 1f          adc     r23, r23
74:  6a 5d          subi   r22, 0xDA      ; 218
76:  7f 4f          sbci   r23, 0xFF      ; 255
78:  42 e0          ldi     r20, 0x02      ; 2
7a:  50 e0          ldi     r21, 0x00      ; 0
7c:  ce 01          movw   r24, r28
7e:  81 96          adiw   r24, 0x21      ; 33
80:  08 d0          rcall  .+16           ; 0x92
```

This code reads the pointer to the desired string from the ROM table array into a register pair.

The value of `i` (in `r22:r23`) is doubled to accommodate for the word offset required to access `array[]`, then the address of array (0x26) is added, by subtracting the negated address (0xffda). The address of variable `p` is computed by adding its offset within the stack frame (33) to the Y pointer register, and `memcpy_P` is called.

```

    strcpy_P(buf, p);
82:  69 a1          ldd     r22, Y+33          ; 0x21
84:  7a a1          ldd     r23, Y+34          ; 0x22
86:  ce 01          movw   r24, r28
88:  01 96          adiw    r24, 0x01          ; 1
8a:  0c d0          rcall   .+24              ; 0xa4

```

This will finally copy the ROM string into the local buffer `buf`.

Variable `p` (located at `Y+33`) is read, and passed together with the address of `buf` (`Y+1`) to `strcpy_P`. This will copy the string from ROM to `buf`.

Note that when using a compile-time constant index, omitting the first step (reading the pointer from ROM via `memcpy_P`) usually remains unnoticed, since the compiler would then optimize the code for accessing `array` at compile-time.

Back to [FAQ Index](#).

11.16 How to use external RAM?

Well, there is no universal answer to this question; it depends on what the external RAM is going to be used for.

Basically, the bit `SRE` (SRAM enable) in the `MCUCR` register needs to be set in order to enable the external memory interface. Depending on the device to be used, and the application details, further registers affecting the external memory operation like `XMCRA` and `XMCRB`, and/or further bits in `MCUCR` might be configured. Refer to the datasheet for details.

If the external RAM is going to be used to store the variables from the C program (i. e., the `.data` and/or `.bss` segment) in that memory area, it is essential to set up the external memory interface early during the [device initialization](#) so the initialization of these variable will take place. Refer to [How to modify MCUCR or WDTCR early?](#) for a description how to do this using few lines of assembler code, or to the chapter about memory sections for an [example written in C](#).

The explanation of [malloc\(\)](#) contains a [discussion](#) about the use of internal RAM vs. external RAM in particular with respect to the various possible locations of the *heap* (area reserved for [malloc\(\)](#)). It also explains the linker command-line options that are required to move the memory regions away from their respective standard locations in internal RAM.

Finally, if the application simply wants to use the additional RAM for private data storage kept outside the domain of the C compiler (e. g. through a `char *` variable initialized directly to a particular address), it would be sufficient to defer the initialization of the external RAM interface to the beginning of `main()`, so no tweaking of the `.init3` section is necessary. The same applies if only the heap is going to be located there, since the application start-up code does not affect the heap.

It is not recommended to locate the stack in external RAM. In general, accessing external RAM is slower than internal RAM, and errata of some AVR devices even prevent this configuration from working properly at all.

Back to [FAQ Index](#).

11.17 Which -O flag to use?

There's a common misconception that larger numbers behind the `-O` option might automatically cause "better" optimization. First, there's no universal definition for "better", with optimization often being a speed vs. code size trade off. See the [detailed discussion](#) for which option affects which part of the code generation.

A test case was run on an ATmega128 to judge the effect of compiling the library itself using different optimization levels. The following table lists the results. The test case consisted of around 2 KB of strings to sort. Test #1 used `qsort()` using the standard library `strcmp()`, test #2 used a function that sorted the strings by their size (thus had two calls to `strlen()` per invocation).

When comparing the resulting code size, it should be noted that a floating point version of `fvprintf()` was linked into the binary (in order to print out the time elapsed) which is entirely not affected by the different optimization levels, and added about 2.5 KB to the code.

Optimization flags	Size of .text	Time for test #1	Time for test #2
-O3	6898	903 μ s	19.7 ms
-O2	6666	972 μ s	20.1 ms
-Os	6618	955 μ s	20.1 ms
-Os -mcall-prologues	6474	972 μ s	20.1 ms

(The difference between 955 μ s and 972 μ s was just a single timer-tick, so take this with a grain of salt.)

So generally, it seems `-Os -mcall-prologues` is the most universal "best" optimization level. Only applications that need to get the last few percent of speed benefit from using `-O3`.

Back to [FAQ Index](#).

11.18 How do I relocate code to a fixed address?

First, the code should be put into a new [named section](#). This is done with a section attribute:

```
__attribute__((section (".bootloader")))
```

In this example, `.bootloader` is the name of the new section. This attribute needs to be placed after the prototype of any function to force the function into the new section.

```
void boot(void) __attribute__((section (".bootloader")));
```

To relocate the section to a fixed address the linker flag `--section-start` is used. This option can be passed to the linker using the [-Wl compiler option](#):

```
-Wl,--section-start=.bootloader=0x1E000
```

The name after section-start is the name of the section to be relocated. The number after the section name is the beginning address of the named section.

Back to [FAQ Index](#).

11.19 My UART is generating nonsense! My ATmega128 keeps crashing! Port F is completely broken!

Well, certain odd problems arise out of the situation that the AVR devices as shipped by Atmel often come with a default fuse bit configuration that doesn't match the user's expectations. Here is a list of things to care for:

- All devices that have an internal RC oscillator ship with the fuse enabled that causes the device to run off this oscillator, instead of an external crystal. This often remains unnoticed until the first attempt is made to use something critical in timing, like UART communication.
- The ATmega128 ships with the fuse enabled that turns this device into ATmega103 compatibility mode. This means that some ports are not fully usable, and in particular that the internal SRAM is located at lower addresses. Since by default, the stack is located at the top of internal SRAM, a program compiled for an ATmega128 running on such a device will immediately crash upon the first function call (or rather, upon the first function return).
- Devices with a JTAG interface have the JTAGEN fuse programmed by default. This will make the respective port pins that are used for the JTAG interface unavailable for regular IO.

Back to [FAQ Index](#).

11.20 Why do all my "foo...bar" strings eat up the SRAM?

By default, all strings are handled as all other initialized variables: they occupy RAM (even though the compiler might warn you when it detects write attempts to these RAM locations), and occupy the same amount of flash ROM so they can be initialized to the actual string by startup code. The compiler can optimize multiple identical strings into a single one, but obviously only for one compilation unit (i. e., a single C source file).

That way, any string literal will be a valid argument to any C function that expects a `const char *` argument.

Of course, this is going to waste a lot of SRAM. In [Program Space String Utilities](#), a method is described how such constant data can be moved out to flash ROM. However, a constant string located in flash ROM is no longer a valid argument to pass to a function that expects a `const char *`-type string, since the AVR processor needs the special instruction `LPM` to access these strings. Thus, separate functions are needed that take this into account. Many of the standard C library functions have equivalents available where one of the string arguments can be located in flash ROM. Private functions in the applications need to handle this, too. For example, the following can be used to implement simple debugging messages that will be sent through a UART:

11.21 Why does the compiler compile an 8-bit operation that uses bitwise operators into a 16-bit operation in assembly?

76

```
#include <inttypes.h>
#include <avr/io.h>
#include <avr/pgmspace.h>

int
uart_putchar(char c)
{
    if (c == '\n')
        uart_putchar('\r');
    loop_until_bit_is_set(USR, UDRE);
    UDR = c;
    return 0; /* so it could be used for fdevopen(), too */
}

void
debug_P(const char *addr)
{
    char c;

    while ((c = pgm_read_byte(addr++)))
        uart_putchar(c);
}

int
main(void)
{
    ioinit(); /* initialize UART, ... */
    debug_P(PSTR("foo was here\n"));
    return 0;
}
```

Note

By convention, the suffix **_P** to the function name is used as an indication that this function is going to accept a "program-space string". Note also the use of the [PSTR\(\)](#) macro.

Back to [FAQ Index](#).

11.21 Why does the compiler compile an 8-bit operation that uses bitwise operators into a 16-bit operation in assembly?

Bitwise operations in Standard C will automatically promote their operands to an int, which is (by default) 16 bits in avr-gcc.

To work around this use typecasts on the operands, including literals, to declare that the values are to be 8 bit operands.

This may be especially important when clearing a bit:

```
var &= ~mask; /* wrong way! */
```

The bitwise "not" operator (~) will also promote the value in `mask` to an int. To keep it an 8-bit value, typecast before the "not" operator:

```
var &= (unsigned char)~mask;
```


Back to [FAQ Index](#).

11.22 How to detect RAM memory and variable overlap problems?

You can simply run `avr-nm` on your output (ELF) file. Run it with the `-n` option, and it will sort the symbols numerically (by default, they are sorted alphabetically).

Look for the symbol `_end`, that's the first address in RAM that is not allocated by a variable. (`avr-gcc` internally adds `0x800000` to all data/bss variable addresses, so please ignore this offset.) Then, the run-time initialization code initializes the stack pointer (by default) to point to the last available address in (internal) SRAM. Thus, the region between `_end` and the end of SRAM is what is available for stack. (If your application uses `malloc()`, which e. g. also can happen inside `printf()`, the heap for dynamic memory is also located there. See [Memory Areas and Using malloc\(\)](#).)

The amount of stack required for your application cannot be determined that easily. For example, if you recursively call a function and forget to break that recursion, the amount of stack required is infinite. :-) You can look at the generated assembler code (`avr-gcc ... -S`), there's a comment in each generated assembler file that tells you the frame size for each generated function. That's the amount of stack required for this function, you have to add up that for all functions where you know that the calls could be nested.

Back to [FAQ Index](#).

11.23 Is it really impossible to program the ATtinyXX in C?

While some small AVR's are not directly supported by the C compiler since they do not have a RAM-based stack (and some do not even have RAM at all), it is possible anyway to use the general-purpose registers as a RAM replacement since they are mapped into the data memory region.

Bruce D. Lightner wrote an excellent description of how to do this, and offers this together with a toolkit on his web page:

<http://lightner.net/avr/ATtinyAvrGcc.html>

Back to [FAQ Index](#).

11.24 What is this "clock skew detected" message?

It's a known problem of the MS-DOS FAT file system. Since the FAT file system has only a granularity of 2 seconds for maintaining a file's timestamp, and it seems that some MS-DOS derivative (Win9x) perhaps rounds up the current time to the next second when calculating the timestamp of an updated file in case the current time cannot be represented in FAT's terms, this causes a situation where `make` sees a "file coming from the future".

Since all `make` decisions are based on file timestamps, and their dependencies, `make` warns about this situation.

Solution: don't use inferior file systems / operating systems. Neither Unix file systems nor HPFS (aka NTFS) do experience that problem.

Workaround: after saving the file, wait a second before starting `make`. Or simply ignore the warning. If you are paranoid, execute a `make clean all` to make sure everything gets rebuilt.

In networked environments where the files are accessed from a file server, this message can also happen if the file server's clock differs too much from the network client's clock. In this case, the solution is to use a proper time keeping protocol on both systems, like NTP. As a workaround, synchronize the client's clock frequently with the server's clock.

Back to [FAQ Index](#).

11.25 Why are (many) interrupt flags cleared by writing a logical 1?

Usually, each interrupt has its own interrupt flag bit in some control register, indicating the specified interrupt condition has been met by representing a logical 1 in the respective bit position. When working with interrupt handlers, this interrupt flag bit usually gets cleared automatically in the course of processing the interrupt, sometimes by just calling the handler at all, sometimes (e. g. for the U[S]ART) by reading a particular hardware register that will normally happen anyway when processing the interrupt.

From the hardware's point of view, an interrupt is asserted as long as the respective bit is set, while global interrupts are enabled. Thus, it is essential to have the bit cleared before interrupts get re-enabled again (which usually happens when returning from an interrupt handler).

Only few subsystems require an explicit action to clear the interrupt request when using interrupt handlers. (The notable exception is the TWI interface, where clearing the interrupt indicates to proceed with the TWI bus hardware handshake, so it's never done automatically.)

However, if no normal interrupt handlers are to be used, or in order to make extra sure any pending interrupt gets cleared before re-activating global interrupts (e. g. an external edge-triggered one), it can be necessary to explicitly clear the respective hardware interrupt bit by software. This is usually done by writing a logical 1 into this bit position. This seems to be illogical at first, the bit position already carries a logical 1 when reading it, so why does writing a logical 1 to it *clear* the interrupt bit?

The solution is simple: writing a logical 1 to it requires only a single `OUT` instruction, and it is clear that only this single interrupt request bit will be cleared. There is no need to perform a read-modify-write cycle (like, an `SBI` instruction), since all bits in these control registers are interrupt bits, and writing a logical 0 to the remaining bits (as it is done by the simple `OUT` instruction) will not alter them, so there is no risk of any race condition that might accidentally clear another interrupt request bit. So instead of writing

```
TIFR |= _BV(TOV0); /* wrong! */
```

simply use

```
TIFR = _BV(TOV0);
```

Back to [FAQ Index](#).

11.26 Why have "programmed" fuses the bit value 0?

Basically, fuses are just a bit in a special EEPROM area. For technical reasons, erased E[EEPROM] cells have all bits set to the value 1, so unprogrammed fuses also have a logical 1. Conversely, programmed fuse cells read out as bit value 0.

Back to [FAQ Index](#).

11.27 Which AVR-specific assembler operators are available?

See [Pseudo-ops and operators](#).

Back to [FAQ Index](#).

11.28 Why are interrupts re-enabled in the middle of writing the stack pointer?

When setting up space for local variables on the stack, the compiler generates code like this:

```
/* prologue: frame size=20 */
    push r28
    push r29
    in r28,__SP_L__
    in r29,__SP_H__
    sbiw r28,20
    in __tmp_reg__,__SREG__
    cli
    out __SP_H__,r29
    out __SREG__,__tmp_reg__
    out __SP_L__,r28
/* prologue end (size=10) */
```

It reads the current stack pointer value, decrements it by the required amount of bytes, then disables interrupts, writes back the high part of the stack pointer, writes back the saved SREG (which will eventually re-enable interrupts if they have been enabled before), and finally writes the low part of the stack pointer.

At the first glance, there's a race between restoring SREG, and writing SPL. However, after enabling interrupts (either explicitly by setting the I flag, or by restoring it as part of the entire SREG), the AVR hardware executes (at least) the next instruction still with interrupts disabled, so the write to SPL is guaranteed to be executed with interrupts disabled still. Thus, the emitted sequence ensures interrupts will be disabled only for the minimum time required to guarantee the integrity of this operation.

Back to [FAQ Index](#).

11.29 Why are there five different linker scripts?

From a comment in the source code:

Which one of the five linker script files is actually used depends on command line options given to ld.

A .x script file is the default script A .xr script is for linking without relocation (-r flag) A .xu script is like .xr but *do* create constructors (-Ur flag) A .xn script is for linking with -n flag (mix text and data on same page). A .xbn script is for linking with -N flag (mix text and data on same page).

Back to [FAQ Index](#).

11.30 How to add a raw binary image to linker output?

The GNU linker `avr-ld` cannot handle binary data directly. However, there's a companion tool called `avr-objcopy`. This is already known from the output side: it's used to extract the contents of the linked ELF file into an Intel Hex load file.

`avr-objcopy` can create a relocatable object file from arbitrary binary input, like

```
avr-objcopy -I binary -O elf32-avr foo.bin foo.o
```

This will create a file named `foo.o`, with the contents of `foo.bin`. The contents will default to section `.data`, and two symbols will be created named `_binary_foo_bin_start` and `_binary_foo_bin_end`. These symbols can be referred to inside a C source to access these data.

If the goal is to have those data go to flash ROM (similar to having used the `PROGMEM` attribute in C source code), the sections have to be renamed while copying, and it's also useful to set the section flags:

```
avr-objcopy --rename-section .data=.progmem.data,contents,alloc,load,readonly,data \
a -I binary -O elf32-avr foo.bin foo.o
```

Note that all this could be conveniently wired into a Makefile, so whenever `foo.bin` changes, it will trigger the recreation of `foo.o`, and a subsequent relink of the final ELF file.

Below are two Makefile fragments that provide rules to convert a .txt file to an object file, and to convert a .bin file to an object file:

```
$(OBJDIR)/%.o : %.txt
    @echo Converting $<
    @cp $(<) $(*) .tmp
    @echo -n 0 | tr 0 '\000' >> $(*) .tmp
    @$(OBJCOPY) -I binary -O elf32-avr \
    --rename-section .data=.progmem.data,contents,alloc,load,readonly,data \
    --redefine-sym _binary_${*_tmp}_start=${*_tmp}_start \
    --redefine-sym _binary_${*_tmp}_end=${*_tmp}_end \
    --redefine-sym _binary_${*_tmp}_size=${*_tmp}_size_sym \
    $(*) .tmp $(@)
    @echo "extern const char" $(*) "[" PROGMEM;" > $(*) .h
```

```

@echo "extern const char" $(*)_end"[] PROGMEM;" >> $(*).h
@echo "extern const char" $(*)_size_sym"[];" >> $(*).h
@echo "#define $(*)_size ((int)$(*_size_sym)" >> $(*).h
@rm $(*).tmp

$(OBJDIR)/%.o : %.bin
@echo Converting $<
@$(OBJCOPY) -I binary -O elf32-avr \
--rename-section .data=.progmem.data,contents,alloc,load,readonly,data \
--redefine-sym _binary_${*_bin_start}=${*_ \
--redefine-sym _binary_${*_bin_end}=${*_end \
--redefine-sym _binary_${*_bin_size}=${*_size_sym \
$(<) $(@)
@echo "extern const char" $(*)"[] PROGMEM;" > $(*).h
@echo "extern const char" $(*)_end"[] PROGMEM;" >> $(*).h
@echo "extern const char" $(*)_size_sym"[];" >> $(*).h
@echo "#define $(*)_size ((int)$(*_size_sym)" >> $(*).h

```

Back to [FAQ Index](#).

11.31 How do I perform a software reset of the AVR?

The canonical way to perform a software reset of non-XMega AVR's is to use the watchdog timer. Enable the watchdog timer to the shortest timeout setting, then go into an infinite, do-nothing loop. The watchdog will then reset the processor.

XMega parts have a specific bit `RST_SWRST_bm` in the `RST_CTRL` register, that generates a hardware reset. `RST_SWRST_bm` is protected by the XMega Configuration Change Protection system.

The reason why using the watchdog timer or `RST_SWRST_bm` is preferable over jumping to the reset vector, is that when the watchdog or `RST_SWRST_bm` resets the AVR, the registers will be reset to their known, default settings. Whereas jumping to the reset vector will leave the registers in their previous state, which is generally not a good idea.

CAUTION! Older AVR's will have the watchdog timer disabled on a reset. For these older AVR's, doing a soft reset by enabling the watchdog is easy, as the watchdog will then be disabled after the reset. On newer AVR's, once the watchdog is enabled, then it **stays enabled, even after a reset!** For these newer AVR's a function needs to be added to the `.init3` section (i.e. during the startup code, before `main()`) to disable the watchdog early enough so it does not continually reset the AVR.

Here is some example code that creates a macro that can be called to perform a soft reset:

```

#include <avr/wdt.h>

...

#define soft_reset() \
do \
{ \
    wdt_enable(WDTO_15MS); \
    for(;;) \
    { \
    } \
} while(0)

```

11.32 I am using floating point math. Why is the compiled code so big? Why does my code not work? 82

For newer AVR's (such as the ATmega1281) also add this function to your code to then disable the watchdog after a reset (e.g., after a soft reset):

```
#include <avr/wdt.h>

...

// Function Pototype
void wdt_init(void) __attribute__((naked)) __attribute__((section(".init3")));

...

// Function Implementation
void wdt_init(void)
{
    MCUSR = 0;
    wdt_disable();

    return;
}
```

Back to [FAQ Index](#).

11.32 I am using floating point math. Why is the compiled code so big? Why does my code not work?

You are not linking in the math library from AVR-LibC. GCC has a library that is used for floating point operations, but it is not optimized for the AVR, and so it generates big code, or it could be incorrect. This can happen even when you are not using any floating point math functions from the Standard C library, but you are just doing floating point math operations.

When you link in the math library from AVR-LibC, those routines get replaced by hand-optimized AVR assembly and it produces much smaller code.

See [I get "undefined reference to..." for functions like "sin\(\)"](#) for more details on how to link in the math library.

Back to [FAQ Index](#).

11.33 What pitfalls exist when writing reentrant code?

Reentrant code means the ability for a piece of code to be called simultaneously from two or more threads. Attention to re-entrability is needed when using a multi-tasking operating system, or when using interrupts since an interrupt is really a temporary thread.

The code generated natively by gcc is reentrant. But, only some of the libraries in avr-libc are explicitly reentrant, and some are known not to be reentrant. In general, any library call that reads and writes global variables (including I/O registers) is not reentrant. This is because more than one thread could read or write the same storage at the same time, unaware that other threads are doing the same, and create inconsistent and/or erroneous results.

A library call that is known not to be reentrant will work if it is used only within one thread *and* no other thread makes use of a library call that shares common storage with it.

Below is a table of library calls with known issues.

Library call	Reentrant Issue	Workaround/Alternative
rand() , random()	Uses global variables to keep state information.	Use special reentrant versions: rand_r() , random_r() .
strtod() , strtol() , strtoul()	Uses the global variable <code>errno</code> to return success/failure.	Ignore <code>errno</code> , or protect calls with cli()/sei() or ATOMIC_BLOCK() if the application can tolerate it. Or use sccanf() or sccanf_P() if possible.
malloc() , realloc() , calloc() , free()	Uses the stack pointer and global variables to allocate and free memory.	Protect calls with cli()/sei() or ATOMIC_BLOCK() if the application can tolerate it. If using an OS, use the OS provided memory allocator since the OS is likely modifying the stack pointer anyway.
fdevopen() , fclose()	Uses calloc() and free() .	Protect calls with cli()/sei() or ATOMIC_BLOCK() if the application can tolerate it. Or use fdev_setup_stream() or FDEV_SETUP_STREAM() . Note: fclose() will only call free() if the stream has been opened with fdevopen() .
eeprom_* (), boot_* ()	Accesses I/O registers.	Protect calls with cli()/sei() , ATOMIC_BLOCK() , or use OS locking.
pgm_*_far()	Accesses I/O register <code>RAMPZ</code> .	Starting with GCC 4.3, <code>RAMPZ</code> is automatically saved for ISRs, so nothing further is needed if only using interrupts. Some OSes may automatically preserve <code>RAMPZ</code> during context switching. Check the OS documentation before assuming it does. Otherwise, protect calls with cli()/sei() , ATOMIC_BLOCK() , or use explicit OS locking.
Generated on Wed May 11 11:27:14 2022 for avr-libc by Doxygen		
printf() , printf_P() , vprintf() , vprintf_P() , puts() , puts_P()	Alters flags and character count in global <code>FILE stdout</code> .	Use only in one thread. Or if returned character count is unimportant, do not use the <code>*_P</code> versions. Note: Formatting to a string output, e.g. sprintf() , sprintf_P() ,

Note

It's not clear one would ever want to do character input simultaneously from more than one thread anyway, but these entries are included for completeness.

An effort will be made to keep this table up to date if any new issues are discovered or introduced.

Back to [FAQ Index](#).

11.34 Why are some addresses of the EEPROM corrupted (usually address zero)?

The two most common reason for EEPROM corruption is either writing to the EEPROM beyond the datasheet endurance specification, or resetting the AVR while an EEPROM write is in progress.

EEPROM writes can take up to tens of milliseconds to complete. So that the CPU is not tied up for that long of time, an internal state-machine handles EEPROM write requests. The EEPROM state-machine expects to have all of the EEPROM registers setup, then an EEPROM write request to start the process. Once the EEPROM state-machine has started, changing EEPROM related registers during an EEPROM write is guaranteed to corrupt the EEPROM write process. The datasheet always shows the proper way to tell when a write is in progress, so that the registers are not changed by the user's program. The EEPROM state-machine will **always** complete the write in progress unless power is removed from the device.

As with all EEPROM technology, if power fails during an EEPROM write the state of the byte being written is undefined.

In older generation AVRs the EEPROM Address Register (EEAR) is initialized to zero on reset, be it from Brown Out Detect, Watchdog or the Reset Pin. If an EEPROM write has just started at the time of the reset, the write will be completed, but now at address zero instead of the requested address. If the reset occurs later in the write process both the requested address and address zero may be corrupted.

To distinguish which AVRs may exhibit the corrupt of address zero while a write is in process during a reset, look at the "initial value" section for the EEPROM Address Register. If EEAR shows the initial value as 0x00 or 0x0000, then address zero and possibly the one being written will be corrupted. Newer parts show the initial value as "undefined", these will not corrupt address zero during a reset (unless it was address zero that was being written).

EEPROMs have limited write endurance. The datasheet specifies the number of EEPROM writes that are guaranteed to function across the full temperature specification of the AVR, for a given byte. A read should always be performed before a write, to see if the value in the EEPROM actually needs to be written, so not to cause unnecessary EEPROM wear.

The failure mechanism for an overwritten byte is generally one of "stuck" bits, i. e. a bit will stay at a one or zero state regardless of the byte written. Also a write followed by a read may return the correct data, but the data will change with the passage of time, due the EEPROM's inability to hold a charge from the excessive write wear.

Back to [FAQ Index](#).

11.35 Why is my baud rate wrong?

Some AVR datasheets give the following formula for calculating baud rates:

$$(F_CPU / (UART_BAUD_RATE * 16L) - 1)$$

Unfortunately that formula does not work with all combinations of clock speeds and baud rates due to integer truncation during the division operator.

When doing integer division it is usually better to round to the nearest integer, rather than to the lowest. To do this add 0.5 (i. e. half the value of the denominator) to the numerator before the division, resulting in the formula:

$$((F_CPU + UART_BAUD_RATE * 8L) / (UART_BAUD_RATE * 16L) - 1)$$

This is also the way it is implemented in [<util/setbaud.h>: Helper macros for baud rate calculations](#).

Back to [FAQ Index](#).

11.36 On a device with more than 128 KiB of flash, how to make function pointers work?

Function pointers beyond the "magical" 128 KiB barrier(s) on larger devices are supposed to be resolved through so-called *trampolines* by the linker, so the actual pointers used in the code can remain 16 bits wide.

In order for this to work, the option `-mrelax` must be given on the compiler command-line that is used to link the final ELF file. (Older compilers did not implement this option for the AVR, use `-Wl,--relax` instead.)

Back to [FAQ Index](#).

11.37 Why is assigning ports in a "chain" a bad idea?

Suppose a number of IO port registers should get the value `0xff` assigned. Conveniently, it is implemented like this:

```
DDRB = DDRD = 0xff;
```

According to the rules of the C language, this causes `0xff` to be assigned to `DDRD`, then `DDRD` is read back, and the value is assigned to `DDRB`. The compiler stands no chance to optimize the readback away, as an IO port register is declared "volatile". Thus, chaining that kind of IO port assignments would better be avoided, using explicit assignments instead:

```
DDRB = 0xff;  
DDRD = 0xff;
```

Even worse is this, e. g. on an ATmega1281:

```
DDRA = DDRB = DDRC = DDRD = DDRE = DDRF = DDRG = 0xff;
```

The same happens as outlined above. However, when reading back register `DDRG`, this register only implements 6 out of the 8 bits, so the two topmost (unimplemented) bits read back as 0! Consequently, all remaining `DDRx` registers get assigned the value `0x3f`, which does not match the intention of the developer in any way.

Back to [FAQ Index](#).

12 Building and Installing the GNU Tool Chain

This chapter shows how to build and install, from source code, a complete development environment for the AVR processors using the GNU toolset. There are two main sections, one for Linux, FreeBSD, and other Unix-like operating systems, and another section for Windows.

12.1 Building and Installing under Linux, FreeBSD, and Others

The default behaviour for most of these tools is to install every thing under the `/usr/local` directory. In order to keep the AVR tools separate from the base system, it is usually better to install everything into `/usr/local/avr`. If the `/usr/local/avr` directory does not exist, you should create it before trying to install anything. You will need `root` access to install there. If you don't have root access to the system, you can alternatively install in your home directory, for example, in `$HOME/local/avr`. Where you install is a completely arbitrary decision, but should be consistent for all the tools.

You specify the installation directory by using the `--prefix=dir` option with the `configure` script. It is important to install all the AVR tools in the same directory or some of the tools will not work correctly. To ensure consistency and simplify the discussion, we will use `$PREFIX` to refer to whatever directory you wish to install in. You can set this as an environment variable if you wish as such (using a Bourne-like shell):

```
$ PREFIX=$HOME/local/avr  
$ export PREFIX
```

Note

Be sure that you have your `PATH` environment variable set to search the directory you install everything in *before* you start installing anything. For example, if you use `--prefix=$PREFIX`, you must have `$PREFIX/bin` in your exported `PATH`. As such:

```
$ PATH=$PATH:$PREFIX/bin
$ export PATH
```

Warning

If you have `CC` set to anything other than `avr-gcc` in your environment, this will cause the configure script to fail. It is best to not have `CC` set at all.

Note

It is usually the best to use the latest released version of each of the tools.

12.2 Required Tools

- **GNU Binutils**

<http://sources.redhat.com/binutils/>
[Installation](#)

- **GCC**

<http://gcc.gnu.org/>
[Installation](#)

- **AVR LibC**

<http://savannah.gnu.org/projects/avr-libc/>
[Installation](#)

12.3 Optional Tools

You can develop programs for AVR devices without the following tools. They may or may not be of use for you.

- **AVRDUDE**

<http://savannah.nongnu.org/projects/avrdude/>
[Installation](#)
[Usage Notes](#)

- **GDB**

<http://sources.redhat.com/gdb/>
[Installation](#)

- **SimulAVR**

<http://savannah.gnu.org/projects/simulavr/>
[Installation](#)

- AVaRICE

<http://avarice.sourceforge.net/>

[Installation](#)

12.4 GNU Binutils for the AVR target

The **binutils** package provides all the low-level utilities needed in building and manipulating object files. Once installed, your environment will have an AVR assembler (`avr-as`), linker (`avr-ld`), and librarian (`avr-ar` and `avr-ranlib`). In addition, you get tools which extract data from object files (`avr-objcopy`), disassemble object file information (`avr-objdump`), and strip information from object files (`avr-strip`). Before we can build the C compiler, these tools need to be in place.

Download and unpack the source files:

```
$ bunzip2 -c binutils-<version>.tar.bz2 | tar xf -
$ cd binutils-<version>
```

Note

Replace `<version>` with the version of the package you downloaded.

If you obtained a gzip compressed file (`.gz`), use `gunzip` instead of `bunzip2`.

It is usually a good idea to configure and build **binutils** in a subdirectory so as not to pollute the source with the compiled files. This is recommended by the **binutils** developers.

```
$ mkdir obj-avr
$ cd obj-avr
```

The next step is to configure and build the tools. This is done by supplying arguments to the `configure` script that enable the AVR-specific options.

```
$ ../configure --prefix=$PREFIX --target=avr --disable-nls
```

If you don't specify the `--prefix` option, the tools will get installed in the `/usr/local` hierarchy (i.e. the binaries will get installed in `/usr/local/bin`, the info pages get installed in `/usr/local/info`, etc.) Since these tools are changing frequently, it is preferable to put them in a location that is easily removed.

When `configure` is run, it generates a lot of messages while it determines what is available on your operating system. When it finishes, it will have created several Makefiles that are custom tailored to your platform. At this point, you can build the project.

```
$ make
```

Note

BSD users should note that the project's Makefile uses GNU make syntax.

This means FreeBSD users may need to build the tools by using `gmake`.

If the tools compiled cleanly, you're ready to install them. If you specified a destination that isn't owned by your account, you'll need `root` access to install them. To install:

```
$ make install
```

You should now have the programs from `binutils` installed into `$PREFIX/bin`. Don't forget to [set your PATH](#) environment variable before going to build `avr-gcc`.

12.5 GCC for the AVR target

Warning

You *must* install [avr-binutils](#) and make sure your [path is set](#) properly before installing `avr-gcc`.

The steps to build `avr-gcc` are essentially same as for [binutils](#):

```
$ bunzip2 -c gcc-<version>.tar.bz2 | tar xf -
$ cd gcc-<version>
$ mkdir obj-avr
$ cd obj-avr
$ ../configure --prefix=$PREFIX --target=avr --enable-languages=c,c++ \
  --disable-nls --disable-libssp --with-dwarf2
$ make
$ make install
```

To save your self some download time, you can alternatively download only the `gcc-core-<version>.tar.bz2` and `gcc-c++-<version>.tar.bz2` parts of the gcc. Also, if you don't need C++ support, you only need the core part and should only enable the C language support. (Starting with GCC 4.7 releases, these split files are no longer available though.)

Note

Early versions of these tools did not support C++.

The `stdc++` libs are not included with C++ for AVR due to the size limitations of the devices.

12.6 AVR LibC

Warning

You *must* install [avr-binutils](#), [avr-gcc](#) and make sure your [path is set](#) properly before installing `avr-libc`.

Note

If you have obtained the latest `avr-libc` from `cvs`, you will have to run the `bootstrap` script before using either of the build methods described below.

To build and install `avr-libc`:

```
$ gunzip -c avr-libc-<version>.tar.gz | tar xf -
$ cd avr-libc-<version>
$ ./configure --prefix=$PREFIX --build='./config.guess' --host=avr
$ make
$ make install
```

Optionally, generation of debug information can be requested with:

```
$ gunzip -c avr-libc-<version>.tar.gz | tar xf -
$ cd avr-libc-<version>
$ ./configure --prefix=$PREFIX --build='./config.guess' --host=avr \
  --with-debug-info=DEBUG_INFO
$ make
$ make install
```

where `DEBUG_INFO` can be one of `stabs`, `dwarf-2`, or `dwarf-4`.

The default is to not generate any debug information, which is suitable for binary distributions of `avr-libc`, where the user does not have the source code installed the debug information would refer to.

12.7 AVRDUDE

Note

It has been ported to windows (via MinGW or cygwin), Linux and Solaris. Other Unix systems should be trivial to port to.

avrdude is part of the FreeBSD ports system. To install it, simply do the following:

```
# cd /usr/ports/devel/avrdude
# make install
```

Note

Installation into the default location usually requires root permissions. However, running the program only requires access permissions to the appropriate `pp1 (4)` device.

Building and installing on other systems should use the `configure` system, as such:

```
$ gunzip -c avrdude-<version>.tar.gz | tar xf -
$ cd avrdude-<version>
$ mkdir obj-avr
$ cd obj-avr
$ ../configure --prefix=$PREFIX
$ make
$ make install
```

12.8 GDB for the AVR target

GDB also uses the `configure` system, so to build and install:

```
$ bunzip2 -c gdb-<version>.tar.bz2 | tar xf -
$ cd gdb-<version>
$ mkdir obj-avr
$ cd obj-avr
$ ../configure --prefix=$PREFIX --target=avr
$ make
$ make install
```

Note

If you are planning on using `avr-gdb`, you will probably want to install either [simulavr](#) or [avarice](#) since `avr-gdb` needs one of these to run as a remote target backend.

12.9 SimulAVR

SimulAVR also uses the `configure` system, so to build and install:

```
$ gunzip -c simulavr-<version>.tar.gz | tar xf -
$ cd simulavr-<version>
$ mkdir obj-avr
$ cd obj-avr
$ ../configure --prefix=$PREFIX
$ make
$ make install
```

Note

You might want to have already installed [avr-binutils](#), [avr-gcc](#) and [avr-libc](#) if you want to have the test programs built in the `simulavr` source.

12.10 AVaRICE

Note

These install notes are not applicable to `avarice-1.5` or older. You probably don't want to use anything that old anyways since there have been many improvements and bug fixes since the 1.5 release.

AVaRICE also uses the `configure` system, so to build and install:

```
$ gunzip -c avarice-<version>.tar.gz | tar xf -
$ cd avarice-<version>
$ mkdir obj-avr
$ cd obj-avr
$ ../configure --prefix=$PREFIX
$ make
$ make install
```

Note

AVaRICE uses the BFD library for accessing various binary file formats. You may need to tell the `configure` script where to find the lib and headers for the link to work. This is usually done by invoking the `configure` script like this (Replace

<hdr_path> with the path to the `bfd.h` file on your system. Replace <lib_path> with the path to `libbfd.a` on your system.):

```
$ CPPFLAGS=-I<hdr_path> LDFLAGS=-L<lib_path> ../configure --prefix=$PREFIX
```

12.11 Building and Installing under Windows

Building and installing the toolchain under Windows requires more effort because all of the tools required for building, and the programs themselves, are mainly designed for running under a POSIX environment such as Unix and Linux. Windows does not natively provide such an environment.

There are two projects available that provide such an environment, Cygwin and MinGW. There are advantages and disadvantages to both. Cygwin provides a very complete POSIX environment that allows one to build many Linux based tools from source with very little or no source modifications. However, POSIX functionality is provided in the form of a DLL that is linked to the application. This DLL has to be redistributed with your application and there are issues if the Cygwin DLL already exists on the installation system and different versions of the DLL. On the other hand, MinGW can compile code as native Win32 applications. However, this means that programs designed for Unix and Linux (i.e. that use POSIX functionality) will not compile as MinGW does not provide that POSIX layer for you. Therefore most programs that compile on both types of host systems, usually must provide some sort of abstraction layer to allow an application to be built cross-platform.

MinGW does provide somewhat of a POSIX environment, called MSYS, that allows you to build Unix and Linux applications as they would normally do, with a `configure` step and a `make` step. Cygwin also provides such an environment. This means that building the AVR toolchain is very similar to how it is built in Linux, described above. The main differences are in what the `PATH` environment variable gets set to, pathname differences, and the tools that are required to build the projects under Windows. We'll take a look at the tools next.

12.12 Tools Required for Building the Toolchain for Windows

These are the tools that are currently used to build an AVR tool chain. This list may change, either the version of the tools, or the tools themselves, as improvements are made.

- **MinGW**

Download the MinGW Automated Installer, 20100909 (or later)

<<http://sourceforge.net/projects/mingw/files/Automated%20MinGW%20Installer>

- Run `mingw-get-inst-20100909.exe`
- In the installation wizard, keep the default values and press the "Next" button for all installer pages except for the pages explicitly listed below.
- In the installer page "Repository Catalogues", select the "Download latest repository catalogues" radio button, and press the "Next" button

- In the installer page "License Agreement", select the "I accept the agreement" radio button, and press the "Next" button
- In the installer page "Select Components", be sure to select these items:
 - * C compiler (default checked)
 - * C++ compiler
 - * Ada compiler
 - * MinGW Developer Toolkit (which includes "MSYS Basic System").
- Install.

• Install Cygwin

- Install everything, all users, UNIX line endings. This will take a *long* time. A fat internet pipe is highly recommended. It is also recommended that you download all to a directory first, and then install from that directory to your machine.

Note

GMP, MPFR, and MPC are required to build GCC.
 GMP is a prerequisite for building MPFR. Build GMP first.
 MPFR is a prerequisite for building MPC. Build MPFR second.

• Build GMP for MinGW

- Latest Version
- <<http://gmplib.org/>>
- Build script:


```
./configure 2>&1 | tee gmp-configure.log
make 2>&1 | tee gmp-make.log
make check 2>&1 | tee gmp-make-check.log
make install 2>&1 | tee gmp-make-install.log
```
- GMP headers will be installed under /usr/local/include and library installed under /usr/local/lib.

• Build MPFR for MinGW

- Latest Version
- <<http://www.mpfr.org/>>
- Build script:


```
./configure --with-gmp=/usr/local --disable-shared 2>&1 | tee mpfr-configure.log
make 2>&1 | tee mpfr-make.log
make check 2>&1 | tee mpfr-make-check.log
make install 2>&1 | tee mpfr-make-install.log
```
- MPFR headers will be installed under /usr/local/include and library installed under /usr/local/lib.

- **Build MPC for MinGW**

- Latest Version
- <<http://www.multiprecision.org/>>
- Build script:

```
./configure --with-gmp=/usr/local --with-mpfr=/usr/local --disable-shared 2>&1 |
make          2>&1 | tee mpfr-make.log
make check    2>&1 | tee mpfr-make-check.log
make install  2>&1 | tee mpfr-make-install.log
```

- MPFR headers will be installed under /usr/local/include and library installed under /usr/local/lib.

Note

Doxygen is required to build AVR-LibC documentation.

- **Install Doxygen**

- Version 1.7.2
- <<http://www.stack.nl/~dimitri/doxygen/>>
- Download and install.

NetPBM is required to build graphics in the AVR-LibC documentation.

- **Install NetPBM**

- Version 10.27.0
- From the GNUWin32 project: <<http://gnuwin32.sourceforge.net/packages.html>>
- Download and install.

fig2dev is required to build graphics in the AVR-LibC documentation.

- **Install fig2dev**

- Version 3.2 patchlevel 5c
- From WinFig 4.62: <<http://www.schmidt-web-berlin.de/winfig/>>
- Download the zip file version of WinFig
- Unzip the download file and install fig2dev.exe in a location of your choice, somewhere in the PATH.
- You may have to unzip and install related DLL files for fig2dev. In the version above, you have to install QtCore4.dll and QtGui4.dll.

MikTeX is required to build various documentation.

- **Install MiKTeX**

- Version 2.9
- <<http://miktex.org/>>
- Download and install.

Ghostscript is required to build various documentation.

- **Install Ghostscript**

- Version 9.00
- <<http://www.ghostscript.com>>

- Download and install.
- In the \bin subdirectory of the installaion, copy gswin32c.exe to gs.exe.
- Set the TEMP and TMP environment variables to **c:\temp** or to the short filename version. This helps to avoid NTVDM errors during building.

12.13 Building the Toolchain for Windows

All directories in the PATH environment variable should be specified using their short filename (8.3) version. This will also help to avoid NTVDM errors during building. These short filenames can be specific to each machine.

Build the tools below in MinGW/MSYS.

• Binutils

- Open source code package and patch as necessary.
- Configure and build in a directory outside of the source code tree.
- Set PATH, in order:
 - * <MikTeX executables>
 - * <ghostscript executables>
 - * /usr/local/bin
 - * /usr/bin
 - * /bin
 - * /mingw/bin
 - * c:/cygwin/bin
 - * <install directory>/bin
- Configure

```
CFLAGS=-D__USE_MINGW_ACCESS \
../$archivedir/configure \
--prefix=$installdir \
--target=avr \
--disable-nls \
--enable-doc \
--datadir=$installdir/doc/binutils \
--with-gmp=/usr/local \
--with-mpfr=/usr/local \
2>&1 | tee binutils-configure.log
```

- Make

```
make all html install install-html 2>&1 | tee binutils-make.log
```

- Manually change documentation location.

• GCC

- Open source code package and patch as necessary.
- Configure and build in a directory outside of the source code tree.

- Set PATH, in order:
 - * <MikTex executables>
 - * <ghostscript executables>
 - * /usr/local/bin
 - * /usr/bin
 - * /bin
 - * /mingw/bin
 - * c:/cygwin/bin
 - * <install directory>/bin
- Configure


```
LDFLAGS='-L /usr/local/lib -R /usr/local/lib' \
CFLAGS='-D__USE_MINGW_ACCESS' \
../gcc-$version/configure \
--with-gmp=/usr/local \
--with-mpfr=/usr/local \
--with-mpc=/usr/local \
--prefix=$installdir \
--target=$target \
--enable-languages=c,c++ \
--with-dwarf2 \
--enable-doc \
--with-docdir=$installdir/doc/$project \
--disable-shared \
--disable-libada \
--disable-libssp \
2>&1 | tee $project-configure.log
```
- Make


```
make all html install 2>&1 | tee $package-make.log
```

- **avr-libc**

- Open source code package.
- Configure and build at the top of the source code tree.
- Set PATH, in order:
 - * /usr/local/bin
 - * /mingw/bin
 - * /bin
 - * <MikTex executables>
 - * <install directory>/bin
 - * <Doxygen executables>
 - * <NetPBM executables>
 - * <fig2dev executable>
 - * <Ghostscript executables>
 - * c:/cygwin/bin
- Configure

```
./configure \
  --host=avr \
  --prefix=$installdir \
  --enable-doc \
  --disable-versioned-doc \
  --enable-html-doc \
  --enable-xml-doc \
  --enable-pdf-doc \
  --enable-man-doc \
  --mandir=$installdir/man \
  --datadir=$installdir \
  2>&1 | tee $package-configure.log
```

– Make

```
make all install 2>&1 | tee $package-make.log
```

- Manually change location of man page documentation.
- Move the examples to the top level of the install tree.
- Convert line endings in examples to Windows line endings.
- Convert line endings in header files to Windows line endings.

• AVRDUDE

- Open source code package.
- Configure and build at the top of the source code tree.
- Set PATH, in order:

```
* <MikTex executables>
* /usr/local/bin
* /usr/bin
* /bin
* /mingw/bin
* c:/cygwin/bin
* <install directory>/bin
```

- Set location of LibUSB headers and libraries

```
export CPPFLAGS="-I../libusb-win32-device-bin-$libusb_version/include"
export CFLAGS="-I../libusb-win32-device-bin-$libusb_version/include"
export LDFLAGS="-L../libusb-win32-device-bin-$libusb_version/lib/gcc"
```

- Configure

```
./configure \
  --prefix=$installdir \
  --datadir=$installdir \
  --sysconfdir=$installdir/bin \
  --enable-doc \
  --disable-versioned-doc \
  2>&1 | tee $package-configure.log
```

- Make

```
make -k all install 2>&1 | tee $package-make.log
```

- Convert line endings in avrdude config file to Windows line endings.
- Delete backup copy of avrdude config file in install directory if exists.

- **Insight/GDB**

- Open source code package and patch as necessary.
- Configure and build in a directory outside of the source code tree.
- Set PATH, in order:
 - * <MikTex executables>
 - * /usr/local/bin
 - * /usr/bin
 - * /bin
 - * /mingw/bin
 - * c:/cygwin/bin
 - * <install directory>/bin

- Configure

```
CFLAGS=-D__USE_MINGW_ACCESS \
LDFLAGS='-static' \
../$archivedir/configure \
  --prefix=$installdir \
  --target=avr \
  --with-gmp=/usr/local \
  --with-mpfr=/usr/local \
  --enable-doc \
2>&1 | tee insight-configure.log
```

- Make

```
make all install 2>&1 | tee $package-make.log
```

- **SRecord**

- Open source code package.
- Configure and build at the top of the source code tree.
- Set PATH, in order:
 - * <MikTex executables>
 - * /usr/local/bin
 - * /usr/bin
 - * /bin
 - * /mingw/bin
 - * c:/cygwin/bin
 - * <install directory>/bin

- Configure

```
./configure \
  --prefix=$installdir \
  --infodir=$installdir/info \
  --mandir=$installdir/man \
2>&1 | tee $package-configure.log
```

- Make

```
make all install 2>&1 | tee $package-make.log
```

Build the tools below in Cygwin.

- **AVaRICE**

- Open source code package.
- Configure and build in a directory outside of the source code tree.
- Set PATH, in order:

```
* <MikTex executables>
* /usr/local/bin
* /usr/bin
* /bin
* <install directory>/bin
```

- Set location of LibUSB headers and libraries

```
export CPPFLAGS=-I$startdir/libusb-win32-device-bin-$libusb_version/include
export CFLAGS=-I$startdir/libusb-win32-device-bin-$libusb_version/include
export LDFLAGS="-static -L$startdir/libusb-win32-device-bin-$libusb_version/lib/g
```

- Configure

```
../$archivedir/configure \
--prefix=$installdir \
--datadir=$installdir/doc \
--mandir=$installdir/man \
--infodir=$installdir/info \
2>&1 | tee avarice-configure.log
```

- Make

```
make all install 2>&1 | tee avarice-make.log
```

- **SimulAVR**

- Open source code package.
- Configure and build in a directory outside of the source code tree.
- Set PATH, in order:

```
* <MikTex executables>
* /usr/local/bin
* /usr/bin
* /bin
* <install directory>/bin
```

- Configure


```
export LDFLAGS="-static"
../$archivedir/configure \
  --prefix=$installdir \
  --datadir=$installdir \
  --disable-tests \
  --disable-versioned-doc \
  2>&1 | tee simulavr-configure.log
```

– Make

```
make -k all install 2>&1 | tee simulavr-make.log
make pdf install-pdf 2>&1 | tee simulavr-pdf-make.log
```

13 Using the GNU tools

This is a short summary of the AVR-specific aspects of using the GNU tools. Normally, the generic documentation of these tools is fairly large and maintained in `texinfo` files. Command-line options are explained in detail in the manual page.

13.1 Options for the C compiler `avr-gcc`

13.1.1 Machine-specific options for the AVR

The following machine-specific options are recognized by the C compiler frontend. In addition to the preprocessor macros indicated in the tables below, the preprocessor will define the macros `__AVR__` and `__AVR__` (to the value 1) when compiling for an AVR target. The macro `AVR` will be defined as well when using the standard levels `gnu89` (default) and `gnu99` but not with `c89` and `c99`.

- `-mmcu=architecture`

Compile code for *architecture*. Currently known architectures are

Architecture	PBSMacros	PBSDescription
avr1	PBS__AVR_ARCH__=1 __AVR_ASM_ONLY__ __AVR_2_BYTE_PC__ [2]	PBSSimple CPU core, only assembler support
avr2	PBS__AVR_ARCH__=2 __AVR_2_BYTE_PC__ [2]	PBS"Classic" CPU core, up to 8 KB of ROM
avr25 [1]	PBS__AVR_ARCH__=25 __AVR_HAVE_MO VW__ [1] __AVR_HAVE_LPMX__ [1] __AVR_2_BYTE_PC__ [2]	PBS"Classic" CPU core with 'MOVW' and 'LPM Rx, Z[+]' instruction, up to 8 KB of ROM
avr3	PBS__AVR_ARCH__=3 __AVR_MEGA__ [5] __AVR_HAVE_JMP_CALL__ [4] __AVR_2_BYTE_PC__ [2]	PBS"Classic" CPU core, 16 KB to 64 KB of ROM

avr31	PBS__AVR_ARCH__=31 __AVR_MEGA__ [5] __AVR_HAVE_JMP_CALL__ [4] __AVR_HAVE_RAMPZ__ [4] __AVR_HAVE_ELPM__ [4] __AVR_2_BYTE_PC__ [2]	PBS"Classic" CPU core, 128 KB of ROM
avr35 [3]	PBS__AVR_ARCH__=35 __AVR_MEGA__ [5] __AVR_HAVE_JMP_CALL__ [4] __AVR_HAVE_MO VW__ [1] __AVR_HAVE_LPMX__ [1] __AVR_2_BYTE_PC__ [2]	PBS"Classic" CPU core with 'MOVW' and 'LPM Rx, Z[+]' instruction, 16 KB to 64 KB of ROM
avr4	PBS__AVR_ARCH__=4 __AVR_ENHANCED__ [5] __AVR_HAVE_MO VW__ [1] __AVR_HAVE_LPMX__ [1] __AVR_HAVE_MUL__ [1] __AVR_2_BYTE_PC__ [2]	PBS"Enhanced" CPU core, up to 8 KB of ROM
avr5	PBS__AVR_ARCH__=5 __AVR_MEGA__ [5] __AVR_ENHANCED__ [5] __AVR_HAVE_JMP_CALL__ [4] __AVR_HAVE_MO VW__ [1] __AVR_HAVE_LPMX__ [1] __AVR_HAVE_MUL__ [1] __AVR_2_BYTE_PC__ [2]	PBS"Enhanced" CPU core, 16 KB to 64 KB of ROM
avr51	PBS__AVR_ARCH__=51 __AVR_MEGA__ [5] __AVR_ENHANCED__ [5] __AVR_HAVE_JMP_CALL__ [4] __AVR_HAVE_MO VW__ [1] __AVR_HAVE_LPMX__ [1] __AVR_HAVE_MUL__ [1] __AVR_HAVE_RAMPZ__ [4] __AVR_HAVE_ELPM__ [4] __AVR_HAVE_ELPMX__ [4] __AVR_2_BYTE_PC__ [2]	PBS"Enhanced" CPU core, 128 KB of ROM
avr6 [2]	PBS__AVR_ARCH__=6 __AVR_MEGA__ [5] __AVR_ENHANCED__ [5] __AVR_HAVE_JMP_CALL__ [4] __AVR_HAVE_MO VW__ [1] __AVR_HAVE_LPMX__ [1] __AVR_HAVE_MUL__ [1] __AVR_HAVE_RAMPZ__ [4] __AVR_HAVE_ELPM__ [4] __AVR_HAVE_ELPMX__ [4] __AVR_3_BYTE_PC__ [2]	PBS"Enhanced" CPU core, 256 KB of ROM

[1] New in GCC 4.2

[2] Unofficial patch for GCC 4.1

[3] New in GCC 4.2.3

[4] New in GCC 4.3

[5] Obsolete.

By default, code is generated for the avr2 architecture.

Note that when only using `-mmcu=architecture` but no `-mmcu=MCU type`, including the file `<avr/io.h>` cannot work since it cannot decide which device's definitions to select.

- `-mmcu=MCU type`

The following MCU types are currently understood by avr-gcc. The table matches them against the corresponding avr-gcc architecture name, and shows the preprocessor symbol declared by the `-mmcu` option.

Architecture	MCU name	Macro
avr1	at90s1200	__AVR_AT90S1200__
avr1	attiny11	__AVR_ATtiny11__
avr1	attiny12	__AVR_ATtiny12__
avr1	attiny15	__AVR_ATtiny15__
avr1	attiny28	__AVR_ATtiny28__
avr2	at90s2313	__AVR_AT90S2313__
avr2	at90s2323	__AVR_AT90S2323__
avr2	at90s2333	__AVR_AT90S2333__
avr2	at90s2343	__AVR_AT90S2343__
avr2	attiny22	__AVR_ATtiny22__
avr2	attiny26	__AVR_ATtiny26__
avr2	at90s4414	__AVR_AT90S4414__
avr2	at90s4433	__AVR_AT90S4433__
avr2	at90s4434	__AVR_AT90S4434__
avr2	at90s8515	__AVR_AT90S8515__
avr2	at90c8534	__AVR_AT90C8534__
avr2	at90s8535	__AVR_AT90S8535__
avr2/avr25 [1]	at86rf401	__AVR_AT86RF401__
avr2/avr25 [1]	ata5272	__AVR_ATA5272__
avr2/avr25 [1]	ata6616c	__AVR_ATA6616C__
avr2/avr25 [1]	attiny13	__AVR_ATtiny13__
avr2/avr25 [1]	attiny13a	__AVR_ATtiny13A__
avr2/avr25 [1]	attiny2313	__AVR_ATtiny2313__
avr2/avr25 [1]	attiny2313a	__AVR_ATtiny2313A__
avr2/avr25 [1]	attiny24	__AVR_ATtiny24__
avr2/avr25 [1]	attiny24a	__AVR_ATtiny24A__
avr2/avr25 [1]	attiny25	__AVR_ATtiny25__
avr2/avr25 [1]	attiny261	__AVR_ATtiny261__
avr2/avr25 [1]	attiny261a	__AVR_ATtiny261A__
avr2/avr25 [1]	attiny4313	__AVR_ATtiny4313__

avr2/avr25 [1]	attiny43u	__AVR_ATtiny43U__
avr2/avr25 [1]	attiny44	__AVR_ATtiny44__
avr2/avr25 [1]	attiny44a	__AVR_ATtiny44A__
avr2/avr25 [1]	attiny441	__AVR_ATtiny441__
avr2/avr25 [1]	attiny45	__AVR_ATtiny45__
avr2/avr25 [1]	attiny461	__AVR_ATtiny461__
avr2/avr25 [1]	attiny461a	__AVR_ATtiny461A__
avr2/avr25 [1]	attiny48	__AVR_ATtiny48__
avr2/avr25 [1]	attiny828	__AVR_ATtiny828__
avr2/avr25 [1]	attiny84	__AVR_ATtiny84__
avr2/avr25 [1]	attiny84a	__AVR_ATtiny84A__
avr2/avr25 [1]	attiny841	__AVR_ATtiny841__
avr2/avr25 [1]	attiny85	__AVR_ATtiny85__
avr2/avr25 [1]	attiny861	__AVR_ATtiny861__
avr2/avr25 [1]	attiny861a	__AVR_ATtiny861A__
avr2/avr25 [1]	attiny87	__AVR_ATtiny87__
avr2/avr25 [1]	attiny88	__AVR_ATtiny88__
avr3	atmega603	__AVR_ATmega603__
avr3	at43usb355	__AVR_AT43USB355__
avr3/avr31 [3]	atmega103	__AVR_ATmega103__
avr3/avr31 [3]	at43usb320	__AVR_AT43USB320__
avr3/avr35 [2]	at90usb82	__AVR_AT90USB82__
avr3/avr35 [2]	at90usb162	__AVR_AT90USB162__
avr3/avr35 [2]	ata5505	__AVR_ATA5505__
avr3/avr35 [2]	ata6617c	__AVR_ATA6617C__
avr3/avr35 [2]	ata664251	__AVR_ATA664251__
avr3/avr35 [2]	atmega8u2	__AVR_ATmega8U2__
avr3/avr35 [2]	atmega16u2	__AVR_ATmega16U2__
avr3/avr35 [2]	atmega32u2	__AVR_ATmega32U2__
avr3/avr35 [2]	attiny167	__AVR_ATtiny167__
avr3/avr35 [2]	attiny1634	__AVR_ATtiny1634__
avr3	at76c711	__AVR_AT76C711__
avr4	ata6285	__AVR_ATA6285__
avr4	ata6286	__AVR_ATA6286__
avr4	ata6289	__AVR_ATA6289__
avr4	ata6612c	__AVR_ATA6612C__
avr4	atmega48	__AVR_ATmega48__
avr4	atmega48a	__AVR_ATmega48A__
avr4	atmega48pa	__AVR_ATmega48PA__
avr4	atmega48pb	__AVR_ATmega48PB__
avr4	atmega48p	__AVR_ATmega48P__
avr4	atmega8	__AVR_ATmega8__
avr4	atmega8a	__AVR_ATmega8A__
avr4	atmega8515	__AVR_ATmega8515__
avr4	atmega8535	__AVR_ATmega8535__
avr4	atmega88	__AVR_ATmega88__
avr4	atmega88a	__AVR_ATmega88A__
avr4	atmega88p	__AVR_ATmega88P__
avr4	atmega88pa	__AVR_ATmega88PA__

avr4	atmega88pb	__AVR_ATmega88PB__
avr4	atmega8hva	__AVR_ATmega8HVA__
avr4	at90pwm1	__AVR_AT90PWM1__
avr4	at90pwm2	__AVR_AT90PWM2__
avr4	at90pwm2b	__AVR_AT90PWM2B__
avr4	at90pwm3	__AVR_AT90PWM3__
avr4	at90pwm3b	__AVR_AT90PWM3B__
avr4	at90pwm81	__AVR_AT90PWM81__
avr5	at90can32	__AVR_AT90CAN32__
avr5	at90can64	__AVR_AT90CAN64__
avr5	at90pwm161	__AVR_AT90PWM161__
avr5	at90pwm216	__AVR_AT90PWM216__
avr5	at90pwm316	__AVR_AT90PWM316__
avr5	at90scr100	__AVR_AT90SCR100__
avr5	at90usb646	__AVR_AT90USB646__
avr5	at90usb647	__AVR_AT90USB647__
avr5	at94k	__AVR_AT94K__
avr5	atmega16	__AVR_ATmega16__
avr5	ata5702m322	__AVR_ATA5702M322__
avr5	ata5782	__AVR_ATA5782__
avr5	ata5790	__AVR_ATA5790__
avr5	ata5790n	__AVR_ATA5790N__
avr5	ata5791	__AVR_ATA5791__
avr5	ata5795	__AVR_ATA5795__
avr5	ata5831	__AVR_ATA5831__
avr5	ata6613c	__AVR_ATA6613C__
avr5	ata6614q	__AVR_ATA6614Q__
avr5	ata8210	__AVR_ATA8210__
avr5	ata8510	__AVR_ATA8510__
avr5	atmega161	__AVR_ATmega161__
avr5	atmega162	__AVR_ATmega162__
avr5	atmega163	__AVR_ATmega163__
avr5	atmega164a	__AVR_ATmega164A__
avr5	atmega164p	__AVR_ATmega164P__
avr5	atmega164pa	__AVR_ATmega164PA__
avr5	atmega165	__AVR_ATmega165__
avr5	atmega165a	__AVR_ATmega165A__
avr5	atmega165p	__AVR_ATmega165P__
avr5	atmega165pa	__AVR_ATmega165PA__
avr5	atmega168	__AVR_ATmega168__
avr5	atmega168a	__AVR_ATmega168A__
avr5	atmega168p	__AVR_ATmega168P__
avr5	atmega168pa	__AVR_ATmega168PA__
avr5	atmega168pb	__AVR_ATmega168PB__
avr5	atmega169	__AVR_ATmega169__
avr5	atmega169a	__AVR_ATmega169A__
avr5	atmega169p	__AVR_ATmega169P__
avr5	atmega169pa	__AVR_ATmega169PA__
avr5	atmega16a	__AVR_ATmega16A__
avr5	atmega16hva	__AVR_ATmega16HVA__

avr5	atmega16hva2	__AVR_ATmega16HVA2__
avr5	atmega16hvb	__AVR_ATmega16HVB__
avr5	atmega16hvbrevb	__AVR_ATmega16HVBREVB__
avr5	atmega16m1	__AVR_ATmega16M1__
avr5	atmega16u4	__AVR_ATmega16U4__
avr5	atmega32	__AVR_ATmega32__
avr5	atmega32a	__AVR_ATmega32A__
avr5	atmega323	__AVR_ATmega323__
avr5	atmega324a	__AVR_ATmega324A__
avr5	atmega324p	__AVR_ATmega324P__
avr5	atmega324pa	__AVR_ATmega324PA__
avr5	atmega325	__AVR_ATmega325__
avr5	atmega325a	__AVR_ATmega325A__
avr5	atmega325p	__AVR_ATmega325P__
avr5	atmega325pa	__AVR_ATmega325PA__
avr5	atmega3250	__AVR_ATmega3250__
avr5	atmega3250a	__AVR_ATmega3250A__
avr5	atmega3250p	__AVR_ATmega3250P__
avr5	atmega3250pa	__AVR_ATmega3250PA__
avr5	atmega328	__AVR_ATmega328__
avr5	atmega328p	__AVR_ATmega328P__
avr5	atmega329	__AVR_ATmega329__
avr5	atmega329a	__AVR_ATmega329A__
avr5	atmega329p	__AVR_ATmega329P__
avr5	atmega329pa	__AVR_ATmega329PA__
avr5	atmega3290	__AVR_ATmega3290__
avr5	atmega3290a	__AVR_ATmega3290A__
avr5	atmega3290p	__AVR_ATmega3290P__
avr5	atmega3290pa	__AVR_ATmega3290PA__
avr5	atmega32c1	__AVR_ATmega32C1__
avr5	atmega32hvb	__AVR_ATmega32HVB__
avr5	atmega32hvbrevb	__AVR_ATmega32HVBREVB__
avr5	atmega32m1	__AVR_ATmega32M1__
avr5	atmega32u4	__AVR_ATmega32U4__
avr5	atmega32u6	__AVR_ATmega32U6__
avr5	atmega406	__AVR_ATmega406__
avr5	atmega644rfr2	__AVR_ATmega644RFR2__
avr5	atmega64rfr2	__AVR_ATmega64RFR2__
avr5	atmega64	__AVR_ATmega64__
avr5	atmega64a	__AVR_ATmega64A__
avr5	atmega640	__AVR_ATmega640__
avr5	atmega644	__AVR_ATmega644__
avr5	atmega644a	__AVR_ATmega644A__
avr5	atmega644p	__AVR_ATmega644P__
avr5	atmega644pa	__AVR_ATmega644PA__
avr5	atmega645	__AVR_ATmega645__
avr5	atmega645a	__AVR_ATmega645A__
avr5	atmega645p	__AVR_ATmega645P__
avr5	atmega6450	__AVR_ATmega6450__

avr5	atmega6450a	__AVR_ATmega6450A__
avr5	atmega6450p	__AVR_ATmega6450P__
avr5	atmega649	__AVR_ATmega649__
avr5	atmega649a	__AVR_ATmega649A__
avr5	atmega6490	__AVR_ATmega6490__
avr5	atmega6490a	__AVR_ATmega6490A__
avr5	atmega6490p	__AVR_ATmega6490P__
avr5	atmega649p	__AVR_ATmega649P__
avr5	atmega64c1	__AVR_ATmega64C1__
avr5	atmega64hve	__AVR_ATmega64HVE__
avr5	atmega64hve2	__AVR_ATmega64HVE2__
avr5	atmega64m1	__AVR_ATmega64M1__
avr5	m3000	__AVR_M3000__
avr5/avr51 [3]	at90can128	__AVR_AT90CAN128__
avr5/avr51 [3]	at90usb1286	__AVR_AT90USB1286__
avr5/avr51 [3]	at90usb1287	__AVR_AT90USB1287__
avr5/avr51 [3]	atmega128	__AVR_ATmega128__
avr5/avr51 [3]	atmega128a	__AVR_ATmega128A__
avr5/avr51 [3]	atmega1280	__AVR_ATmega1280__
avr5/avr51 [3]	atmega1281	__AVR_ATmega1281__
avr5/avr51 [3]	atmega1284	__AVR_ATmega1284__
avr5/avr51 [3]	atmega1284p	__AVR_ATmega1284P__
avr5/avr51 [3]	atmega1284rfr2	__AVR_ATmega1284RFR2__
avr5/avr51 [3]	atmega128rfr2	__AVR_ATmega128RFR2__
avr6	atmega2560	__AVR_ATmega2560__
avr6	atmega2561	__AVR_ATmega2561__
avr6	atmega2564rfr2	__AVR_ATmega2564RFR2__
avr6	atmega256rfr2	__AVR_ATmega256RFR2__
avrxcmega2	atxmega8e5	__AVR_ATxmega8E5__
avrxcmega2	atxmega16a4	__AVR_ATxmega16A4__
avrxcmega2	atxmega16a4u	__AVR_ATxmega16A4U__
avrxcmega2	atxmega16c4	__AVR_ATxmega16C4__
avrxcmega2	atxmega16d4	__AVR_ATxmega16D4__
avrxcmega2	atxmega16e5	__AVR_ATxmega16E5__
avrxcmega2	atxmega32a4	__AVR_ATxmega32A4__
avrxcmega2	atxmega32a4u	__AVR_ATxmega32A4U__
avrxcmega2	atxmega32c3	__AVR_ATxmega32C3__
avrxcmega2	atxmega32c4	__AVR_ATxmega32C4__
avrxcmega2	atxmega32d3	__AVR_ATxmega32D3__
avrxcmega2	atxmega32d4	__AVR_ATxmega32D4__
avrxcmega2	atxmega32e5	__AVR_ATxmega32E5__
avrxcmega4	atxmega64a3	__AVR_ATxmega64A3__
avrxcmega4	atxmega64a3u	__AVR_ATxmega64A3U__
avrxcmega4	atxmega64a4u	__AVR_ATxmega64A4U__
avrxcmega4	atxmega64b1	__AVR_ATxmega64B1__
avrxcmega4	atxmega64b3	__AVR_ATxmega64B3__
avrxcmega4	atxmega64c3	__AVR_ATxmega64C3__
avrxcmega4	atxmega64d3	__AVR_ATxmega64D3__
avrxcmega4	atxmega64d4	__AVR_ATxmega64D4__
avrxcmega5	atxmega64a1	__AVR_ATxmega64A1__

avrxtmega5	atxmega64a1u	__AVR_ATxmega64A1U__
avrxtmega6	atxmega128a3	__AVR_ATxmega128A3__
avrxtmega6	atxmega128a3u	__AVR_ATxmega128A3U__
avrxtmega6	atxmega128b1	__AVR_ATxmega128B1__
avrxtmega6	atxmega128b3	__AVR_ATxmega128B3__
avrxtmega6	atxmega128c3	__AVR_ATxmega128C3__
avrxtmega6	atxmega128d3	__AVR_ATxmega128D3__
avrxtmega6	atxmega128d4	__AVR_ATxmega128D4__
avrxtmega6	atxmega192a3	__AVR_ATxmega192A3__
avrxtmega6	atxmega192a3u	__AVR_ATxmega192A3U__
avrxtmega6	atxmega192c3	__AVR_ATxmega192C3__
avrxtmega6	atxmega192d3	__AVR_ATxmega192D3__
avrxtmega6	atxmega256a3	__AVR_ATxmega256A3__
avrxtmega6	atxmega256a3u	__AVR_ATxmega256A3U__
avrxtmega6	atxmega256a3b	__AVR_ATxmega256A3B__
avrxtmega6	atxmega256a3bu	__AVR_ATxmega256A3BU__
avrxtmega6	atxmega256c3	__AVR_ATxmega256C3__
avrxtmega6	atxmega256d3	__AVR_ATxmega256D3__
avrxtmega6	atxmega384c3	__AVR_ATxmega384C3__
avrxtmega6	atxmega384d3	__AVR_ATxmega384D3__
avrxtmega7	atxmega128a1	__AVR_ATxmega128A1__
avrxtmega7	atxmega128a1u	__AVR_ATxmega128A1U__
avrxtmega7	atxmega128a4u	__AVR_ATxmega128A4U__
avrtiny10	attiny4	__AVR_ATtiny4__
avrtiny10	attiny5	__AVR_ATtiny5__
avrtiny10	attiny9	__AVR_ATtiny9__
avrtiny10	attiny10	__AVR_ATtiny10__
avrtiny10	attiny20	__AVR_ATtiny20__
avrtiny10	attiny40	__AVR_ATtiny40__

[1] 'avr25' architecture is new in GCC 4.2

[2] 'avr35' architecture is new in GCC 4.2.3

[3] 'avr31' and 'avr51' architectures is new in GCC 4.3

- `-morder1`
- `-morder2`

Change the order of register assignment. The default is

r24, r25, r18, r19, r20, r21, r22, r23, r30, r31, r26, r27, r28, r29, r17, r16, r15, r14, r13, r12, r11, r10, r9, r8, r7, r6, r5, r4, r3, r2, r0, r1

Order 1 uses

r18, r19, r20, r21, r22, r23, r24, r25, r30, r31, r26, r27, r28, r29, r17, r16, r15, r14, r13, r12, r11, r10, r9, r8, r7, r6, r5, r4, r3, r2, r0, r1

Order 2 uses

`r25, r24, r23, r22, r21, r20, r19, r18, r30, r31, r26, r27, r28, r29, r17, r16, r15, r14, r13, r12, r11, r10, r9, r8, r7, r6, r5, r4, r3, r2, r1, r0`

- `-mint8`

Assume `int` to be an 8-bit integer. Note that this is not really supported by `avr-libc`, so it should normally not be used. The default is to use 16-bit integers.

- `-mno-interrupts`

Generates code that changes the stack pointer without disabling interrupts. Normally, the state of the status register `SREG` is saved in a temporary register, interrupts are disabled while changing the stack pointer, and `SREG` is restored.

Specifying this option will define the preprocessor macro `__NO_INTERRUPTS__` to the value 1.

- `-mcall-prologues`

Use subroutines for function prologue/epilogue. For complex functions that use many registers (that needs to be saved/restored on function entry/exit), this saves some space at the cost of a slightly increased execution time.

- `-mtiny-stack`

Change only the low 8 bits of the stack pointer.

- `-mno-tablejump`

Deprecated, use `-fno-jump-tables` instead.

- `-mshort-calls`

Use `rjmp/rcall` (limited range) on `>8K` devices. On `avr2` and `avr4` architectures (less than 8 KB or flash memory), this is always the case. On `avr3` and `avr5` architectures, calls and jumps to targets outside the current function will by default use `jmp/call` instructions that can cover the entire address range, but that require more flash ROM and execution time.

- `-mrtl`

Dump the internal compilation result called "RTL" into comments in the generated assembler code. Used for debugging `avr-gcc`.

- `-msize`

Dump the address, size, and relative cost of each statement into comments in the generated assembler code. Used for debugging `avr-gcc`.

- `-mdeb`

Generate lots of debugging information to `stderr`.

13.1.2 Selected general compiler options

The following general gcc options might be of some interest to AVR users.

- `-On`

Optimization level *n*. Increasing *n* is meant to optimize more, an optimization level of 0 means no optimization at all, which is the default if no `-O` option is present. The special option `-Os` is meant to turn on all `-O2` optimizations that are not expected to increase code size.

Note that at `-O3`, gcc attempts to inline all "simple" functions. For the AVR target, this will normally constitute a large pessimization due to the code increasement. The only other optimization turned on with `-O3` is `-frename-registers`, which could rather be enabled manually instead.

A simple `-O` option is equivalent to `-O1`.

Note also that turning off all optimizations will prevent some warnings from being issued since the generation of those warnings depends on code analysis steps that are only performed when optimizing (unreachable code, unused variables).

See also the [appropriate FAQ entry](#) for issues regarding debugging optimized code.

- `-Wa, assembler-options`
- `-Wl, linker-options`

Pass the listed options to the assembler, or linker, respectively.

- `-g`

Generate debugging information that can be used by avr-gdb.

- `-ffreestanding`

Assume a "freestanding" environment as per the C standard. This turns off automatic builtin functions (though they can still be reached by prepending `__builtin_` to the actual function name). It also makes the compiler not complain when `main()` is declared with a `void` return type which makes some sense in a microcontroller environment where the application cannot meaningfully provide a return value to its environment (in most cases, `main()` won't even return anyway). However, this also turns off all optimizations normally done by the compiler which assume that functions known by a certain name behave as described by the standard. E. g., applying the function [strlen\(\)](#) to a literal string will normally cause the compiler to immediately replace that call by the actual length of the string, while with `-ffreestanding`, it will always call [strlen\(\)](#) at run-time.

- `-funsigned-char`

Make any unqualified `char` type an unsigned char. Without this option, they default to a signed char.

- `-funsigned-bitfields`

Make any unqualified bitfield type unsigned. By default, they are signed.

- `-fshort-enums`

Allocate to an `enum` type only as many bytes as it needs for the declared range of possible values. Specifically, the `enum` type will be equivalent to the smallest integer type which has enough room.

- `-fpack-struct`

Pack all structure members together without holes.

- `-fno-jump-tables`

Do not generate `tablejump` instructions. By default, jump tables can be used to optimize `switch` statements. When turned off, sequences of compare statements are used instead. Jump tables are usually faster to execute on average, but in particular for `switch` statements, where most of the jumps would go to the default label, they might waste a bit of flash memory.

NOTE: The `tablejump` instructions use the `LPM` assembler instruction for access to jump tables. Always use `-fno-jump-tables` switch, if compiling a bootloader for devices with more than 64 KB of code memory.

13.2 Options for the assembler `avr-as`

13.2.1 Machine-specific assembler options

- `-mmcu=architecture`
- `-mmcu=MCU name`

`avr-as` understands the same `-mmcu=` options as [avr-gcc](#). By default, `avr2` is assumed, but this can be altered by using the appropriate `.arch` pseudo-instruction inside the assembler source file.

- `-mall-opcodes`

Turns off opcode checking for the actual MCU type, and allows any possible AVR opcode to be assembled.

- `-mno-skip-bug`

Don't emit a warning when trying to skip a 2-word instruction with a `CPSE/SBIC/SBIS/SBRC/SBRS` instruction. Early AVR devices suffered from a hardware bug where these instructions could not be properly skipped.

- `-mno-wrap`

For `RJMP/RCALL` instructions, don't allow the target address to wrap around for devices that have more than 8 KB of memory.

- `--gstabs`

Generate `.stabs` debugging symbols for assembler source lines. This enables `avr-gdb` to trace through assembler source files. This option *must not* be used when assembling sources that have been generated by the C compiler; these files already contain the appropriate line number information from the C source files.

- `-a[cdhlmns=file]`

Turn on the assembler listing. The sub-options are:

- `c` omit false conditionals
- `d` omit debugging directives
- `h` include high-level source
- `l` include assembly
- `m` include macro expansions
- `n` omit forms processing
- `s` include symbols
- `=file` set the name of the listing file

The various sub-options can be combined into a single `-a` option list; `=file` must be the last one in that case.

13.2.2 Examples for assembler options passed through the C compiler

Remember that assembler options can be passed from the C compiler frontend using `-Wa` (see [above](#)), so in order to include the C source code into the assembler listing in file `foo.lst`, when compiling `foo.c`, the following compiler command-line can be used:

```
$ avr-gcc -c -O foo.c -o foo.o -Wa,-ahls=foo.lst
```

In order to pass an assembler file through the C preprocessor first, and have the assembler generate line number debugging information for it, the following command can be used:

```
$ avr-gcc -c -x assembler-with-cpp -o foo.o foo.S -Wa,--gstabs
```

Note that on Unix systems that have case-distinguishing file systems, specifying a file name with the suffix `.S` (upper-case letter `S`) will make the compiler automatically assume `-x assembler-with-cpp`, while using `.s` would pass the file directly to the assembler (no preprocessing done).

13.3 Controlling the linker `avr-ld`

13.3.1 Selected linker options

While there are no machine-specific options for `avr-ld`, a number of the standard options might be of interest to AVR users.

- `-lname`

Locate the archive library named `libname.a`, and use it to resolve currently unresolved symbols from it. The library is searched along a path that consists of builtin pathname entries that have been specified at compile time (e. g. `/usr/local/avr/lib` on Unix systems), possibly extended by pathname entries as specified by `-L` options (that must precede the `-l` options on the command-line).

- `-Lpath`

Additional location to look for archive libraries requested by `-l` options.

- `--defsym symbol=expr`

Define a global symbol *symbol* using *expr* as the value.

- `-M`

Print a linker map to `stdout`.

- `-Map mapfile`

Print a linker map to *mapfile*.

- `--cref`

Output a cross reference table to the map file (in case `-Map` is also present), or to `stdout`.

- `--section-start sectionname=org`

Start section *sectionname* at absolute address *org*.

- `-Tbss org`
- `-Tdata org`
- `-Ttext org`

Start the `bss`, `data`, or `text` section at `org`, respectively.

- `-T scriptfile`

Use *scriptfile* as the linker script, replacing the default linker script. Default linker scripts are stored in a system-specific location (e. g. under `/usr/local/avr/lib/ldscripts` on Unix systems), and consist of the AVR architecture name (`avr2` through `avr5`) with the suffix `.x` appended. They describe how the various [memory sections](#) will be linked together.

13.3.2 Passing linker options from the C compiler

By default, all unknown non-option arguments on the `avr-gcc` command-line (i. e., all filename arguments that don't have a suffix that is handled by `avr-gcc`) are passed straight to the linker. Thus, all files ending in `.o` (object files) and `.a` (object libraries) are provided to the linker.

System libraries are usually not passed by their explicit filename but rather using the `-l` option which uses an abbreviated form of the archive filename (see above). `avr-libc` ships two system libraries, `libc.a`, and `libm.a`. While the standard library `libc.a` will always be searched for unresolved references when the linker is started using the C compiler frontend (i. e., there's always at least one implied `-lc` option), the mathematics library `libm.a` needs to be explicitly requested using `-lm`. See also the [entry in the FAQ](#) explaining this.

Conventionally, Makefiles use the `make` macro `LDLIBS` to keep track of `-l` (and possibly `-L`) options that should only be appended to the C compiler command-line when linking the final binary. In contrast, the macro `LD_FLAGS` is used to store other command-line options to the C compiler that should be passed as options during the linking stage. The difference is that options are placed early on the command-line, while libraries are put at the end since they are to be used to resolve global symbols that are still unresolved at this point.

Specific linker flags can be passed from the C compiler command-line using the `-Wl` compiler option, see [above](#). This option requires that there be no spaces in the appended linker option, while some of the linker options above (like `-Map` or `--defsym`) would require a space. In these situations, the space can be replaced by an equal sign as well. For example, the following command-line can be used to compile `foo.c` into an executable, and also produce a link map that contains a cross-reference list in the file `foo.map`:

```
$ avr-gcc -O -o foo.out -Wl,-Map=foo.map -Wl,--cref foo.c
```

Alternatively, a comma as a placeholder will be replaced by a space before passing the option to the linker. So for a device with external SRAM, the following command-line would cause the linker to place the data segment at address `0x2000` in the SRAM:

```
$ avr-gcc -mmcu=atmega128 -o foo.out -Wl,-Tdata,0x802000
```

See the explanation of the [data section](#) for why 0x800000 needs to be added to the actual value. Note that the stack will still remain in internal RAM, through the symbol `__stack` that is provided by the run-time startup code. This is probably a good idea anyway (since internal RAM access is faster), and even required for some early devices that had hardware bugs preventing them from using a stack in external RAM. Note also that the heap for `malloc()` will still be placed after all the variables in the data section, so in this situation, no stack/heap collision can occur.

In order to relocate the stack from its default location at the top of internal RAM, the value of the symbol `__stack` can be changed on the linker command-line. As the linker is typically called from the compiler frontend, this can be achieved using a compiler option like

```
-Wl,--defsym=__stack=0x8003ff
```

The above will make the code use stack space from RAM address 0x3ff downwards. The amount of stack space available then depends on the bottom address of internal RAM for a particular device. It is the responsibility of the application to ensure the stack does not grow out of bounds, as well as to arrange for the stack to not collide with variable allocations made by the compiler (sections `.data` and `.bss`).

14 Compiler optimization

14.1 Problems with reordering code

Author

Jan Waclawek

Programs contain sequences of statements, and a naive compiler would execute them exactly in the order as they are written. But an optimizing compiler is free to *re-order* the statements - or even parts of them - if the resulting "net effect" is the same. The "measure" of the "net effect" is what the standard calls "side effects", and is accomplished exclusively through accesses (reads and writes) to variables qualified as `volatile`. So, as long as all volatile reads and writes are to the same addresses and in the same order (and writes write the same values), the program is correct, regardless of other operations in it. (One important point to note here is, that time duration between consecutive volatile accesses is not considered at all.)

Unfortunately, there are also operations which are not covered by volatile accesses. An example of this in `avr-gcc/avr-libc` are the `cli()` and `sei()` macros defined in `<avr/interrupt.h>`, which convert directly to the respective assembler mnemonics through the `__asm__()` statement. These don't constitute a variable access at all, not even volatile, so the compiler is free to move them around. Although there is a "volatile" qualifier which can be attached to the `__asm__()` statement, its effect on (re)ordering is not clear from the documentation (and is more likely only to prevent complete removal by the optimiser), as it (among other) states:

Note that even a volatile asm instruction can be moved relative to other code, including across jump instructions. [...] Similarly, you can't expect a sequence of volatile asm instructions to remain perfectly consecutive.

See also

<http://gcc.gnu.org/onlinedocs/gcc-4.3.4/gcc/Extended-Asm.html>

There is another mechanism which can be used to achieve something similar: *memory barriers*. This is accomplished through adding a special "memory" clobber to the inline asm statement, and ensures that all variables are flushed from registers to memory before the statement, and then re-read after the statement. The purpose of memory barriers is slightly different than to enforce code ordering: it is supposed to ensure that there are no variables "cached" in registers, so that it is safe to change the content of registers e.g. when switching context in a multitasking OS (on "big" processors with out-of-order execution they also imply usage of special instructions which force the processor into "in-order" state (this is not the case of AVR's)).

However, memory barrier works well in ensuring that all volatile accesses before and after the barrier occur in the given order with respect to the barrier. However, it does not ensure the compiler moving non-volatile-related statements across the barrier. Peter Dannegger provided a nice example of this effect:

```
#define cli() __asm volatile( "cli" ::: "memory" )
#define sei() __asm volatile( "sei" ::: "memory" )

unsigned int ivar;

void test2( unsigned int val )
{
    val = 65535U / val;

    cli();

    ivar = val;

    sei();
}
```

compiles with optimisations switched on (-Os) to

```
00000112 <test2>:
112: bc 01      movw r22, r24
114: f8 94      cli
116: 8f ef      ldi r24, 0xFF ; 255
118: 9f ef      ldi r25, 0xFF ; 255
11a: 0e 94 96 00 call 0x12c ; 0x12c <__udivmodhi4>
11e: 70 93 01 02 sts 0x0201, r23
122: 60 93 00 02 sts 0x0200, r22
126: 78 94      sei
128: 08 95      ret
```

where the potentially slow division is moved across `cli()`, resulting in interrupts to be disabled longer than intended. Note, that the volatile access occurs in order with respect to `cli()` or `sei()`; so the "net effect" required by the standard is achieved as intended, it

is "only" the timing which is off. However, for most of embedded applications, timing is an important, sometimes critical factor.

See also

<https://www.mikrocontroller.net/topic/65923>

Unfortunately, at the moment, in avr-gcc (nor in the C standard), there is no mechanism to enforce complete match of written and executed code ordering - except maybe of switching the optimization completely off (-O0), or writing all the critical code in assembly.

To sum it up:

- memory barriers ensure proper ordering of volatile accesses
- memory barriers don't ensure statements with no volatile accesses to be re-ordered across the barrier

15 Using the avrdude program

Note

This section was contributed by Brian Dean [bsd@bsdhome.com].

The avrdude program was previously called avrprog. The name was changed to avoid confusion with the avrprog program that Atmel ships with AvrStudio.

avrdude is a program that is used to update or read the flash and EEPROM memories of Atmel AVR microcontrollers on FreeBSD Unix. It supports the Atmel serial programming protocol using the PC's parallel port and can upload either a raw binary file or an Intel Hex format file. It can also be used in an interactive mode to individually update EEPROM cells, fuse bits, and/or lock bits (if their access is supported by the Atmel serial programming protocol.) The main flash instruction memory of the AVR can also be programmed in interactive mode, however this is not very useful because one can only turn bits off. The only way to turn flash bits on is to erase the entire memory (using avrdude's -e option).

avrdude is part of the FreeBSD ports system. To install it, simply do the following:

```
# cd /usr/ports/devel/avrdude
# make install
```

Once installed, avrdude can program processors using the contents of the .hex file specified on the command line. In this example, the file main.hex is burned into the flash memory:

```
# avrdude -p 2313 -e -m flash -i main.hex

avrdude: AVR device initialized and ready to accept instructions
```

```
avrdude: Device signature = 0x1e9101

avrdude: erasing chip
avrdude: done.
avrdude: reading input file "main.hex"
avrdude: input file main.hex auto detected as Intel Hex

avrdude: writing flash:
1749 0x00
avrdude: 1750 bytes of flash written
avrdude: verifying flash memory against main.hex:
avrdude: reading on-chip flash data:
1749 0x00
avrdude: verifying ...
avrdude: 1750 bytes of flash verified

avrdude done. Thank you.
```

The `-p 2313` option lets avrdude know that we are operating on an AT90S2313 chip. This option specifies the device id and is matched up with the device of the same id in avrdude's configuration file (`/usr/local/etc/avrdude.conf`). To list valid parts, specify the `-v` option. The `-e` option instructs avrdude to perform a chip-erase before programming; this is almost always necessary before programming the flash. The `-m flash` option indicates that we want to upload data into the flash memory, while `-i main.hex` specifies the name of the input file.

The EEPROM is uploaded in the same way, the only difference is that you would use `-m eeprom` instead of `-m flash`.

To use interactive mode, use the `-t` option:

```
# avrdude -p 2313 -t
avrdude: AVR device initialized and ready to accept instructions
avrdude: Device signature = 0x1e9101
avrdude>

The '?' command displays a list of valid
commands:

avrdude> ?
>>> ?
Valid commands:

dump      : dump memory : dump <memtype> <addr> <N-Bytes>
read      : alias for dump
write     : write memory : write <memtype> <addr> <b1> <b2> ... <bN>
erase     : perform a chip erase
sig       : display device signature bytes
part      : display the current part information
send      : send a raw command : send <b1> <b2> <b3> <b4>
help      : help
?         : help
quit      : quit
```

Use the 'part' command to display valid memory types for use with the 'dump' and 'write' commands.

```
avrdude>
```

16 Release Numbering and Methodology

16.1 Release Version Numbering Scheme

Release numbers consist of three parts, a major number, a minor number, and a revision number, each separated by a dot.

The major number is currently 2, to indicate the multilib layout has been adapted to the fairly different one used starting with AVR-GCC version 5. Nevertheless, it is still believed to be generally API-compatible with release versions 1.x.

In the past (up to 1.6.x), even minor numbers have been used to indicate "stable" releases, and odd minor numbers have been reserved for development branches/versions. As the latter has never really been used, and maintaining a stable branch that eventually became effectively the same as the development version has proven to be just a cumbersome and tedious job, this scheme has given up in early 2010, so starting with 1.7.0, every minor number will be used. Minor numbers will be bumped upon judgement of the development team, whenever it seems appropriate, but at least in cases where some API was changed.

Starting with version 1.4.0, a file `<avr/version.h>` indicates the library version of an installed library tree.

16.2 Releasing AVR Libc

The information in this section is only relevant to AVR Libc developers and can be ignored by end users.

Note

In what follows, I assume you know how to use SVN and how to checkout multiple source trees in a single directory without having them clobber each other. If you don't know how to do this, you probably shouldn't be making releases or cutting branches.

16.2.1 Creating an SVN branch

The following steps should be taken to cut a branch in SVN (assuming \$username is set to your savannah username):

1. Check out a fresh source tree from SVN trunk.
2. Update the NEWS file with pending release number and commit to SVN trunk:
Change *Changes since avr-libc-<last_release>*: to *Changes in avr-libc-<this_release>*.
3. Set the branch-point tag (setting <major> and <minor> accordingly):

```
svn copy svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/trunk  
svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/tags/avr-libc-<major>_  
<minor>-branchpoint
```

4. Create the branch:

```
svn copy svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/trunk
svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/branches/avr-libc-<major>
<minor>-branch
```

5. Update the package version in configure.ac and commit configure.ac to SVN trunk:

Change minor number to next odd value.

6. Update the NEWS file and commit to SVN trunk:

Add *Changes since avr-libc-<this_release>*:

7. Check out a new tree for the branch:

```
svn co svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/branches/avr-libc-<major>
<minor>-branch
```

8. Update the package version in configure.ac and commit configure.ac to SVN branch:

Change the patch number to 90 to denote that this now a branch leading up to a release. Be sure to leave the <date> part of the version.

9. Bring the build system up to date by running bootstrap and configure.

10. Perform a 'make distcheck' and make sure it succeeds. This will create the snapshot source tarball. This should be considered the first release candidate.

11. Upload the snapshot tarball to savannah.

12. Update the bug tracker interface on Savannah: Bugs —> Edit field values —> Release / Fixed Release

13. Announce the branch and the branch tag to the avr-libc-dev list so other developers can checkout the branch.

16.2.2 Making a release

A stable release will only be done on a branch, not from the SVN trunk.

The following steps should be taken when making a release:

1. Make sure the source tree you are working from is on the correct branch:

```
svn switch svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/branches/avr-libc-<major>
<minor>-branch
```

2. Update the package version in configure.ac and commit it to SVN.

3. Update the gnu tool chain version requirements in the README and commit to SVN.

4. Update the ChangeLog file to note the release and commit to SVN on the branch:

Add *Released avr-libc-<this_release>*.

5. Update the NEWS file with pending release number and commit to SVN:

Change *Changes since avr-libc-<last_release>:* to *Changes in avr-libc-<this_release>:*.

6. Bring the build system up to date by running bootstrap and configure.
7. Perform a 'make distcheck' and make sure it succeeds. This will create the source tarball.
8. Tag the release:

```
svn copy . svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/tags/avr-libc-  
<minor>_<patch>-release
```

or

```
svn copy svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/branches/avr-libc-  
<minor>-branch svn+ssh://$username@svn.savannah.nongnu.org/avr-libc/tags/avr-libc-  
<minor>_<patch>-release
```

9. Upload the tarball to savannah.
10. Update the NEWS file, and commit to SVN:
Add *Changes since avr-libc-<major>_<minor>_<patch>:*
11. Update the bug tracker interface on Savannah: Bugs —> Edit field values —> Release / Fixed Release
12. Generate the latest documentation and upload to savannah.
13. Announce the release.

The following hypothetical diagram should help clarify version and branch relationships.




Figure 4: Release tree

17 Acknowledgments

This document tries to tie together the labors of a large group of people. Without these individuals' efforts, we wouldn't have a terrific, *free* set of tools to develop AVR projects. We all owe thanks to:

- The GCC Team, which produced a very capable set of development tools for an amazing number of platforms and processors.
- Denis Chertykov [denisc@overta.ru] for making the AVR-specific changes to the GNU tools.
- Denis Chertykov and Marek Michalkiewicz [marekm@linux.org.pl] for developing the standard libraries and startup code for **AVR-GCC**.
- Uros Platise for developing the AVR programmer tool, **uisp**.

- Joerg Wunsch [joerg@FreeBSD.ORG] for adding all the AVR development tools to the FreeBSD [<http://www.freebsd.org>] ports tree and for providing the basics for the [demo project](#).
- Brian Dean [bsd@bsdhome.com] for developing **avrdude** (an alternative to **uisp**) and for contributing [documentation](#) which describes how to use it. **Avr-dude** was previously called **avrprog**.
- Eric Weddington [eweddington@cs0.atmel.com] for maintaining the **WinAVR** package and thus making the continued improvements to the open source AVR toolchain available to many users.
- Rich Neswold for writing the original avr-tools document (which he graciously allowed to be merged into this document) and his improvements to the [demo project](#).
- Theodore A. Roth for having been a long-time maintainer of many of the tools (**AVR-Libc**, the AVR port of **GDB**, **AVaRICE**, **uisp**, **avrdude**).
- All the people who currently maintain the tools, and/or have submitted suggestions, patches and bug reports. (See the AUTHORS files of the various tools.)
- And lastly, all the users who use the software. If nobody used the software, we would probably not be very motivated to continue to develop it. Keep those bug reports coming. ;-)

18 Todo List

Group [avr_boot](#) From email with Marek: On smaller devices (all except AT-mega64/128), `__SPM_REG` is in the I/O space, accessible with the shorter "in" and "out" instructions - since the boot loader has a limited size, this could be an important optimization.

19 Deprecated List

Global [cbi](#)(port, bit)

Global [enable_external_int](#)(mask)

Global [inb](#)(port)

Global [inp](#)(port)

Global **[INTERRUPT](#)**(signame)

Global **[ISR_ALIAS](#)**(vector, target_vector) For new code, the use of [ISR\(..., ISR_ALIASOF\(...\)\)](#) is recommended.

Global **[outb](#)**(port, val)

Global **[outp](#)**(val, port)

Global **[sbi](#)**(port, bit)

Global **[SIGNAL](#)**(vector) Do not use [SIGNAL\(\)](#) in new code. Use [ISR\(\)](#) instead.

Global **[timer_enable_int](#)**(unsigned char ints)

20 Module Index

20.1 Modules

Here is a list of all modules:

<alloca.h> : Allocate space in the stack	128
<assert.h> : Diagnostics	129
<ctype.h> : Character Operations	130
<errno.h> : System Errors	132
<inttypes.h> : Integer Type conversions	133
<math.h> : Mathematics	147
<setjmp.h> : Non-local goto	160
<stdint.h> : Standard Integer Types	162
<stdio.h> : Standard IO facilities	174
<stdlib.h> : General utilities	192
<string.h> : Strings	203
<time.h> : Time	216

<avr/boot.h>: Bootloader Support Utilities	225
<avr/cpufunc.h>: Special AVR CPU functions	231
<avr/eeprom.h>: EEPROM handling	232
<avr/fuse.h>: Fuse Support	237
<avr/interrupt.h>: Interrupts	240
<avr/io.h>: AVR device-specific IO definitions	260
<avr/lock.h>: Lockbit Support	261
<avr/pgmspace.h>: Program Space Utilities	264
<avr/power.h>: Power Reduction Management	288
<avr/sfr_defs.h>: Special function registers	293
Additional notes from <avr/sfr_defs.h>	292
<avr/signature.h>: Signature Support	295
<avr/sleep.h>: Power Management and Sleep Modes	296
<avr/version.h>: avr-libc version macros	298
<avr/wdt.h>: Watchdog timer handling	299
<util/atomic.h> Atomically and Non-Atomically Executed Code Blocks	303
<util/crc16.h>: CRC Computations	306
<util/delay.h>: Convenience functions for busy-wait delay loops	310
<util/delay_basic.h>: Basic busy-wait delay loops	313
<util/parity.h>: Parity bit generation	314
<util/setbaud.h>: Helper macros for baud rate calculations	314
<util/twi.h>: TWI bit mask definitions	317
<compat/deprecated.h>: Deprecated items	321
<compat/ina90.h>: Compatibility with IAR EWB 3.x	324
Demo projects	325
Combining C and assembly source files	326
A simple project	329

A more sophisticated project	343
Using the standard IO facilities	350
Example using the two-wire interface (TWI)	357

21 Data Structure Index

21.1 Data Structures

Here are the data structures with brief descriptions:

div_t	362
ldiv_t	363
tm	363
week_date	365

22 File Index

22.1 File List

Here is a list of all documented files with brief descriptions:

alloca.h	??
assert.h	365
atoi.S	366
atol.S	366
atomic.h	366
boot.h	366
cpufunc.h	367
crc16.h	367
ctype.h	367
defines.h	??
delay.h	368
delay_basic.h	368

deprecated.h	??
dtoa_conv.h	??
eedef.h	??
eprom.h	??
ephemera_common.h	??
errno.h	368
eu_dst.h	??
fdevopen.c	369
fuse.h	369
hd44780.h	??
ina90.h	??
interrupt.h	369
inttypes.h	370
io.h	372
iocompat.h	??
lcd.h	??
lock.h	372
math.h	372
parity.h	375
pgmspace.h	375
portpins.h	??
power.h	377
project.h	??
setbaud.h	378
setjmp.h	378
sfr_defs.h	??
signal.h	??

signature.h	378
sleep.h	378
stdint.h	378
stdio.h	381
stdio_private.h	??
stdlib.h	383
stdlib_private.h	??
string.h	384
time.h	385
util/twi.h	387
compat/twi.h	??
uart.h	??
usa_dst.h	??
version.h	??
wdt.h	388
xmega.h	??
xtoa_fast.h	??

23 Module Documentation

23.1 <alloca.h>: Allocate space in the stack

Functions

- void * [alloca](#) (size_t __size)

23.1.1 Detailed Description

23.1.2 Function Documentation

23.1.2.1 void* [alloca](#) (size_t __size)

Allocate *__size* bytes of space in the stack frame of the caller.

This temporary space is automatically freed when the function that called `alloca()` returns to its caller. Avr-libc defines the `alloca()` as a macro, which is translated into the inlined `__builtin_alloca()` function. The fact that the code is inlined, means that it is impossible to take the address of this function, or to change its behaviour by linking with a different library.

Returns

`alloca()` returns a pointer to the beginning of the allocated space. If the allocation causes stack overflow, program behaviour is undefined.

Warning

Avoid use `alloca()` inside the list of arguments of a function call.

23.2 <assert.h>: Diagnostics

Defines

- #define `assert(expression)`

23.2.1 Detailed Description

```
#include <assert.h>
```

This header file defines a debugging aid.

As there is no standard error output stream available for many applications using this library, the generation of a printable error message is not enabled by default. These messages will only be generated if the application defines the macro

```
__ASSERT_USE_STDERR
```

before including the `<assert.h>` header file. By default, only `abort()` will be called to halt the application.

23.2.2 Define Documentation

23.2.2.1 #define `assert(expression)`

Parameters

expression Expression to test for.

The `assert()` macro tests the given expression and if it is false, the calling process is terminated. A diagnostic message is written to `stderr` and the function `abort()` is called, effectively terminating the program.

If expression is true, the `assert()` macro does nothing.

The `assert()` macro may be removed at compile time by defining `NDEBUG` as a macro (e.g., by using the compiler option `-DNDEBUG`).

23.3 <ctype.h>: Character Operations

Character classification routines

These functions perform character classification. They return true or false status depending whether the character passed to the function falls into the function's classification (i.e. `isdigit()` returns true if its argument is any value '0' through '9', inclusive). If the input is not an unsigned char value, all of these functions return false.

- int `isalnum` (int __c)
- int `isalpha` (int __c)
- int `isascii` (int __c)
- int `isblank` (int __c)
- int `isctrl` (int __c)
- int `isdigit` (int __c)
- int `isgraph` (int __c)
- int `islower` (int __c)
- int `isprint` (int __c)
- int `ispunct` (int __c)
- int `isspace` (int __c)
- int `isupper` (int __c)
- int `isxdigit` (int __c)

Character conversion routines

This realization permits all possible values of integer argument. The `toascii()` function clears all highest bits. The `tolower()` and `toupper()` functions return an input argument as is, if it is not an unsigned char value.

- int `toascii` (int __c)
- int `tolower` (int __c)
- int `toupper` (int __c)

23.3.1 Detailed Description

These functions perform various operations on characters.

```
#include <ctype.h>
```

23.3.2 Function Documentation

23.3.2.1 `int isalnum (int __c)`

Checks for an alphanumeric character. It is equivalent to `(isalpha(c) || isdigit(c))`.

23.3.2.2 `int isalpha (int __c)`

Checks for an alphabetic character. It is equivalent to `(isupper(c) || islower(c))`.

23.3.2.3 `int isascii (int __c)`

Checks whether `c` is a 7-bit unsigned char value that fits into the ASCII character set.

23.3.2.4 `int isblank (int __c)`

Checks for a blank character, that is, a space or a tab.

23.3.2.5 `int iscntrl (int __c)`

Checks for a control character.

23.3.2.6 `int isdigit (int __c)`

Checks for a digit (0 through 9).

23.3.2.7 `int isgraph (int __c)`

Checks for any printable character except space.

23.3.2.8 `int islower (int __c)`

Checks for a lower-case character.

23.3.2.9 `int isprint (int __c)`

Checks for any printable character including space.

23.3.2.10 int ispunct (int __c)

Checks for any printable character which is not a space or an alphanumeric character.

23.3.2.11 int isspace (int __c)

Checks for white-space characters. For the avr-libc library, these are: space, form-feed ('`\f`'), newline ('`\n`'), carriage return ('`\r`'), horizontal tab ('`\t`'), and vertical tab ('`\v`').

23.3.2.12 int isupper (int __c)

Checks for an uppercase letter.

23.3.2.13 int isxdigit (int __c)

Checks for a hexadecimal digits, i.e. one of 0 1 2 3 4 5 6 7 8 9 a b c d e f A B C D E F.

23.3.2.14 int toascii (int __c)

Converts `c` to a 7-bit unsigned char value that fits into the ASCII character set, by clearing the high-order bits.

Warning

Many people will be unhappy if you use this function. This function will convert accented letters into random characters.

23.3.2.15 int tolower (int __c)

Converts the letter `c` to lower case, if possible.

23.3.2.16 int toupper (int __c)

Converts the letter `c` to upper case, if possible.

23.4 <errno.h>: System Errors**Defines**

- #define [EDOM](#) 33
- #define [ERANGE](#) 34

Variables

- int `errno`

23.4.1 Detailed Description

```
#include <errno.h>
```

Some functions in the library set the global variable `errno` when an error occurs. The file, <errno.h>, provides symbolic names for various error codes.

23.4.2 Define Documentation

23.4.2.1 #define EDOM 33

Domain error.

23.4.2.2 #define ERANGE 34

Range error.

23.4.3 Variable Documentation

23.4.3.1 int `errno`

Error code for last error encountered by library.

The variable `errno` holds the last error code encountered by a library function. This variable must be cleared by the user prior to calling a library function.

Warning

The `errno` global variable is not safe to use in a threaded or multi-task system. A race condition can occur if a task is interrupted between the call which sets `error` and when the task examines `errno`. If another task changes `errno` during this time, the result will be incorrect for the interrupted task.

23.5 <inttypes.h>: Integer Type conversions

Far pointers for memory access >64K

- typedef `int32_t int_farptr_t`
- typedef `uint32_t uint_farptr_t`

macros for printf and scanf format specifiers

For C++, these are only included if `__STDC_LIMIT_MACROS` is defined before including `<inttypes.h>`.

- `#define PRId8 "d"`
- `#define PRIdLEAST8 "d"`
- `#define PRIdFAST8 "d"`
- `#define PRIi8 "i"`
- `#define PRIiLEAST8 "i"`
- `#define PRIiFAST8 "i"`
- `#define PRId16 "d"`
- `#define PRIdLEAST16 "d"`
- `#define PRIdFAST16 "d"`
- `#define PRIi16 "i"`
- `#define PRIiLEAST16 "i"`
- `#define PRIiFAST16 "i"`
- `#define PRId32 "ld"`
- `#define PRIdLEAST32 "ld"`
- `#define PRIdFAST32 "ld"`
- `#define PRIi32 "li"`
- `#define PRIiLEAST32 "li"`
- `#define PRIiFAST32 "li"`
- `#define PRIdPTR PRId16`
- `#define PRIiPTR PRIi16`
- `#define PRIo8 "o"`
- `#define PRIoLEAST8 "o"`
- `#define PRIoFAST8 "o"`
- `#define PRIu8 "u"`
- `#define PRIuLEAST8 "u"`
- `#define PRIuFAST8 "u"`
- `#define PRIx8 "x"`
- `#define PRIxLEAST8 "x"`
- `#define PRIxFAST8 "x"`
- `#define PRIX8 "X"`
- `#define PRIXLEAST8 "X"`
- `#define PRIXFAST8 "X"`
- `#define PRIo16 "o"`
- `#define PRIoLEAST16 "o"`
- `#define PRIoFAST16 "o"`
- `#define PRIu16 "u"`
- `#define PRIuLEAST16 "u"`
- `#define PRIuFAST16 "u"`
- `#define PRIx16 "x"`
- `#define PRIxLEAST16 "x"`
- `#define PRIxFAST16 "x"`
- `#define PRIX16 "X"`

- #define PRXLEAST16 "X"
- #define PRXFAST16 "X"
- #define PRIo32 "lo"
- #define PRIoLEAST32 "lo"
- #define PRIoFAST32 "lo"
- #define PRlu32 "lu"
- #define PRluLEAST32 "lu"
- #define PRluFAST32 "lu"
- #define PRlx32 "lx"
- #define PRlxLEAST32 "lx"
- #define PRlxFAST32 "lx"
- #define PRIX32 "IX"
- #define PRXLEAST32 "IX"
- #define PRXFAST32 "IX"
- #define PRIoPTR PRIo16
- #define PRluPTR PRlu16
- #define PRlxPTR PRlx16
- #define PRIXPTR PRIX16
- #define SCNd8 "hhd"
- #define SCNdLEAST8 "hhd"
- #define SCNdFAST8 "hhd"
- #define SCNi8 "hhi"
- #define SCNiLEAST8 "hhi"
- #define SCNiFAST8 "hhi"
- #define SCNd16 "d"
- #define SCNdLEAST16 "d"
- #define SCNdFAST16 "d"
- #define SCNi16 "i"
- #define SCNiLEAST16 "i"
- #define SCNiFAST16 "i"
- #define SCNd32 "ld"
- #define SCNdLEAST32 "ld"
- #define SCNdFAST32 "ld"
- #define SCNi32 "li"
- #define SCNiLEAST32 "li"
- #define SCNiFAST32 "li"
- #define SCNdPTR SCNd16
- #define SCNiPTR SCNi16
- #define SCNo8 "hho"
- #define SCNoLEAST8 "hho"
- #define SCNoFAST8 "hho"
- #define SCNu8 "hhu"
- #define SCNuLEAST8 "hhu"
- #define SCNuFAST8 "hhu"
- #define SCNx8 "hhx"
- #define SCNxLEAST8 "hhx"

- #define [SCNxFAST8](#) "hhx"
- #define [SCNo16](#) "o"
- #define [SCNoLEAST16](#) "o"
- #define [SCNoFAST16](#) "o"
- #define [SCNu16](#) "u"
- #define [SCNuLEAST16](#) "u"
- #define [SCNuFAST16](#) "u"
- #define [SCNx16](#) "x"
- #define [SCNxLEAST16](#) "x"
- #define [SCNxFAST16](#) "x"
- #define [SCNo32](#) "lo"
- #define [SCNoLEAST32](#) "lo"
- #define [SCNoFAST32](#) "lo"
- #define [SCNu32](#) "lu"
- #define [SCNuLEAST32](#) "lu"
- #define [SCNuFAST32](#) "lu"
- #define [SCNx32](#) "lx"
- #define [SCNxLEAST32](#) "lx"
- #define [SCNxFAST32](#) "lx"
- #define [SCNoPTR](#) SCNo16
- #define [SCNuPTR](#) SCNu16
- #define [SCNxPTR](#) SCNx16

23.5.1 Detailed Description

```
#include <inttypes.h>
```

This header file includes the exact-width integer definitions from [<stdint.h>](#), and extends them with additional facilities provided by the implementation.

Currently, the extensions include two additional integer types that could hold a "far" pointer (i.e. a code pointer that can address more than 64 KB), as well as standard names for all printf and scanf formatting options that are supported by the [<stdio.h>: Standard IO facilities](#). As the library does not support the full range of conversion specifiers from ISO 9899:1999, only those conversions that are actually implemented will be listed here.

The idea behind these conversion macros is that, for each of the types defined by [<stdint.h>](#), a macro will be supplied that portably allows formatting an object of that type in [printf\(\)](#) or [scanf\(\)](#) operations. Example:

```
#include <inttypes.h>

uint8_t smallval;
int32_t longval;
...
printf("The hexadecimal value of smallval is %" PRIx8
      ", the decimal value of longval is %" PRId32 ".\n",
      smallval, longval);
```

23.5.2 Define Documentation

23.5.2.1 #define PRId16 "d"

decimal printf format for int16_t

23.5.2.2 #define PRId32 "ld"

decimal printf format for int32_t

23.5.2.3 #define PRId8 "d"

decimal printf format for int8_t

23.5.2.4 #define PRIdFAST16 "d"

decimal printf format for int_fast16_t

23.5.2.5 #define PRIdFAST32 "ld"

decimal printf format for int_fast32_t

23.5.2.6 #define PRIdFAST8 "d"

decimal printf format for int_fast8_t

23.5.2.7 #define PRIdLEAST16 "d"

decimal printf format for int_least16_t

23.5.2.8 #define PRIdLEAST32 "ld"

decimal printf format for int_least32_t

23.5.2.9 #define PRIdLEAST8 "d"

decimal printf format for int_least8_t

23.5.2.10 #define PRIdPTR PRId16

decimal printf format for intptr_t

23.5.2.11 `#define PRIi16 "i"`

integer printf format for int16_t

23.5.2.12 `#define PRIi32 "li"`

integer printf format for int32_t

23.5.2.13 `#define PRIi8 "i"`

integer printf format for int8_t

23.5.2.14 `#define PRIiFAST16 "i"`

integer printf format for int_fast16_t

23.5.2.15 `#define PRIiFAST32 "li"`

integer printf format for int_fast32_t

23.5.2.16 `#define PRIiFAST8 "i"`

integer printf format for int_fast8_t

23.5.2.17 `#define PRIiLEAST16 "i"`

integer printf format for int_least16_t

23.5.2.18 `#define PRIiLEAST32 "li"`

integer printf format for int_least32_t

23.5.2.19 `#define PRIiLEAST8 "i"`

integer printf format for int_least8_t

23.5.2.20 `#define PRIiPTR PRIi16`

integer printf format for intptr_t

23.5.2.21 `#define PRIo16 "o"`

octal printf format for uint16_t

23.5.2.22 #define PRIo32 "lo"

octal printf format for uint32_t

23.5.2.23 #define PRIo8 "o"

octal printf format for uint8_t

23.5.2.24 #define PRIoFAST16 "o"

octal printf format for uint_fast16_t

23.5.2.25 #define PRIoFAST32 "lo"

octal printf format for uint_fast32_t

23.5.2.26 #define PRIoFAST8 "o"

octal printf format for uint_fast8_t

23.5.2.27 #define PRIoLEAST16 "o"

octal printf format for uint_least16_t

23.5.2.28 #define PRIoLEAST32 "lo"

octal printf format for uint_least32_t

23.5.2.29 #define PRIoLEAST8 "o"

octal printf format for uint_least8_t

23.5.2.30 #define PRIoPTR PRIo16

octal printf format for uintptr_t

23.5.2.31 #define PRIu16 "u"

decimal printf format for uint16_t

23.5.2.32 #define PRIu32 "lu"

decimal printf format for uint32_t

23.5.2.33 `#define PRIu8 "u"`

decimal printf format for `uint8_t`

23.5.2.34 `#define PRIuFAST16 "u"`

decimal printf format for `uint_fast16_t`

23.5.2.35 `#define PRIuFAST32 "lu"`

decimal printf format for `uint_fast32_t`

23.5.2.36 `#define PRIuFAST8 "u"`

decimal printf format for `uint_fast8_t`

23.5.2.37 `#define PRIuLEAST16 "u"`

decimal printf format for `uint_least16_t`

23.5.2.38 `#define PRIuLEAST32 "lu"`

decimal printf format for `uint_least32_t`

23.5.2.39 `#define PRIuLEAST8 "u"`

decimal printf format for `uint_least8_t`

23.5.2.40 `#define PRIuPTR PRIu16`

decimal printf format for `uintptr_t`

23.5.2.41 `#define PRIx16 "X"`

uppercase hexadecimal printf format for `uint16_t`

23.5.2.42 `#define PRIx16 "x"`

hexadecimal printf format for `uint16_t`

23.5.2.43 `#define PRIx32 "IX"`

uppercase hexadecimal printf format for `uint32_t`

23.5.2.44 `#define PRIx32 "lx"`

hexadecimal printf format for uint32_t

23.5.2.45 `#define PRIx8 "X"`

uppercase hexadecimal printf format for uint8_t

23.5.2.46 `#define PRIx8 "x"`

hexadecimal printf format for uint8_t

23.5.2.47 `#define PRIxFAST16 "X"`

uppercase hexadecimal printf format for uint_fast16_t

23.5.2.48 `#define PRIxFAST16 "x"`

hexadecimal printf format for uint_fast16_t

23.5.2.49 `#define PRIxFAST32 "IX"`

uppercase hexadecimal printf format for uint_fast32_t

23.5.2.50 `#define PRIxFAST32 "lx"`

hexadecimal printf format for uint_fast32_t

23.5.2.51 `#define PRIxFAST8 "X"`

uppercase hexadecimal printf format for uint_fast8_t

23.5.2.52 `#define PRIxFAST8 "x"`

hexadecimal printf format for uint_fast8_t

23.5.2.53 `#define PRIxLEAST16 "X"`

uppercase hexadecimal printf format for uint_least16_t

23.5.2.54 `#define PRIxLEAST16 "x"`

hexadecimal printf format for uint_least16_t

23.5.2.55 #define PRIxLEAST32 "IX"

uppercase hexadecimal printf format for uint_least32_t

23.5.2.56 #define PRIxLEAST32 "lx"

hexadecimal printf format for uint_least32_t

23.5.2.57 #define PRIxLEAST8 "X"

uppercase hexadecimal printf format for uint_least8_t

23.5.2.58 #define PRIxLEAST8 "x"

hexadecimal printf format for uint_least8_t

23.5.2.59 #define PRIXPTR PRIx16

uppercase hexadecimal printf format for uintptr_t

23.5.2.60 #define PRIXPTR PRIx16

hexadecimal printf format for uintptr_t

23.5.2.61 #define SCNd16 "d"

decimal scanf format for int16_t

23.5.2.62 #define SCNd32 "ld"

decimal scanf format for int32_t

23.5.2.63 #define SCNd8 "hhd"

decimal scanf format for int8_t

23.5.2.64 #define SCNdFAST16 "d"

decimal scanf format for int_fast16_t

23.5.2.65 #define SCNdFAST32 "ld"

decimal scanf format for int_fast32_t

23.5.2.66 `#define SCNdFAST8 "hhd"`

decimal scanf format for `int_fast8_t`

23.5.2.67 `#define SCNdLEAST16 "d"`

decimal scanf format for `int_least16_t`

23.5.2.68 `#define SCNdLEAST32 "ld"`

decimal scanf format for `int_least32_t`

23.5.2.69 `#define SCNdLEAST8 "hhd"`

decimal scanf format for `int_least8_t`

23.5.2.70 `#define SCNdPTR SCNd16`

decimal scanf format for `intptr_t`

23.5.2.71 `#define SCNi16 "i"`

generic-integer scanf format for `int16_t`

23.5.2.72 `#define SCNi32 "li"`

generic-integer scanf format for `int32_t`

23.5.2.73 `#define SCNi8 "hhi"`

generic-integer scanf format for `int8_t`

23.5.2.74 `#define SCNiFAST16 "i"`

generic-integer scanf format for `int_fast16_t`

23.5.2.75 `#define SCNiFAST32 "li"`

generic-integer scanf format for `int_fast32_t`

23.5.2.76 `#define SCNiFAST8 "hhi"`

generic-integer scanf format for `int_fast8_t`

23.5.2.77 **#define SCNiLEAST16 "i"**

generic-integer scanf format for int_least16_t

23.5.2.78 **#define SCNiLEAST32 "li"**

generic-integer scanf format for int_least32_t

23.5.2.79 **#define SCNiLEAST8 "hi"**

generic-integer scanf format for int_least8_t

23.5.2.80 **#define SCNiPTR SCNi16**

generic-integer scanf format for intptr_t

23.5.2.81 **#define SCNo16 "o"**

octal scanf format for uint16_t

23.5.2.82 **#define SCNo32 "lo"**

octal scanf format for uint32_t

23.5.2.83 **#define SCNo8 "hho"**

octal scanf format for uint8_t

23.5.2.84 **#define SCNoFAST16 "o"**

octal scanf format for uint_fast16_t

23.5.2.85 **#define SCNoFAST32 "lo"**

octal scanf format for uint_fast32_t

23.5.2.86 **#define SCNoFAST8 "hho"**

octal scanf format for uint_fast8_t

23.5.2.87 **#define SCNoLEAST16 "o"**

octal scanf format for uint_least16_t

23.5.2.88 **#define SCNoLEAST32 "lo"**

octal scanf format for uint_least32_t

23.5.2.89 **#define SCNoLEAST8 "hho"**

octal scanf format for uint_least8_t

23.5.2.90 **#define SCNoPTR SCNo16**

octal scanf format for uintptr_t

23.5.2.91 **#define SCNu16 "u"**

decimal scanf format for uint16_t

23.5.2.92 **#define SCNu32 "lu"**

decimal scanf format for uint32_t

23.5.2.93 **#define SCNu8 "hhu"**

decimal scanf format for uint8_t

23.5.2.94 **#define SCNuFAST16 "u"**

decimal scanf format for uint_fast16_t

23.5.2.95 **#define SCNuFAST32 "lu"**

decimal scanf format for uint_fast32_t

23.5.2.96 **#define SCNuFAST8 "hhu"**

decimal scanf format for uint_fast8_t

23.5.2.97 **#define SCNuLEAST16 "u"**

decimal scanf format for uint_least16_t

23.5.2.98 **#define SCNuLEAST32 "lu"**

decimal scanf format for uint_least32_t

23.5.2.99 #define SCNuLEAST8 "hhu"

decimal scanf format for uint_least8_t

23.5.2.100 #define SCNuPTR SCNu16

decimal scanf format for uintptr_t

23.5.2.101 #define SCNx16 "x"

hexadecimal scanf format for uint16_t

23.5.2.102 #define SCNx32 "lx"

hexadecimal scanf format for uint32_t

23.5.2.103 #define SCNx8 "hxx"

hexadecimal scanf format for uint8_t

23.5.2.104 #define SCNxFAST16 "x"

hexadecimal scanf format for uint_fast16_t

23.5.2.105 #define SCNxFAST32 "lx"

hexadecimal scanf format for uint_fast32_t

23.5.2.106 #define SCNxFAST8 "hxx"

hexadecimal scanf format for uint_fast8_t

23.5.2.107 #define SCNxLEAST16 "x"

hexadecimal scanf format for uint_least16_t

23.5.2.108 #define SCNxLEAST32 "lx"

hexadecimal scanf format for uint_least32_t

23.5.2.109 #define SCNxLEAST8 "hxx"

hexadecimal scanf format for uint_least8_t

23.5.2.110 #define SCNxPTR SCNx16

hexadecimal scanf format for uintptr_t

23.5.3 Typedef Documentation**23.5.3.1 typedef int32_t int_farptr_t**

signed integer type that can hold a pointer > 64 KB

23.5.3.2 typedef uint32_t uint_farptr_t

unsigned integer type that can hold a pointer > 64 KB

23.6 <math.h>: Mathematics**Defines**

- #define [M_E](#) 2.7182818284590452354
- #define [M_LOG2E](#) 1.4426950408889634074
- #define [M_LOG10E](#) 0.43429448190325182765
- #define [M_LN2](#) 0.69314718055994530942
- #define [M_LN10](#) 2.30258509299404568402
- #define [M_PI](#) 3.14159265358979323846
- #define [M_PI_2](#) 1.57079632679489661923
- #define [M_PI_4](#) 0.78539816339744830962
- #define [M_1_PI](#) 0.31830988618379067154
- #define [M_2_PI](#) 0.63661977236758134308
- #define [M_2_SQRTPI](#) 1.12837916709551257390
- #define [M_SQRT2](#) 1.41421356237309504880
- #define [M_SQRT1_2](#) 0.70710678118654752440
- #define [NAN](#) __builtin_nan("")
- #define [INFINITY](#) __builtin_inf()
- #define [cosf](#) cos
- #define [sinf](#) sin
- #define [tanf](#) tan
- #define [fabsf](#) fabs
- #define [fmodf](#) fmod
- #define [cbrtf](#) cbrt
- #define [hypotf](#) hypot
- #define [squaref](#) square
- #define [floorf](#) floor
- #define [ceilf](#) ceil
- #define [frexpf](#) frexp
- #define [ldexpf](#) ldexp

- #define `expf` `exp`
- #define `coshf` `cosh`
- #define `sinhf` `sinh`
- #define `tanhf` `tanh`
- #define `acosf` `acos`
- #define `asinf` `asin`
- #define `atanf` `atan`
- #define `atan2f` `atan2`
- #define `logf` `log`
- #define `log10f` `log10`
- #define `powf` `pow`
- #define `isnanf` `isnan`
- #define `isinf` `isinf`
- #define `isfinitef` `isfinite`
- #define `copysignf` `copysign`
- #define `signbitf` `signbit`
- #define `fdimf` `fdim`
- #define `fmaf` `fma`
- #define `fmaxf` `fmax`
- #define `fminf` `fmin`
- #define `truncf` `trunc`
- #define `roundf` `round`
- #define `lroundf` `lround`
- #define `rintf` `rint`

Functions

- double `cos` (double __x)
- double `sin` (double __x)
- double `tan` (double __x)
- double `fabs` (double __x)
- double `fmod` (double __x, double __y)
- double `modf` (double __x, double *__iptr)
- float `modff` (float __x, float *__iptr)
- double `sqrt` (double __x)
- float `sqrtf` (float)
- double `cbrt` (double __x)
- double `hypot` (double __x, double __y)
- double `square` (double __x)
- double `floor` (double __x)
- double `ceil` (double __x)
- double `frexp` (double __x, int *__pexp)
- double `ldexp` (double __x, int __exp)
- double `exp` (double __x)
- double `cosh` (double __x)
- double `sinh` (double __x)

- double [tanh](#) (double __x)
- double [acos](#) (double __x)
- double [asin](#) (double __x)
- double [atan](#) (double __x)
- double [atan2](#) (double __y, double __x)
- double [log](#) (double __x)
- double [log10](#) (double __x)
- double [pow](#) (double __x, double __y)
- int [isnan](#) (double __x)
- int [isinf](#) (double __x)
- static int [isfinite](#) (double __x)
- static double [copysign](#) (double __x, double __y)
- int [signbit](#) (double __x)
- double [fdim](#) (double __x, double __y)
- double [fma](#) (double __x, double __y, double __z)
- double [fmax](#) (double __x, double __y)
- double [fmin](#) (double __x, double __y)
- double [trunc](#) (double __x)
- double [round](#) (double __x)
- long [lround](#) (double __x)
- long [lrint](#) (double __x)

23.6.1 Detailed Description

```
#include <math.h>
```

This header file declares basic mathematics constants and functions.

Notes:

- In order to access the functions declared herein, it is usually also required to additionally link against the library `libm.a`. See also the related [FAQ entry](#).
- Math functions do not raise exceptions and do not change the `errno` variable. Therefore the majority of them are declared with `const` attribute, for better optimization by GCC.

23.6.2 Define Documentation

23.6.2.1 `#define acosf acos`

The alias for [acos\(\)](#).

23.6.2.2 `#define asinf asin`

The alias for [asin\(\)](#).

23.6.2.3 #define atan2f atan2

The alias for [atan2\(\)](#).

23.6.2.4 #define atanf atan

The alias for [atan\(\)](#).

23.6.2.5 #define cbrtf cbrt

The alias for [cbrt\(\)](#).

23.6.2.6 #define ceilf ceil

The alias for [ceil\(\)](#).

23.6.2.7 #define copysignf copysign

The alias for [copysign\(\)](#).

23.6.2.8 #define cosf cos

The alias for [cos\(\)](#).

23.6.2.9 #define coshf cosh

The alias for [cosh\(\)](#).

23.6.2.10 #define expf exp

The alias for [exp\(\)](#).

23.6.2.11 #define fabsf fabs

The alias for [fabs\(\)](#).

23.6.2.12 #define fdimf fdim

The alias for [fdim\(\)](#).

23.6.2.13 #define floorf floor

The alias for [floor\(\)](#).

23.6.2.14 #define fmaf fma

The alias for [fma\(\)](#).

23.6.2.15 #define fmaxf fmax

The alias for [fmax\(\)](#).

23.6.2.16 #define fminf fmin

The alias for [fmin\(\)](#).

23.6.2.17 #define fmodf fmod

The alias for [fmod\(\)](#).

23.6.2.18 #define frexpf frexp

The alias for [frexp\(\)](#).

23.6.2.19 #define hypotf hypot

The alias for [hypot\(\)](#).

23.6.2.20 #define INFINITY __builtin_inf()

INFINITY constant.

23.6.2.21 #define isfinitef isfinite

The alias for [isfinite\(\)](#).

23.6.2.22 #define isinff isinf

The alias for [isinf\(\)](#).

23.6.2.23 #define isnanf isnan

The alias for [isnan\(\)](#).

23.6.2.24 #define ldexpf ldexp

The alias for [ldexp\(\)](#).

23.6.2.25 #define log10f log10

The alias for [log10\(\)](#).

23.6.2.26 #define logf log

The alias for [log\(\)](#).

23.6.2.27 #define lrintf lrint

The alias for [lrint\(\)](#).

23.6.2.28 #define lroundf lround

The alias for [lround\(\)](#).

23.6.2.29 #define M_1_PI 0.31830988618379067154

The constant $1/\pi$.

23.6.2.30 #define M_2_PI 0.63661977236758134308

The constant $2/\pi$.

23.6.2.31 #define M_2_SQRTPI 1.12837916709551257390

The constant $2/\sqrt{\pi}$.

23.6.2.32 #define M_E 2.7182818284590452354

The constant e .

23.6.2.33 #define M_LN10 2.30258509299404568402

The natural logarithm of the 10.

23.6.2.34 #define M_LN2 0.69314718055994530942

The natural logarithm of the 2.

23.6.2.35 #define M_LOG10E 0.43429448190325182765

The logarithm of the e to base 10.

23.6.2.36 **#define M_LOG2E 1.4426950408889634074**

The logarithm of the e to base 2.

23.6.2.37 **#define M_PI 3.14159265358979323846**

The constant π .

23.6.2.38 **#define M_PI_2 1.57079632679489661923**

The constant $\pi/2$.

23.6.2.39 **#define M_PI_4 0.78539816339744830962**

The constant $\pi/4$.

23.6.2.40 **#define M_SQRT1_2 0.70710678118654752440**

The constant $1/\sqrt{2}$.

23.6.2.41 **#define M_SQRT2 1.41421356237309504880**

The square root of 2.

23.6.2.42 **#define NAN __builtin_nan("")**

NAN constant.

23.6.2.43 **#define powf pow**

The alias for [pow\(\)](#).

23.6.2.44 **#define roundf round**

The alias for [round\(\)](#).

23.6.2.45 **#define signbitf signbit**

The alias for [signbit\(\)](#).

23.6.2.46 **#define sinf sin**

The alias for [sin\(\)](#).

23.6.2.47 #define sinhf sinh

The alias for [sinh\(\)](#).

23.6.2.48 #define squaref square

The alias for [square\(\)](#).

23.6.2.49 #define tanf tan

The alias for [tan\(\)](#).

23.6.2.50 #define tanhf tanh

The alias for [tanh\(\)](#).

23.6.2.51 #define truncf trunc

The alias for [trunc\(\)](#).

23.6.3 Function Documentation**23.6.3.1 double acos (double __x)**

The [acos\(\)](#) function computes the principal value of the arc cosine of `__x`. The returned value is in the range $[0, \pi]$ radians. A domain error occurs for arguments not in the range $[-1, +1]$.

23.6.3.2 double asin (double __x)

The [asin\(\)](#) function computes the principal value of the arc sine of `__x`. The returned value is in the range $[-\pi/2, \pi/2]$ radians. A domain error occurs for arguments not in the range $[-1, +1]$.

23.6.3.3 double atan (double __x)

The [atan\(\)](#) function computes the principal value of the arc tangent of `__x`. The returned value is in the range $[-\pi/2, \pi/2]$ radians.

23.6.3.4 double atan2 (double __y, double __x)

The [atan2\(\)](#) function computes the principal value of the arc tangent of `__y` / `__x`, using the signs of both arguments to determine the quadrant of the return value. The returned value is in the range $[-\pi, +\pi]$ radians.

23.6.3.5 double cbrt (double __x)

The `cbrt()` function returns the cube root of `__x`.

23.6.3.6 double ceil (double __x)

The `ceil()` function returns the smallest integral value greater than or equal to `__x`, expressed as a floating-point number.

23.6.3.7 static double copysign (double __x, double __y) [static]

The `copysign()` function returns `__x` but with the sign of `__y`. They work even if `__x` or `__y` are NaN or zero.

23.6.3.8 double cos (double __x)

The `cos()` function returns the cosine of `__x`, measured in radians.

23.6.3.9 double cosh (double __x)

The `cosh()` function returns the hyperbolic cosine of `__x`.

23.6.3.10 double exp (double __x)

The `exp()` function returns the exponential value of `__x`.

23.6.3.11 double fabs (double __x)

The `fabs()` function computes the absolute value of a floating-point number `__x`.

23.6.3.12 double fdim (double __x, double __y)

The `fdim()` function returns $\max(\text{__x} - \text{__y}, 0)$. If `__x` or `__y` or both are NaN, NaN is returned.

23.6.3.13 double floor (double __x)

The `floor()` function returns the largest integral value less than or equal to `__x`, expressed as a floating-point number.

23.6.3.14 double fma (double __x, double __y, double __z)

The `fma()` function performs floating-point multiply-add. This is the operation $(\text{__x} * \text{__y}) + \text{__z}$, but the intermediate result is not rounded to the destination type. This can sometimes improve the precision of a calculation.

23.6.3.15 double fmax (double __x, double __y)

The `fmax()` function returns the greater of the two values `__x` and `__y`. If an argument is NaN, the other argument is returned. If both arguments are NaN, NaN is returned.

23.6.3.16 double fmin (double __x, double __y)

The `fmin()` function returns the lesser of the two values `__x` and `__y`. If an argument is NaN, the other argument is returned. If both arguments are NaN, NaN is returned.

23.6.3.17 double fmod (double __x, double __y)

The function `fmod()` returns the floating-point remainder of $\text{__x} / \text{__y}$.

23.6.3.18 double frexp (double __x, int * __pexp)

The `frexp()` function breaks a floating-point number into a normalized fraction and an integral power of 2. It stores the integer in the `int` object pointed to by `__pexp`.

If `__x` is a normal float point number, the `frexp()` function returns the value v , such that v has a magnitude in the interval $[1/2, 1)$ or zero, and `__x` equals v times 2 raised to the power `__pexp`. If `__x` is zero, both parts of the result are zero. If `__x` is not a finite number, the `frexp()` returns `__x` as is and stores 0 by `__pexp`.

Note

This implementation permits a zero pointer as a directive to skip a storing the exponent.

23.6.3.19 double hypot (double __x, double __y)

The `hypot()` function returns $\text{sqrt}(\text{__x} * \text{__x} + \text{__y} * \text{__y})$. This is the length of the hypotenuse of a right triangle with sides of length `__x` and `__y`, or the distance of the point $(\text{__x}, \text{__y})$ from the origin. Using this function instead of the direct formula is wise, since the error is much smaller. No underflow with small `__x` and `__y`. No overflow if result is in range.

23.6.3.20 static int isfinite (double __x) [static]

The `isfinite()` function returns a nonzero value if `__x` is finite: not plus or minus infinity, and not NaN.

23.6.3.21 int isinf (double __x)

The function `isinf()` returns 1 if the argument `__x` is positive infinity, -1 if `__x` is negative infinity, and 0 otherwise.

Note

The GCC 4.3 can replace this function with inline code that returns the 1 value for both infinities (gcc bug #35509).

23.6.3.22 int isnan (double __x)

The function `isnan()` returns 1 if the argument `__x` represents a "not-a-number" (NaN) object, otherwise 0.

23.6.3.23 double ldexp (double __x, int __exp)

The `ldexp()` function multiplies a floating-point number by an integral power of 2. It returns the value of `__x` times 2 raised to the power `__exp`.

23.6.3.24 double log (double __x)

The `log()` function returns the natural logarithm of argument `__x`.

23.6.3.25 double log10 (double __x)

The `log10()` function returns the logarithm of argument `__x` to base 10.

23.6.3.26 long lrint (double __x)

The `lrint()` function rounds `__x` to the nearest integer, rounding the halfway cases to the even integer direction. (That is both 1.5 and 2.5 values are rounded to 2). This function is similar to `rint()` function, but it differs in type of return value and in that an overflow is possible.

Returns

The rounded long integer value. If `__x` is not a finite number or an overflow was, this realization returns the `LONG_MIN` value (0x80000000).

23.6.3.27 long lround (double __x)

The `lround()` function rounds `__x` to the nearest integer, but rounds halfway cases away from zero (instead of to the nearest even integer). This function is similar to `round()` function, but it differs in type of return value and in that an overflow is possible.

Returns

The rounded long integer value. If `__x` is not a finite number or an overflow was, this realization returns the `LONG_MIN` value (0x80000000).

23.6.3.28 double modf (double __x, double * __iptr)

The `modf()` function breaks the argument `__x` into integral and fractional parts, each of which has the same sign as the argument. It stores the integral part as a double in the object pointed to by `__iptr`.

The `modf()` function returns the signed fractional part of `__x`.

Note

This implementation skips writing by zero pointer. However, the GCC 4.3 can replace this function with inline code that does not permit to use NULL address for the avoiding of storing.

23.6.3.29 float modff (float __x, float * __iptr)

An alias for `modf()`.

23.6.3.30 double pow (double __x, double __y)

The function `pow()` returns the value of `__x` to the exponent `__y`.

23.6.3.31 double round (double __x)

The `round()` function rounds `__x` to the nearest integer, but rounds halfway cases away from zero (instead of to the nearest even integer). Overflow is impossible.

Returns

The rounded value. If `__x` is an integral or infinite, `__x` itself is returned. If `__x` is NaN, then NaN is returned.

23.6.3.32 int signbit (double __x)

The `signbit()` function returns a nonzero value if the value of `__x` has its sign bit set. This is not the same as '`__x < 0.0`', because IEEE 754 floating point allows zero to be signed. The comparison '`-0.0 < 0.0`' is false, but '`signbit (-0.0)`' will return a nonzero value.

23.6.3.33 double sin (double __x)

The `sin()` function returns the sine of `__x`, measured in radians.

23.6.3.34 double sinh (double __x)

The `sinh()` function returns the hyperbolic sine of `__x`.

23.6.3.35 double sqrt (double __x)

The `sqrt()` function returns the non-negative square root of `__x`.

23.6.3.36 float sqrtf (float)

An alias for `sqrt()`.

23.6.3.37 double square (double __x)

The function `square()` returns `__x * __x`.

Note

This function does not belong to the C standard definition.

23.6.3.38 double tan (double __x)

The `tan()` function returns the tangent of `__x`, measured in radians.

23.6.3.39 double tanh (double __x)

The `tanh()` function returns the hyperbolic tangent of `__x`.

23.6.3.40 double trunc (double __x)

The `trunc()` function rounds `__x` to the nearest integer not larger in absolute value.

23.7 <setjmp.h>: Non-local goto

Functions

- int [setjmp](#) (jmp_buf __jmpb)
- void [longjmp](#) (jmp_buf __jmpb, int __ret) __ATTR_NORETURN__

23.7.1 Detailed Description

While the C language has the dreaded `goto` statement, it can only be used to jump to a label in the same (local) function. In order to jump directly to another (non-local) function, the C library provides the [setjmp\(\)](#) and [longjmp\(\)](#) functions. [setjmp\(\)](#) and [longjmp\(\)](#) are useful for dealing with errors and interrupts encountered in a low-level subroutine of a program.

Note

[setjmp\(\)](#) and [longjmp\(\)](#) make programs hard to understand and maintain. If possible, an alternative should be used.

[longjmp\(\)](#) can destroy changes made to global register variables (see [How to permanently bind a variable to a register?](#)).

For a very detailed discussion of [setjmp\(\)](#)/[longjmp\(\)](#), see Chapter 7 of *Advanced Programming in the UNIX Environment*, by W. Richard Stevens.

Example:

```
#include <setjmp.h>

jmp_buf env;

int main (void)
{
    if (setjmp (env))
    {
        ... handle error ...
    }

    while (1)
    {
        ... main processing loop which calls foo() some where ...
    }
}

...

void foo (void)
{
    ... blah, blah, blah ...

    if (err)
    {
        longjmp (env, 1);
    }
}
```

23.7.2 Function Documentation

23.7.2.1 void longjmp (jmp_buf __jmpb, int __ret)

Non-local jump to a saved stack context.

```
#include <setjmp.h>
```

[longjmp\(\)](#) restores the environment saved by the last call of [setjmp\(\)](#) with the corresponding `__jmpb` argument. After [longjmp\(\)](#) is completed, program execution continues as if the corresponding call of [setjmp\(\)](#) had just returned the value `__ret`.

Note

[longjmp\(\)](#) cannot cause 0 to be returned. If [longjmp\(\)](#) is invoked with a second argument of 0, 1 will be returned instead.

Parameters

`__jmpb` Information saved by a previous call to [setjmp\(\)](#).

`__ret` Value to return to the caller of [setjmp\(\)](#).

Returns

This function never returns.

23.7.2.2 int setjmp (jmp_buf __jmpb)

Save stack context for non-local goto.

```
#include <setjmp.h>
```

[setjmp\(\)](#) saves the stack context/environment in `__jmpb` for later use by [longjmp\(\)](#). The stack context will be invalidated if the function which called [setjmp\(\)](#) returns.

Parameters

`__jmpb` Variable of type `jmp_buf` which holds the stack information such that the environment can be restored.

Returns

[setjmp\(\)](#) returns 0 if returning directly, and non-zero when returning from [longjmp\(\)](#) using the saved context.

23.8 <stdint.h>: Standard Integer Types

Exact-width integer types

Integer types having exactly the specified width

- typedef signed char [int8_t](#)
- typedef unsigned char [uint8_t](#)
- typedef signed int [int16_t](#)
- typedef unsigned int [uint16_t](#)
- typedef signed long int [int32_t](#)
- typedef unsigned long int [uint32_t](#)
- typedef signed long long int [int64_t](#)
- typedef unsigned long long int [uint64_t](#)

Integer types capable of holding object pointers

These allow you to declare variables of the same size as a pointer.

- typedef [int16_t](#) [intptr_t](#)
- typedef [uint16_t](#) [uintptr_t](#)

Minimum-width integer types

Integer types having at least the specified width

- typedef [int8_t](#) [int_least8_t](#)
- typedef [uint8_t](#) [uint_least8_t](#)
- typedef [int16_t](#) [int_least16_t](#)
- typedef [uint16_t](#) [uint_least16_t](#)
- typedef [int32_t](#) [int_least32_t](#)
- typedef [uint32_t](#) [uint_least32_t](#)
- typedef [int64_t](#) [int_least64_t](#)
- typedef [uint64_t](#) [uint_least64_t](#)

Fastest minimum-width integer types

Integer types being usually fastest having at least the specified width

- typedef [int8_t](#) [int_fast8_t](#)
- typedef [uint8_t](#) [uint_fast8_t](#)
- typedef [int16_t](#) [int_fast16_t](#)
- typedef [uint16_t](#) [uint_fast16_t](#)
- typedef [int32_t](#) [int_fast32_t](#)
- typedef [uint32_t](#) [uint_fast32_t](#)
- typedef [int64_t](#) [int_fast64_t](#)
- typedef [uint64_t](#) [uint_fast64_t](#)

Greatest-width integer types

Types designating integer data capable of representing any value of any integer type in the corresponding signed or unsigned category

- typedef [int64_t](#) [intmax_t](#)
- typedef [uint64_t](#) [uintmax_t](#)

Limits of specified-width integer types

C++ implementations should define these macros only when `__STDC_LIMIT_MACROS` is defined before `<stdint.h>` is included

- #define [INT8_MAX](#) 0x7f
- #define [INT8_MIN](#) (-INT8_MAX - 1)
- #define [UINT8_MAX](#) (INT8_MAX * 2 + 1)
- #define [INT16_MAX](#) 0x7fff
- #define [INT16_MIN](#) (-INT16_MAX - 1)
- #define [UINT16_MAX](#) (__CONCAT(INT16_MAX, U) * 2U + 1U)
- #define [INT32_MAX](#) 0x7fffffffL
- #define [INT32_MIN](#) (-INT32_MAX - 1L)
- #define [UINT32_MAX](#) (__CONCAT(INT32_MAX, U) * 2UL + 1UL)
- #define [INT64_MAX](#) 0x7fffffffffffffffLL
- #define [INT64_MIN](#) (-INT64_MAX - 1LL)
- #define [UINT64_MAX](#) (__CONCAT(INT64_MAX, U) * 2ULL + 1ULL)

Limits of minimum-width integer types

- #define [INT_LEAST8_MAX](#) INT8_MAX
- #define [INT_LEAST8_MIN](#) INT8_MIN
- #define [UINT_LEAST8_MAX](#) UINT8_MAX
- #define [INT_LEAST16_MAX](#) INT16_MAX
- #define [INT_LEAST16_MIN](#) INT16_MIN
- #define [UINT_LEAST16_MAX](#) UINT16_MAX
- #define [INT_LEAST32_MAX](#) INT32_MAX
- #define [INT_LEAST32_MIN](#) INT32_MIN
- #define [UINT_LEAST32_MAX](#) UINT32_MAX
- #define [INT_LEAST64_MAX](#) INT64_MAX
- #define [INT_LEAST64_MIN](#) INT64_MIN
- #define [UINT_LEAST64_MAX](#) UINT64_MAX

Limits of fastest minimum-width integer types

- #define `INT_FAST8_MAX` `INT8_MAX`
- #define `INT_FAST8_MIN` `INT8_MIN`
- #define `UINT_FAST8_MAX` `UINT8_MAX`
- #define `INT_FAST16_MAX` `INT16_MAX`
- #define `INT_FAST16_MIN` `INT16_MIN`
- #define `UINT_FAST16_MAX` `UINT16_MAX`
- #define `INT_FAST32_MAX` `INT32_MAX`
- #define `INT_FAST32_MIN` `INT32_MIN`
- #define `UINT_FAST32_MAX` `UINT32_MAX`
- #define `INT_FAST64_MAX` `INT64_MAX`
- #define `INT_FAST64_MIN` `INT64_MIN`
- #define `UINT_FAST64_MAX` `UINT64_MAX`

Limits of integer types capable of holding object pointers

- #define `INTPTR_MAX` `INT16_MAX`
- #define `INTPTR_MIN` `INT16_MIN`
- #define `UINTPTR_MAX` `UINT16_MAX`

Limits of greatest-width integer types

- #define `INTMAX_MAX` `INT64_MAX`
- #define `INTMAX_MIN` `INT64_MIN`
- #define `UINTMAX_MAX` `UINT64_MAX`

Limits of other integer types

C++ implementations should define these macros only when `__STDC_LIMIT_MACROS` is defined before `<stdint.h>` is included

- #define `PTRDIFF_MAX` `INT16_MAX`
- #define `PTRDIFF_MIN` `INT16_MIN`
- #define `SIG_ATOMIC_MAX` `INT8_MAX`
- #define `SIG_ATOMIC_MIN` `INT8_MIN`
- #define `SIZE_MAX` `UINT16_MAX`
- #define `WCHAR_MAX` `__WCHAR_MAX__`
- #define `WCHAR_MIN` `__WCHAR_MIN__`
- #define `WINT_MAX` `__WINT_MAX__`
- #define `WINT_MIN` `__WINT_MIN__`

Macros for integer constants

C++ implementations should define these macros only when `__STDC_CONSTANT_MACROS` is defined before `<stdint.h>` is included.

These definitions are valid for integer constants without suffix and for macros defined as integer constant without suffix

- `#define INT8_C(value) ((int8_t) value)`
- `#define UINT8_C(value) ((uint8_t) __CONCAT(value, U))`
- `#define INT16_C(value) value`
- `#define UINT16_C(value) __CONCAT(value, U)`
- `#define INT32_C(value) __CONCAT(value, L)`
- `#define UINT32_C(value) __CONCAT(value, UL)`
- `#define INT64_C(value) __CONCAT(value, LL)`
- `#define UINT64_C(value) __CONCAT(value, ULL)`
- `#define INTMAX_C(value) __CONCAT(value, LL)`
- `#define UINTMAX_C(value) __CONCAT(value, ULL)`

23.8.1 Detailed Description

```
#include <stdint.h>
```

Use `[u]intN_t` if you need exactly N bits.

Since these typedefs are mandated by the C99 standard, they are preferred over rolling your own typedefs.

23.8.2 Define Documentation

23.8.2.1 `#define INT16_C(value) value`

define a constant of type `int16_t`

23.8.2.2 `#define INT16_MAX 0x7fff`

largest positive value an `int16_t` can hold.

23.8.2.3 `#define INT16_MIN (-INT16_MAX - 1)`

smallest negative value an `int16_t` can hold.

23.8.2.4 `#define INT32_C(value) __CONCAT(value, L)`

define a constant of type `int32_t`

23.8.2.5 #define INT32_MAX 0x7fffffffL

largest positive value an int32_t can hold.

23.8.2.6 #define INT32_MIN (-INT32_MAX - 1L)

smallest negative value an int32_t can hold.

23.8.2.7 #define INT64_C(value) __CONCAT(value, LL)

define a constant of type int64_t

23.8.2.8 #define INT64_MAX 0x7fffffffffffffffLL

largest positive value an int64_t can hold.

23.8.2.9 #define INT64_MIN (-INT64_MAX - 1LL)

smallest negative value an int64_t can hold.

23.8.2.10 #define INT8_C(value) ((int8_t) value)

define a constant of type int8_t

23.8.2.11 #define INT8_MAX 0x7f

largest positive value an int8_t can hold.

23.8.2.12 #define INT8_MIN (-INT8_MAX - 1)

smallest negative value an int8_t can hold.

23.8.2.13 #define INT_FAST16_MAX INT16_MAX

largest positive value an int_fast16_t can hold.

23.8.2.14 #define INT_FAST16_MIN INT16_MIN

smallest negative value an int_fast16_t can hold.

23.8.2.15 #define INT_FAST32_MAX INT32_MAX

largest positive value an int_fast32_t can hold.

23.8.2.16 #define INT_FAST32_MIN INT32_MIN

smallest negative value an int_fast32_t can hold.

23.8.2.17 #define INT_FAST64_MAX INT64_MAX

largest positive value an int_fast64_t can hold.

23.8.2.18 #define INT_FAST64_MIN INT64_MIN

smallest negative value an int_fast64_t can hold.

23.8.2.19 #define INT_FAST8_MAX INT8_MAX

largest positive value an int_fast8_t can hold.

23.8.2.20 #define INT_FAST8_MIN INT8_MIN

smallest negative value an int_fast8_t can hold.

23.8.2.21 #define INT_LEAST16_MAX INT16_MAX

largest positive value an int_least16_t can hold.

23.8.2.22 #define INT_LEAST16_MIN INT16_MIN

smallest negative value an int_least16_t can hold.

23.8.2.23 #define INT_LEAST32_MAX INT32_MAX

largest positive value an int_least32_t can hold.

23.8.2.24 #define INT_LEAST32_MIN INT32_MIN

smallest negative value an int_least32_t can hold.

23.8.2.25 #define INT_LEAST64_MAX INT64_MAX

largest positive value an int_least64_t can hold.

23.8.2.26 #define INT_LEAST64_MIN INT64_MIN

smallest negative value an int_least64_t can hold.

23.8.2.27 `#define INT_LEAST8_MAX INT8_MAX`

largest positive value an `int_least8_t` can hold.

23.8.2.28 `#define INT_LEAST8_MIN INT8_MIN`

smallest negative value an `int_least8_t` can hold.

23.8.2.29 `#define INTMAX_C(value) __CONCAT(value, LL)`

define a constant of type `intmax_t`

23.8.2.30 `#define INTMAX_MAX INT64_MAX`

largest positive value an `intmax_t` can hold.

23.8.2.31 `#define INTMAX_MIN INT64_MIN`

smallest negative value an `intmax_t` can hold.

23.8.2.32 `#define INTPTR_MAX INT16_MAX`

largest positive value an `intptr_t` can hold.

23.8.2.33 `#define INTPTR_MIN INT16_MIN`

smallest negative value an `intptr_t` can hold.

23.8.2.34 `#define PTRDIFF_MAX INT16_MAX`

largest positive value a `ptrdiff_t` can hold.

23.8.2.35 `#define PTRDIFF_MIN INT16_MIN`

smallest negative value a `ptrdiff_t` can hold.

23.8.2.36 `#define SIG_ATOMIC_MAX INT8_MAX`

largest positive value a `sig_atomic_t` can hold.

23.8.2.37 `#define SIG_ATOMIC_MIN INT8_MIN`

smallest negative value a `sig_atomic_t` can hold.

23.8.2.38 `#define SIZE_MAX UINT16_MAX`

largest value a `size_t` can hold.

23.8.2.39 `#define UINT16_C(value) __CONCAT(value, U)`

define a constant of type `uint16_t`

23.8.2.40 `#define UINT16_MAX (__CONCAT(INT16_MAX, U) * 2U + 1U)`

largest value an `uint16_t` can hold.

23.8.2.41 `#define UINT32_C(value) __CONCAT(value, UL)`

define a constant of type `uint32_t`

23.8.2.42 `#define UINT32_MAX (__CONCAT(INT32_MAX, U) * 2UL + 1UL)`

largest value an `uint32_t` can hold.

23.8.2.43 `#define UINT64_C(value) __CONCAT(value, ULL)`

define a constant of type `uint64_t`

23.8.2.44 `#define UINT64_MAX (__CONCAT(INT64_MAX, U) * 2ULL + 1ULL)`

largest value an `uint64_t` can hold.

23.8.2.45 `#define UINT8_C(value) ((uint8_t) __CONCAT(value, U))`

define a constant of type `uint8_t`

23.8.2.46 `#define UINT8_MAX (INT8_MAX * 2 + 1)`

largest value an `uint8_t` can hold.

23.8.2.47 `#define UINT_FAST16_MAX UINT16_MAX`

largest value an `uint_fast16_t` can hold.

23.8.2.48 #define UINT_FAST32_MAX UINT32_MAX

largest value an uint_fast32_t can hold.

23.8.2.49 #define UINT_FAST64_MAX UINT64_MAX

largest value an uint_fast64_t can hold.

23.8.2.50 #define UINT_FAST8_MAX UINT8_MAX

largest value an uint_fast8_t can hold.

23.8.2.51 #define UINT_LEAST16_MAX UINT16_MAX

largest value an uint_least16_t can hold.

23.8.2.52 #define UINT_LEAST32_MAX UINT32_MAX

largest value an uint_least32_t can hold.

23.8.2.53 #define UINT_LEAST64_MAX UINT64_MAX

largest value an uint_least64_t can hold.

23.8.2.54 #define UINT_LEAST8_MAX UINT8_MAX

largest value an uint_least8_t can hold.

23.8.2.55 #define UINTMAX_C(value) __CONCAT(value, ULL)

define a constant of type uintmax_t

23.8.2.56 #define UINTMAX_MAX UINT64_MAX

largest value an uintmax_t can hold.

23.8.2.57 #define UINTPTR_MAX UINT16_MAX

largest value an uintptr_t can hold.

23.8.3 Typedef Documentation

23.8.3.1 typedef signed int int16_t

16-bit signed type.

23.8.3.2 typedef signed long int int32_t

32-bit signed type.

23.8.3.3 typedef signed long long int int64_t

64-bit signed type.

Note

This type is not available when the compiler option `-mint8` is in effect.

23.8.3.4 typedef signed char int8_t

8-bit signed type.

23.8.3.5 typedef int16_t int_fast16_t

fastest signed int with at least 16 bits.

23.8.3.6 typedef int32_t int_fast32_t

fastest signed int with at least 32 bits.

23.8.3.7 typedef int64_t int_fast64_t

fastest signed int with at least 64 bits.

Note

This type is not available when the compiler option `-mint8` is in effect.

23.8.3.8 typedef int8_t int_fast8_t

fastest signed int with at least 8 bits.

23.8.3.9 typedef int16_t int_least16_t

signed int with at least 16 bits.

23.8.3.10 typedef int32_t int_least32_t

signed int with at least 32 bits.

23.8.3.11 typedef int64_t int_least64_t

signed int with at least 64 bits.

Note

This type is not available when the compiler option -mint8 is in effect.

23.8.3.12 typedef int8_t int_least8_t

signed int with at least 8 bits.

23.8.3.13 typedef int64_t intmax_t

largest signed int available.

23.8.3.14 typedef int16_t intptr_t

Signed pointer compatible type.

23.8.3.15 typedef unsigned int uint16_t

16-bit unsigned type.

23.8.3.16 typedef unsigned long int uint32_t

32-bit unsigned type.

23.8.3.17 typedef unsigned long long int uint64_t

64-bit unsigned type.

Note

This type is not available when the compiler option -mint8 is in effect.

23.8.3.18 typedef unsigned char uint8_t

8-bit unsigned type.

23.8.3.19 typedef uint16_t uint_fast16_t

fastest unsigned int with at least 16 bits.

23.8.3.20 typedef uint32_t uint_fast32_t

fastest unsigned int with at least 32 bits.

23.8.3.21 typedef uint64_t uint_fast64_t

fastest unsigned int with at least 64 bits.

Note

This type is not available when the compiler option `-mint8` is in effect.

23.8.3.22 typedef uint8_t uint_fast8_t

fastest unsigned int with at least 8 bits.

23.8.3.23 typedef uint16_t uint_least16_t

unsigned int with at least 16 bits.

23.8.3.24 typedef uint32_t uint_least32_t

unsigned int with at least 32 bits.

23.8.3.25 typedef uint64_t uint_least64_t

unsigned int with at least 64 bits.

Note

This type is not available when the compiler option `-mint8` is in effect.

23.8.3.26 typedef uint8_t uint_least8_t

unsigned int with at least 8 bits.

23.8.3.27 typedef uint64_t uintmax_t

largest unsigned int available.

23.8.3.28 typedef uint16_t uintptr_t

Unsigned pointer compatible type.

23.9 <stdio.h>: Standard IO facilities

Defines

- #define [stdin](#) (__iob[0])
- #define [stdout](#) (__iob[1])
- #define [stderr](#) (__iob[2])
- #define [EOF](#) (-1)
- #define [fdev_set_udata](#)(stream, u) do { (stream)->udata = u; } while(0)
- #define [fdev_get_udata](#)(stream) ((stream)->udata)
- #define [fdev_setup_stream](#)(stream, put, get, rflag)
- #define [_FDEV_SETUP_READ](#) __SRD
- #define [_FDEV_SETUP_WRITE](#) __SWR
- #define [_FDEV_SETUP_RW](#) (__SRD|__SWR)
- #define [_FDEV_ERR](#) (-1)
- #define [_FDEV_EOF](#) (-2)
- #define [FDEV_SETUP_STREAM](#)(put, get, rflag)
- #define [fdev_close](#)()
- #define [putc](#)(__c, __stream) fputc(__c, __stream)
- #define [putchar](#)(__c) fputc(__c, stdout)
- #define [getc](#)(__stream) fgetc(__stream)
- #define [getchar](#)() fgetc(stdin)

Typedefs

- typedef struct __file [FILE](#)

Functions

- int [fclose](#) ([FILE](#) *__stream)
- int [vfprintf](#) ([FILE](#) *__stream, const char *__fmt, va_list __ap)
- int [vfprintf_P](#) ([FILE](#) *__stream, const char *__fmt, va_list __ap)
- int [fputc](#) (int __c, [FILE](#) *__stream)
- int [printf](#) (const char *__fmt,...)
- int [printf_P](#) (const char *__fmt,...)
- int [vprintf](#) (const char *__fmt, va_list __ap)
- int [sprintf](#) (char *__s, const char *__fmt,...)
- int [sprintf_P](#) (char *__s, const char *__fmt,...)
- int [snprintf](#) (char *__s, size_t __n, const char *__fmt,...)
- int [snprintf_P](#) (char *__s, size_t __n, const char *__fmt,...)
- int [vsprintf](#) (char *__s, const char *__fmt, va_list __ap)
- int [vsprintf_P](#) (char *__s, const char *__fmt, va_list __ap)

- int `vsnprintf` (char *__s, size_t __n, const char *__fmt, va_list ap)
- int `vsnprintf_P` (char *__s, size_t __n, const char *__fmt, va_list ap)
- int `fprintf` (FILE *__stream, const char *__fmt,...)
- int `fprintf_P` (FILE *__stream, const char *__fmt,...)
- int `fputs` (const char *__str, FILE *__stream)
- int `fputs_P` (const char *__str, FILE *__stream)
- int `puts` (const char *__str)
- int `puts_P` (const char *__str)
- size_t `fwrite` (const void *__ptr, size_t __size, size_t __nmemb, FILE *__stream)
- int `fgetc` (FILE *__stream)
- int `ungetc` (int __c, FILE *__stream)
- char * `fgets` (char *__str, int __size, FILE *__stream)
- char * `gets` (char *__str)
- size_t `fread` (void *__ptr, size_t __size, size_t __nmemb, FILE *__stream)
- void `clearerr` (FILE *__stream)
- int `feof` (FILE *__stream)
- int `ferror` (FILE *__stream)
- int `vfscanf` (FILE *__stream, const char *__fmt, va_list __ap)
- int `vfscanf_P` (FILE *__stream, const char *__fmt, va_list __ap)
- int `fscanf` (FILE *__stream, const char *__fmt,...)
- int `fscanf_P` (FILE *__stream, const char *__fmt,...)
- int `scanf` (const char *__fmt,...)
- int `scanf_P` (const char *__fmt,...)
- int `vscanf` (const char *__fmt, va_list __ap)
- int `sscanf` (const char *__buf, const char *__fmt,...)
- int `sscanf_P` (const char *__buf, const char *__fmt,...)
- int `fflush` (FILE *stream)
- FILE * `fdevopen` (int(*put)(char, FILE *), int(*get)(FILE *))

23.9.1 Detailed Description

```
#include <stdio.h>
```

Introduction to the Standard IO facilities This file declares the standard IO facilities that are implemented in `avr-libc`. Due to the nature of the underlying hardware, only a limited subset of standard IO is implemented. There is no actual file implementation available, so only device IO can be performed. Since there's no operating system, the application needs to provide enough details about their devices in order to make them usable by the standard IO facilities.

Due to space constraints, some functionality has not been implemented at all (like some of the `printf` conversions that have been left out). Nevertheless, potential users of this implementation should be warned: the `printf` and `scanf` families of functions, although usually associated with presumably simple things like the famous "Hello, world!" program, are actually fairly complex which causes their inclusion to eat up a fair amount of code space. Also, they are not fast due to the nature of interpreting

the format string at run-time. Whenever possible, resorting to the (sometimes non-standard) predetermined conversion facilities that are offered by avr-libc will usually cost much less in terms of speed and code size.

Tunable options for code size vs. feature set In order to allow programmers a code size vs. functionality tradeoff, the function `vfprintf()` which is the heart of the `printf` family can be selected in different flavours using linker options. See the documentation of `vfprintf()` for a detailed description. The same applies to `vfscanf()` and the `scanf` family of functions.

Outline of the chosen API The standard streams `stdin`, `stdout`, and `stderr` are provided, but contrary to the C standard, since avr-libc has no knowledge about applicable devices, these streams are not already pre-initialized at application startup. Also, since there is no notion of "file" whatsoever to avr-libc, there is no function `fopen()` that could be used to associate a stream to some device. (See [note 1](#).) Instead, the function `fdevopen()` is provided to associate a stream to a device, where the device needs to provide a function to send a character, to receive a character, or both. There is no differentiation between "text" and "binary" streams inside avr-libc. Character `\n` is sent literally down to the device's `put()` function. If the device requires a carriage return (`\r`) character to be sent before the linefeed, its `put()` routine must implement this (see [note 2](#)).

As an alternative method to `fdevopen()`, the macro `fdev_setup_stream()` might be used to setup a user-supplied FILE structure.

It should be noted that the automatic conversion of a newline character into a carriage return - newline sequence breaks binary transfers. If binary transfers are desired, no automatic conversion should be performed, but instead any string that aims to issue a CR-LF sequence must use `"\r\n"` explicitly.

For convenience, the first call to `fdevopen()` that opens a stream for reading will cause the resulting stream to be aliased to `stdin`. Likewise, the first call to `fdevopen()` that opens a stream for writing will cause the resulting stream to be aliased to both, `stdout`, and `stderr`. Thus, if the open was done with both, read and write intent, all three standard streams will be identical. Note that these aliases are indistinguishable from each other, thus calling `fclose()` on such a stream will also effectively close all of its aliases ([note 3](#)).

It is possible to tie additional user data to a stream, using `fdev_set_udata()`. The backend `put` and `get` functions can then extract this user data using `fdev_get_udata()`, and act appropriately. For example, a single `put` function could be used to talk to two different UARTs that way, or the `put` and `get` functions could keep internal state between calls there.

Format strings in flash ROM All the `printf` and `scanf` family functions come in two flavours: the standard name, where the format string is expected to be in SRAM, as well as a version with the suffix `"_P"` where the format string is expected to reside in the flash ROM. The macro `PSTR` (explained in [<avr/pgmspace.h>: Program Space Utilities](#)) becomes very handy for declaring these format strings.

Running stdio without `malloc()` By default, `fdevopen()` requires `malloc()`. As this is often not desired in the limited environment of a microcontroller, an alternative option is provided to run completely without `malloc()`.

The macro `fdev_setup_stream()` is provided to prepare a user-supplied FILE buffer for operation with stdio.

Example

```
#include <stdio.h>

static int uart_putchar(char c, FILE *stream);

static FILE mystdout = FDEV_SETUP_STREAM(uart_putchar, NULL,
                                         _FDEV_SETUP_WRITE);

static int
uart_putchar(char c, FILE *stream)
{
    if (c == '\n')
        uart_putchar('\r', stream);
    loop_until_bit_is_set(UCSRA, UDRE);
    UDR = c;
    return 0;
}

int
main(void)
{
    init_uart();
    stdout = &mystdout;
    printf("Hello, world!\n");

    return 0;
}
```

This example uses the initializer form `FDEV_SETUP_STREAM()` rather than the function-like `fdev_setup_stream()`, so all data initialization happens during C start-up.

If streams initialized that way are no longer needed, they can be destroyed by first calling the macro `fdev_close()`, and then destroying the object itself. No call to `fclose()` should be issued for these streams. While calling `fclose()` itself is harmless, it will cause an undefined reference to `free()` and thus cause the linker to link the malloc module into the application.

Notes

Note 1:

It might have been possible to implement a device abstraction that is compatible with `fopen()` but since this would have required to parse a string, and to take all the information needed either out of this string, or out of an additional table that would need to be provided by the application, this approach was not taken.

Note 2:

This basically follows the Unix approach: if a device such as a terminal needs special handling, it is in the domain of the terminal device driver to provide this functionality. Thus, a simple function suitable as `put()` for `fdevopen()` that talks to a UART interface might look like this:

```
int
uart_putchar(char c, FILE *stream)
{
    if (c == '\n')
        uart_putchar('\r');
    loop_until_bit_is_set(UCSRA, UDRE);
    UDR = c;
    return 0;
}
```

Note 3:

This implementation has been chosen because the cost of maintaining an alias is considerably smaller than the cost of maintaining full copies of each stream. Yet, providing an implementation that offers the complete set of standard streams was deemed to be useful. Not only that writing `printf()` instead of `fprintf(mystream, ...)` saves typing work, but since `avr-gcc` needs to resort to pass all arguments of variadic functions on the stack (as opposed to passing them in registers for functions that take a fixed number of parameters), the ability to pass one parameter less by implying `stdin` or `stdout` will also save some execution time.

23.9.2 Define Documentation**23.9.2.1 #define _FDEV_EOF (-2)**

Return code for an end-of-file condition during device read.

To be used in the get function of `fdevopen()`.

23.9.2.2 #define _FDEV_ERR (-1)

Return code for an error condition during device read.

To be used in the get function of `fdevopen()`.

23.9.2.3 #define _FDEV_SETUP_READ __SRD

`fdev_setup_stream()` with read intent

23.9.2.4 #define _FDEV_SETUP_RW (__SRD|__SWR)

`fdev_setup_stream()` with read/write intent

23.9.2.5 #define _FDEV_SETUP_WRITE __SWR

[fdev_setup_stream\(\)](#) with write intent

23.9.2.6 #define EOF (-1)

EOF declares the value that is returned by various standard IO functions in case of an error. Since the AVR platform (currently) doesn't contain an abstraction for actual files, its origin as "end of file" is somewhat meaningless here.

23.9.2.7 #define fdev_close()

This macro frees up any library resources that might be associated with `stream`. It should be called if `stream` is no longer needed, right before the application is going to destroy the `stream` object itself.

(Currently, this macro evaluates to nothing, but this might change in future versions of the library.)

23.9.2.8 #define fdev_get_udata(stream) ((stream)->udata)

This macro retrieves a pointer to user defined data from a FILE stream object.

23.9.2.9 #define fdev_set_udata(stream, u) do { (stream)->udata = u; } while(0)

This macro inserts a pointer to user defined data into a FILE stream object. The user data can be useful for tracking state in the put and get functions supplied to the [fdevopen\(\)](#) function.

23.9.2.10 #define FDEV_SETUP_STREAM(put, get, rwflag)

Initializer for a user-supplied stdio stream.

This macro acts similar to [fdev_setup_stream\(\)](#), but it is to be used as the initializer of a variable of type FILE.

The remaining arguments are to be used as explained in [fdev_setup_stream\(\)](#).

23.9.2.11 #define fdev_setup_stream(stream, put, get, rwflag)

Setup a user-supplied buffer as an stdio stream.

This macro takes a user-supplied buffer `stream`, and sets it up as a stream that is valid for stdio operations, similar to one that has been obtained dynamically from `fdevopen()`. The buffer to setup must be of type `FILE`.

The arguments `put` and `get` are identical to those that need to be passed to `fdevopen()`.

The `rwflag` argument can take one of the values `_FDEV_SETUP_READ`, `_FDEV_SETUP_WRITE`, or `_FDEV_SETUP_RW`, for read, write, or read/write intent, respectively.

Note

No assignments to the standard streams will be performed by `fdev_setup_stream()`. If standard streams are to be used, these need to be assigned by the user. See also under [Running stdio without malloc\(\)](#).

23.9.2.12 #define getc(__stream) fgetc(__stream)

The macro `getc` used to be a "fast" macro implementation with a functionality identical to `fgetc()`. For space constraints, in `avr-libc`, it is just an alias for `fgetc`.

23.9.2.13 #define getchar(void) fgetc(stdin)

The macro `getchar` reads a character from `stdin`. Return values and error handling is identical to `fgetc()`.

23.9.2.14 #define putc(__c, __stream) fputc(__c, __stream)

The macro `putc` used to be a "fast" macro implementation with a functionality identical to `fputc()`. For space constraints, in `avr-libc`, it is just an alias for `fputc`.

23.9.2.15 #define putchar(__c) fputc(__c, stdout)

The macro `putchar` sends character `c` to `stdout`.

23.9.2.16 #define stderr (__iob[2])

Stream destined for error output. Unless specifically assigned, identical to `stdout`. If `stderr` should point to another stream, the result of another `fdevopen()` must be explicitly assigned to it without closing the previous `stderr` (since this would also close `stdout`).

23.9.2.17 #define stdin (__iob[0])

Stream that will be used as an input stream by the simplified functions that don't take a `stream` argument.

The first stream opened with read intent using `fdevopen()` will be assigned to `stdin`.

23.9.2.18 #define stdout (__iob[1])

Stream that will be used as an output stream by the simplified functions that don't take a `stream` argument.

The first stream opened with write intent using `fdevopen()` will be assigned to both, `stdin`, and `stderr`.

23.9.3 Typedef Documentation

23.9.3.1 typedef struct __file FILE

`FILE` is the opaque structure that is passed around between the various standard IO functions.

23.9.4 Function Documentation

23.9.4.1 void clearerr (FILE * __stream)

Clear the error and end-of-file flags of `stream`.

23.9.4.2 int fclose (FILE * __stream)

This function closes `stream`, and disallows and further IO to and from it.

When using `fdevopen()` to setup the stream, a call to `fclose()` is needed in order to free the internal resources allocated.

If the stream has been set up using `fdev_setup_stream()` or `FDEV_SETUP_STREAM()`, use `fdev_close()` instead.

It currently always returns 0 (for success).

23.9.4.3 FILE* fdevopen (int(*) (char, FILE *) put, int(*) (FILE *) get)

This function is a replacement for `fopen()`.

It opens a stream for a device where the actual device implementation needs to be provided by the application. If successful, a pointer to the structure for the opened stream is returned. Reasons for a possible failure currently include that neither the `put` nor the `get` argument have been provided, thus attempting to open a stream with

no IO intent at all, or that insufficient dynamic memory is available to establish a new stream.

If the `put` function pointer is provided, the stream is opened with write intent. The function passed as `put` shall take two arguments, the first a character to write to the device, and the second a pointer to `FILE`, and shall return 0 if the output was successful, and a nonzero value if the character could not be sent to the device.

If the `get` function pointer is provided, the stream is opened with read intent. The function passed as `get` shall take a pointer to `FILE` as its single argument, and return one character from the device, passed as an `int` type. If an error occurs when trying to read from the device, it shall return `_FDEV_ERR`. If an end-of-file condition was reached while reading from the device, `_FDEV_EOF` shall be returned.

If both functions are provided, the stream is opened with read and write intent.

The first stream opened with read intent is assigned to `stdin`, and the first one opened with write intent is assigned to both, `stdout` and `stderr`.

`fdevopen()` uses `calloc()` (and thus `malloc()`) in order to allocate the storage for the new stream.

Note

If the macro `__STDIO_FDEVOPEN_COMPAT_12` is declared before including `<stdio.h>`, a function prototype for `fdevopen()` will be chosen that is backwards compatible with `avr-libc` version 1.2 and before. This is solely intended for providing a simple migration path without the need to immediately change all source code. Do not use for new code.

23.9.4.4 `int feof (FILE * __stream)`

Test the end-of-file flag of `stream`. This flag can only be cleared by a call to `clearerr()`.

23.9.4.5 `int ferror (FILE * __stream)`

Test the error flag of `stream`. This flag can only be cleared by a call to `clearerr()`.

23.9.4.6 `int fflush (FILE * stream)`

Flush `stream`.

This is a null operation provided for source-code compatibility only, as the standard IO implementation currently does not perform any buffering.

23.9.4.7 `int fgetc (FILE * __stream)`

The function `fgetc` reads a character from `stream`. It returns the character, or `EOF` in case end-of-file was encountered or an error occurred. The routines `feof()` or `ferror()` must be used to distinguish between both situations.

23.9.4.8 char* fgets (char * __str, int __size, FILE * __stream)

Read at most `size - 1` bytes from `stream`, until a newline character was encountered, and store the characters in the buffer pointed to by `str`. Unless an error was encountered while reading, the string will then be terminated with a NUL character.

If an error was encountered, the function returns NULL and sets the error flag of `stream`, which can be tested using `ferror()`. Otherwise, a pointer to the string will be returned.

23.9.4.9 int fprintf (FILE * __stream, const char * __fmt, ...)

The function `fprintf` performs formatted output to `stream`. See `vfprintf()` for details.

23.9.4.10 int fprintf_P (FILE * __stream, const char * __fmt, ...)

Variant of `fprintf()` that uses a `fmt` string that resides in program memory.

23.9.4.11 int fputc (int __c, FILE * __stream)

The function `fputc` sends the character `c` (though given as type `int`) to `stream`. It returns the character, or EOF in case an error occurred.

23.9.4.12 int fputs (const char * __str, FILE * __stream)

Write the string pointed to by `str` to stream `stream`.
Returns 0 on success and EOF on error.

23.9.4.13 int fputs_P (const char * __str, FILE * __stream)

Variant of `fputs()` where `str` resides in program memory.

23.9.4.14 size_t fread (void * __ptr, size_t __size, size_t __nmemb, FILE * __stream)

Read `nmemb` objects, `size` bytes each, from `stream`, to the buffer pointed to by `ptr`.

Returns the number of objects successfully read, i. e. `nmemb` unless an input error occurred or end-of-file was encountered. `feof()` and `ferror()` must be used to distinguish between these two conditions.

23.9.4.15 `int fscanf (FILE * __stream, const char * __fmt, ...)`

The function `fscanf` performs formatted input, reading the input data from `stream`.

See [vfscanf\(\)](#) for details.

23.9.4.16 `int fscanf_P (FILE * __stream, const char * __fmt, ...)`

Variant of [fscanf\(\)](#) using a `fmt` string in program memory.

23.9.4.17 `size_t fwrite (const void * __ptr, size_t __size, size_t __nmemb, FILE * __stream)`

Write `nmemb` objects, `size` bytes each, to `stream`. The first byte of the first object is referenced by `ptr`.

Returns the number of objects successfully written, i. e. `nmemb` unless an output error occurred.

23.9.4.18 `char* gets (char * __str)`

Similar to [fgets\(\)](#) except that it will operate on stream `stdin`, and the trailing newline (if any) will not be stored in the string. It is the caller's responsibility to provide enough storage to hold the characters read.

23.9.4.19 `int printf (const char * __fmt, ...)`

The function `printf` performs formatted output to stream `stdout`. See [vfprintf\(\)](#) for details.

23.9.4.20 `int printf_P (const char * __fmt, ...)`

Variant of [printf\(\)](#) that uses a `fmt` string that resides in program memory.

23.9.4.21 `int puts (const char * __str)`

Write the string pointed to by `str`, and a trailing newline character, to `stdout`.

23.9.4.22 `int puts_P (const char * __str)`

Variant of [puts\(\)](#) where `str` resides in program memory.

23.9.4.23 int scanf (const char * __fmt, ...)

The function `scanf` performs formatted input from stream `stdin`. See `vfscanf()` for details.

23.9.4.24 int scanf_P (const char * __fmt, ...)

Variant of `scanf()` where `fmt` resides in program memory.

23.9.4.25 int snprintf (char * __s, size_t __n, const char * __fmt, ...)

Like `sprintf()`, but instead of assuming `s` to be of infinite size, no more than `n` characters (including the trailing NUL character) will be converted to `s`.

Returns the number of characters that would have been written to `s` if there were enough space.

23.9.4.26 int snprintf_P (char * __s, size_t __n, const char * __fmt, ...)

Variant of `snprintf()` that uses a `fmt` string that resides in program memory.

23.9.4.27 int sprintf (char * __s, const char * __fmt, ...)

Variant of `printf()` that sends the formatted characters to string `s`.

23.9.4.28 int sprintf_P (char * __s, const char * __fmt, ...)

Variant of `sprintf()` that uses a `fmt` string that resides in program memory.

23.9.4.29 int sscanf (const char * __buf, const char * __fmt, ...)

The function `sscanf` performs formatted input, reading the input data from the buffer pointed to by `buf`.

See `vfscanf()` for details.

23.9.4.30 int sscanf_P (const char * __buf, const char * __fmt, ...)

Variant of `sscanf()` using a `fmt` string in program memory.

23.9.4.31 int ungetc (int __c, FILE * __stream)

The `ungetc()` function pushes the character `c` (converted to an unsigned char) back onto the input stream pointed to by `stream`. The pushed-back character will be returned by a subsequent read on the stream.

Currently, only a single character can be pushed back onto the stream.

The `ungetc()` function returns the character pushed back after the conversion, or `EOF` if the operation fails. If the value of the argument `c` character equals `EOF`, the operation will fail and the stream will remain unchanged.

23.9.4.32 `int vfprintf(FILE * __stream, const char * __fmt, va_list __ap)`

`vfprintf` is the central facility of the `printf` family of functions. It outputs values to `stream` under control of a format string passed in `fmt`. The actual values to print are passed as a variable argument list `ap`.

`vfprintf` returns the number of characters written to `stream`, or `EOF` in case of an error. Currently, this will only happen if `stream` has not been opened with write intent.

The format string is composed of zero or more directives: ordinary characters (not `%`), which are copied unchanged to the output stream; and conversion specifications, each of which results in fetching zero or more subsequent arguments. Each conversion specification is introduced by the `%` character. The arguments must properly correspond (after type promotion) with the conversion specifier. After the `%`, the following appear in sequence:

- Zero or more of the following flags:
 - `#` The value should be converted to an "alternate form". For `c`, `d`, `i`, `s`, and `u` conversions, this option has no effect. For `o` conversions, the precision of the number is increased to force the first character of the output string to a zero (except if a zero value is printed with an explicit precision of zero). For `x` and `X` conversions, a non-zero result has the string `'0x'` (or `'0X'` for `X` conversions) prepended to it.
 - `0` (zero) Zero padding. For all conversions, the converted value is padded on the left with zeros rather than blanks. If a precision is given with a numeric conversion (`d`, `i`, `o`, `u`, `i`, `x`, and `X`), the `0` flag is ignored.
 - `-` A negative field width flag; the converted value is to be left adjusted on the field boundary. The converted value is padded on the right with blanks, rather than on the left with blanks or zeros. A `-` overrides a `0` if both are given.
 - `' '` (space) A blank should be left before a positive number produced by a signed conversion (`d`, or `i`).
 - `+` A sign must always be placed before a number produced by a signed conversion. A `+` overrides a space if both are used.
- An optional decimal digit string specifying a minimum field width. If the converted value has fewer characters than the field width, it will be padded with spaces on the left (or right, if the left-adjustment flag has been given) to fill out the field width.
- An optional precision, in the form of a period `.` followed by an optional digit string. If the digit string is omitted, the precision is taken as zero. This gives the

minimum number of digits to appear for d, i, o, u, x, and X conversions, or the maximum number of characters to be printed from a string for s conversions.

- An optional l or h length modifier, that specifies that the argument for the d, i, o, u, x, or X conversion is a "long int" rather than int. The h is ignored, as "short int" is equivalent to int.
- A character that specifies the type of conversion to be applied.

The conversion specifiers and their meanings are:

- **dioxX** The int (or appropriate variant) argument is converted to signed decimal (d and i), unsigned octal (o), unsigned decimal (u), or unsigned hexadecimal (x and X) notation. The letters "abcdef" are used for x conversions; the letters "ABCDEF" are used for X conversions. The precision, if any, gives the minimum number of digits that must appear; if the converted value requires fewer digits, it is padded on the left with zeros.
- **p** The void * argument is taken as an unsigned integer, and converted similarly as a %#x command would do.
- **c** The int argument is converted to an "unsigned char", and the resulting character is written.
- **s** The "char *" argument is expected to be a pointer to an array of character type (pointer to a string). Characters from the array are written up to (but not including) a terminating NUL character; if a precision is specified, no more than the number specified are written. If a precision is given, no null character need be present; if the precision is not specified, or is greater than the size of the array, the array must contain a terminating NUL character.
- **%** A % is written. No argument is converted. The complete conversion specification is "%%".
- **eE** The double argument is rounded and converted in the format "[-]d.ddde±dd" where there is one digit before the decimal-point character and the number of digits after it is equal to the precision; if the precision is missing, it is taken as 6; if the precision is zero, no decimal-point character appears. An E conversion uses the letter 'E' (rather than 'e') to introduce the exponent. The exponent always contains two digits; if the value is zero, the exponent is 00.
- **fF** The double argument is rounded and converted to decimal notation in the format "[-]ddd.ddd", where the number of digits after the decimal-point character is equal to the precision specification. If the precision is missing, it is taken as 6; if the precision is explicitly zero, no decimal-point character appears. If a decimal point appears, at least one digit appears before it.
- **gG** The double argument is converted in style f or e (or F or E for G conversions). The precision specifies the number of significant digits. If the precision is missing, 6 digits are given; if the precision is zero, it is treated as 1. Style e is used if the exponent from its conversion is less than -4 or greater than or equal to the precision. Trailing zeros are removed from the fractional part of the result; a decimal point appears only if it is followed by at least one digit.

- S Similar to the s format, except the pointer is expected to point to a program-memory (ROM) string instead of a RAM string.

In no case does a non-existent or small field width cause truncation of a numeric field; if the result of a conversion is wider than the field width, the field is expanded to contain the conversion result.

Since the full implementation of all the mentioned features becomes fairly large, three different flavours of `fprintf()` can be selected using linker options. The default `fprintf()` implements all the mentioned functionality except floating point conversions. A minimized version of `fprintf()` is available that only implements the very basic integer and string conversion facilities, but only the # additional option can be specified using conversion flags (these flags are parsed correctly from the format specification, but then simply ignored). This version can be requested using the following [compiler options](#):

```
-Wl,-u,vfprintf -lprintf_min
```

If the full functionality including the floating point conversions is required, the following options should be used:

```
-Wl,-u,vfprintf -lprintf_flt -lm
```

Limitations:

- The specified width and precision can be at most 255.

Notes:

- For floating-point conversions, if you link default or minimized version of `fprintf()`, the symbol ? will be output and double argument will be skipped. So you output below will not be crashed. For default version the width field and the "pad to left" (symbol minus) option will work in this case.
- The hh length modifier is ignored (char argument is promoted to int). More exactly, this realization does not check the number of h symbols.
- But the ll length modifier will abort the output, as this realization does not operate long long arguments.
- The variable width or precision field (an asterisk * symbol) is not realized and will abort the output.

23.9.4.33 int fprintf_P (FILE * __stream, const char * __fmt, va_list __ap)

Variant of `fprintf()` that uses a `fmt` string that resides in program memory.

23.9.4.34 int fscanf (FILE * stream, const char * fmt, va_list ap)

Formatted input. This function is the heart of the **scanf** family of functions.

Characters are read from *stream* and processed in a way described by *fmt*. Conversion results will be assigned to the parameters passed via *ap*.

The format string *fmt* is scanned for conversion specifications. Anything that doesn't comprise a conversion specification is taken as text that is matched literally against the input. White space in the format string will match any white space in the data (including none), all other characters match only itself. Processing is aborted as soon as the data and format string no longer match, or there is an error or end-of-file condition on *stream*.

Most conversions skip leading white space before starting the actual conversion.

Conversions are introduced with the character `%`. Possible options can follow the `%`:

- a `*` indicating that the conversion should be performed but the conversion result is to be discarded; no parameters will be processed from *ap*,
- the character `h` indicating that the argument is a pointer to `short int` (rather than `int`),
- the 2 characters `hh` indicating that the argument is a pointer to `char` (rather than `int`),
- the character `l` indicating that the argument is a pointer to `long int` (rather than `int`, for integer type conversions), or a pointer to `double` (for floating point conversions),

In addition, a maximal field width may be specified as a nonzero positive decimal integer, which will restrict the conversion to at most this many characters from the input stream. This field width is limited to at most 255 characters which is also the default value (except for the `c` conversion that defaults to 1).

The following conversion flags are supported:

- `%` Matches a literal `%` character. This is not a conversion.
- `d` Matches an optionally signed decimal integer; the next pointer must be a pointer to `int`.
- `i` Matches an optionally signed integer; the next pointer must be a pointer to `int`. The integer is read in base 16 if it begins with `0x` or `0X`, in base 8 if it begins with `0`, and in base 10 otherwise. Only characters that correspond to the base are used.
- `o` Matches an octal integer; the next pointer must be a pointer to `unsigned int`.
- `u` Matches an optionally signed decimal integer; the next pointer must be a pointer to `unsigned int`.
- `x` Matches an optionally signed hexadecimal integer; the next pointer must be a pointer to `unsigned int`.
- `f` Matches an optionally signed floating-point number; the next pointer must be a pointer to `float`.

- `e`, `g`, `F`, `E`, `G` Equivalent to `f`.
- `s` Matches a sequence of non-white-space characters; the next pointer must be a pointer to `char`, and the array must be large enough to accept all the sequence and the terminating `NUL` character. The input string stops at white space or at the maximum field width, whichever occurs first.
- `c` Matches a sequence of width count characters (default 1); the next pointer must be a pointer to `char`, and there must be enough room for all the characters (no terminating `NUL` is added). The usual skip of leading white space is suppressed. To skip white space first, use an explicit space in the format.
- `[` Matches a nonempty sequence of characters from the specified set of accepted characters; the next pointer must be a pointer to `char`, and there must be enough room for all the characters in the string, plus a terminating `NUL` character. The usual skip of leading white space is suppressed. The string is to be made up of characters in (or not in) a particular set; the set is defined by the characters between the open bracket `[` character and a close bracket `]` character. The set excludes those characters if the first character after the open bracket is a circumflex `^`. To include a close bracket in the set, make it the first character after the open bracket or the circumflex; any other position will end the set. The hyphen character `-` is also special; when placed between two other characters, it adds all intervening characters to the set. To include a hyphen, make it the last character before the final close bracket. For instance, `[^]0-9-]` means the set of *everything except close bracket, zero through nine, and hyphen*. The string ends with the appearance of a character not in the (or, with a circumflex, in) set or when the field width runs out. Note that usage of this conversion enlarges the stack expense.
- `p` Matches a pointer value (as printed by `p` in `printf()`); the next pointer must be a pointer to `void`.
- `n` Nothing is expected; instead, the number of characters consumed thus far from the input is stored through the next pointer, which must be a pointer to `int`. This is not a conversion, although it can be suppressed with the `*` flag.

These functions return the number of input items assigned, which can be fewer than provided for, or even zero, in the event of a matching failure. Zero indicates that, while there was input available, no conversions were assigned; typically this is due to an invalid input character, such as an alphabetic character for a `d` conversion. The value `EOF` is returned if an input failure occurs before any conversion such as an end-of-file occurs. If an error or end-of-file occurs after conversion has begun, the number of conversions which were successfully completed is returned.

By default, all the conversions described above are available except the floating-point conversions and the width is limited to 255 characters. The float-point conversion will be available in the extended version provided by the library `libscanf_flt.a`. Also in this case the width is not limited (exactly, it is limited to 65535 characters). To link a program against the extended version, use the following compiler flags in the link stage:

```
-Wl,-u,vfscanf -lscanf_flt -lm
```

A third version is available for environments that are tight on space. In addition to the restrictions of the standard one, this version implements no `%[` specification. This version is provided in the library `libscanf_min.a`, and can be requested using the following options in the link stage:

```
-Wl,-u,vfscanf -lscanf_min -lm
```

23.9.4.35 `int vfscanf_P (FILE * __stream, const char * __fmt, va_list __ap)`

Variant of `vfscanf()` using a `fmt` string in program memory.

23.9.4.36 `int vprintf (const char * __fmt, va_list __ap)`

The function `vprintf` performs formatted output to stream `stdout`, taking a variable argument list as in `fprintf()`.

See `fprintf()` for details.

23.9.4.37 `int vscanf (const char * __fmt, va_list __ap)`

The function `vscanf` performs formatted input from stream `stdin`, taking a variable argument list as in `vfscanf()`.

See `vfscanf()` for details.

23.9.4.38 `int vsnprintf (char * __s, size_t __n, const char * __fmt, va_list ap)`

Like `vsprintf()`, but instead of assuming `s` to be of infinite size, no more than `n` characters (including the trailing NUL character) will be converted to `s`.

Returns the number of characters that would have been written to `s` if there were enough space.

23.9.4.39 `int vsnprintf_P (char * __s, size_t __n, const char * __fmt, va_list ap)`

Variant of `vsnprintf()` that uses a `fmt` string that resides in program memory.

23.9.4.40 `int vsprintf (char * __s, const char * __fmt, va_list ap)`

Like `sprintf()` but takes a variable argument list for the arguments.

23.9.4.41 `int vsprintf_P (char * __s, const char * __fmt, va_list ap)`

Variant of `vsprintf()` that uses a `fmt` string that resides in program memory.

23.10 <stdlib.h>: General utilities

Data Structures

- struct [div_t](#)
- struct [ldiv_t](#)

Defines

- #define [RAND_MAX](#) 0x7FFF

Typedefs

- typedef int(* [__compar_fn_t](#))(const void *, const void *)

Functions

- void [abort](#) (void) [__ATTR_NORETURN__](#)
- int [abs](#) (int __i)
- long [labs](#) (long __i)
- void * [bsearch](#) (const void *__key, const void *__base, size_t __nmemb, size_t __size, int(*__compar)(const void *, const void *))
- [div_t](#) [div](#) (int __num, int __denom) [__asm__](#)("__divmodhi4")
- [ldiv_t](#) [ldiv](#) (long __num, long __denom) [__asm__](#)("__divmodsi4")
- void [qsort](#) (void *__base, size_t __nmemb, size_t __size, [__compar_fn_t](#) __compar)
- long [strtol](#) (const char *__nptr, char **__endptr, int __base)
- unsigned long [strtoul](#) (const char *__nptr, char **__endptr, int __base)
- long [atol](#) (const char *__s) [__ATTR_PURE__](#)
- int [atoi](#) (const char *__s) [__ATTR_PURE__](#)
- void [exit](#) (int __status) [__ATTR_NORETURN__](#)
- void * [malloc](#) (size_t __size) [__ATTR_MALLOC__](#)
- void [free](#) (void *__ptr)
- void * [calloc](#) (size_t __nele, size_t __size) [__ATTR_MALLOC__](#)
- void * [realloc](#) (void *__ptr, size_t __size) [__ATTR_MALLOC__](#)
- double [strtod](#) (const char *__nptr, char **__endptr)
- double [atof](#) (const char *__nptr)
- int [rand](#) (void)
- void [srand](#) (unsigned int __seed)
- int [rand_r](#) (unsigned long *__ctx)

Variables

- size_t [__malloc_margin](#)
- char * [__malloc_heap_start](#)
- char * [__malloc_heap_end](#)

Non-standard (i.e. non-ISO C) functions.

- char * [ltoa](#) (long val, char *s, int radix)
- char * [utoa](#) (unsigned int val, char *s, int radix)
- char * [ultoa](#) (unsigned long val, char *s, int radix)
- long [random](#) (void)
- void [srandom](#) (unsigned long __seed)
- long [random_r](#) (unsigned long *__ctx)
- char * [itoa](#) (int val, char *s, int radix)
- #define [RANDOM_MAX](#) 0x7FFFFFFF

Conversion functions for double arguments.

Note that these functions are not located in the default library, `libc.a`, but in the mathematical library, `libm.a`. So when linking the application, the `-lm` option needs to be specified.

- char * [dtostre](#) (double __val, char *__s, unsigned char __prec, unsigned char __flags)
- char * [dtostrf](#) (double __val, signed char __width, unsigned char __prec, char *__s)
- #define [DTOSTR_ALWAYS_SIGN](#) 0x01
- #define [DTOSTR_PLUS_SIGN](#) 0x02
- #define [DTOSTR_UPPERCASE](#) 0x04
- #define [EXIT_SUCCESS](#) 0
- #define [EXIT_FAILURE](#) 1

23.10.1 Detailed Description

```
#include <stdlib.h>
```

This file declares some basic C macros and functions as defined by the ISO standard, plus some AVR-specific extensions.

23.10.2 Define Documentation**23.10.2.1 #define DTOSTR_ALWAYS_SIGN 0x01**

Bit value that can be passed in `flags` to [dtostre\(\)](#).

23.10.2.2 #define DTOSTR_PLUS_SIGN 0x02

Bit value that can be passed in `flags` to [dtostre\(\)](#).

23.10.2.3 #define DTOSTR_UPPERCASE 0x04

Bit value that can be passed in `flags` to `dtostre()`.

23.10.2.4 #define EXIT_FAILURE 1

Unsuccessful termination for `exit()`; evaluates to a non-zero value.

23.10.2.5 #define EXIT_SUCCESS 0

Successful termination for `exit()`; evaluates to 0.

23.10.2.6 #define RAND_MAX 0x7FFF

Highest number that can be generated by `rand()`.

23.10.2.7 #define RANDOM_MAX 0x7FFFFFFF

Highest number that can be generated by `random()`.

23.10.3 Typedef Documentation**23.10.3.1 typedef int(* __compar_fn_t)(const void *, const void *)**

Comparison function type for `qsort()`, just for convenience.

23.10.4 Function Documentation**23.10.4.1 void abort (void)**

The `abort()` function causes abnormal program termination to occur. This realization disables interrupts and jumps to `_exit()` function with argument equal to 1. In the limited AVR environment, execution is effectively halted by entering an infinite loop.

23.10.4.2 int abs (int __i)

The `abs()` function computes the absolute value of the integer `i`.

Note

The `abs()` and `labs()` functions are builtins of gcc.

23.10.4.3 double atof (const char * *nptr*)

The `atof()` function converts the initial portion of the string pointed to by *nptr* to double representation.

It is equivalent to calling

```
strtod(nptr, (char **)0);
```

23.10.4.4 int atoi (const char * *s*)

Convert a string to an integer.

The `atoi()` function converts the initial portion of the string pointed to by *s* to integer representation. In contrast to

```
(int)strtol(s, (char **)NULL, 10);
```

this function does not detect overflow (`errno` is not changed and the result value is not predictable), uses smaller memory (flash and stack) and works more quickly.

23.10.4.5 long atol (const char * *s*)

Convert a string to a long integer.

The `atol()` function converts the initial portion of the string pointed to by *s* to long integer representation. In contrast to

```
strtol(s, (char **)NULL, 10);
```

this function does not detect overflow (`errno` is not changed and the result value is not predictable), uses smaller memory (flash and stack) and works more quickly.

23.10.4.6 void* bsearch (const void * *__key*, const void * *__base*, size_t *__nmemb*, size_t *__size*, int(*) (const void *, const void *) *__compar*)

The `bsearch()` function searches an array of `nmemb` objects, the initial member of which is pointed to by `base`, for a member that matches the object pointed to by `key`. The size of each member of the array is specified by `size`.

The contents of the array should be in ascending sorted order according to the comparison function referenced by `compar`. The `compar` routine is expected to have two arguments which point to the key object and to an array member, in that order, and should return an integer less than, equal to, or greater than zero if the key object is found, respectively, to be less than, to match, or be greater than the array member.

The `bsearch()` function returns a pointer to a matching member of the array, or a null pointer if no match is found. If two members compare as equal, which member is matched is unspecified.

23.10.4.7 `void* calloc (size_t __nele, size_t __size)`

Allocate `nele` elements of `size` each. Identical to calling `malloc()` using `nele * size` as argument, except the allocated memory will be cleared to zero.

23.10.4.8 `div_t div (int __num, int __denom)`

The `div()` function computes the value `num/denom` and returns the quotient and remainder in a structure named `div_t` that contains two `int` members named `quot` and `rem`.

23.10.4.9 `char* dtostre (double __val, char * __s, unsigned char __prec, unsigned char __flags)`

The `dtostre()` function converts the double value passed in `val` into an ASCII representation that will be stored under `s`. The caller is responsible for providing sufficient storage in `s`.

Conversion is done in the format "`[-]d.ddde±dd`" where there is one digit before the decimal-point character and the number of digits after it is equal to the precision `prec`; if the precision is zero, no decimal-point character appears. If `flags` has the `DTOSTR_UPPERCASE` bit set, the letter '`E`' (rather than '`e`') will be used to introduce the exponent. The exponent always contains two digits; if the value is zero, the exponent is "`00`".

If `flags` has the `DTOSTR_ALWAYS_SIGN` bit set, a space character will be placed into the leading position for positive numbers.

If `flags` has the `DTOSTR_PLUS_SIGN` bit set, a plus sign will be used instead of a space character in this case.

The `dtostre()` function returns the pointer to the converted string `s`.

23.10.4.10 `char* dtostrf (double __val, signed char __width, unsigned char __prec, char * __s)`

The `dtostrf()` function converts the double value passed in `val` into an ASCII representation that will be stored under `s`. The caller is responsible for providing sufficient storage in `s`.

Conversion is done in the format "`[-]d.ddd`". The minimum field width of the output string (including the possible '`.`' and the possible sign for negative values) is given in `width`, and `prec` determines the number of digits after the decimal sign. `width` is signed value, negative for left adjustment.

The `dtostrf()` function returns the pointer to the converted string `s`.

23.10.4.11 void exit (int *__status*)

The `exit()` function terminates the application. Since there is no environment to return to, `status` is ignored, and code execution will eventually reach an infinite loop, thereby effectively halting all code processing. Before entering the infinite loop, interrupts are globally disabled.

In a C++ context, global destructors will be called before halting execution.

23.10.4.12 void free (void * *__ptr*)

The `free()` function causes the allocated memory referenced by `ptr` to be made available for future allocations. If `ptr` is NULL, no action occurs.

23.10.4.13 char* itoa (int *val*, char * *s*, int *radix*)

Convert an integer to a string.

The function `itoa()` converts the integer value from `val` into an ASCII representation that will be stored under `s`. The caller is responsible for providing sufficient storage in `s`.

Note

The minimal size of the buffer `s` depends on the choice of `radix`. For example, if the `radix` is 2 (binary), you need to supply a buffer with a minimal length of `8 * sizeof (int) + 1` characters, i.e. one character for each bit plus one for the string terminator. Using a larger `radix` will require a smaller minimal buffer size.

Warning

If the buffer is too small, you risk a buffer overflow.

Conversion is done using the `radix` as base, which may be a number between 2 (binary conversion) and up to 36. If `radix` is greater than 10, the next digit after '9' will be the letter 'a'.

If `radix` is 10 and `val` is negative, a minus sign will be prepended.

The `itoa()` function returns the pointer passed as `s`.

23.10.4.14 long labs (long *__i*)

The `labs()` function computes the absolute value of the long integer `i`.

Note

The `abs()` and `labs()` functions are builtins of gcc.

23.10.4.15 `ldiv_t ldiv (long __num, long __denom)`

The `ldiv()` function computes the value `num/denom` and returns the quotient and remainder in a structure named `ldiv_t` that contains two long integer members named `quot` and `rem`.

23.10.4.16 `char* ltoa (long val, char * s, int radix)`

Convert a long integer to a string.

The function `ltoa()` converts the long integer value from `val` into an ASCII representation that will be stored under `s`. The caller is responsible for providing sufficient storage in `s`.

Note

The minimal size of the buffer `s` depends on the choice of `radix`. For example, if the `radix` is 2 (binary), you need to supply a buffer with a minimal length of `8 * sizeof (long int) + 1` characters, i.e. one character for each bit plus one for the string terminator. Using a larger `radix` will require a smaller minimal buffer size.

Warning

If the buffer is too small, you risk a buffer overflow.

Conversion is done using the `radix` as base, which may be a number between 2 (binary conversion) and up to 36. If `radix` is greater than 10, the next digit after '9' will be the letter 'a'.

If `radix` is 10 and `val` is negative, a minus sign will be prepended.

The `ltoa()` function returns the pointer passed as `s`.

23.10.4.17 `void* malloc (size_t __size)`

The `malloc()` function allocates `size` bytes of memory. If `malloc()` fails, a NULL pointer is returned.

Note that `malloc()` does *not* initialize the returned memory to zero bytes.

See the chapter about `malloc() usage` for implementation details.

23.10.4.18 `void qsort (void * __base, size_t __nmemb, size_t __size, __compar_fn_t __compar)`

The `qsort()` function is a modified partition-exchange sort, or quicksort.

The `qsort()` function sorts an array of `nmemb` objects, the initial member of which is pointed to by `base`. The size of each object is specified by `size`. The contents of the

array base are sorted in ascending order according to a comparison function pointed to by `compar`, which requires two arguments pointing to the objects being compared.

The comparison function must return an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second.

23.10.4.19 `int rand (void)`

The `rand()` function computes a sequence of pseudo-random integers in the range of 0 to `RAND_MAX` (as defined by the header file <stdlib.h>).

The `srand()` function sets its argument `seed` as the seed for a new sequence of pseudo-random numbers to be returned by `rand()`. These sequences are repeatable by calling `srand()` with the same seed value.

If no seed value is provided, the functions are automatically seeded with a value of 1.

In compliance with the C standard, these functions operate on `int` arguments. Since the underlying algorithm already uses 32-bit calculations, this causes a loss of precision. See `random()` for an alternate set of functions that retains full 32-bit precision.

23.10.4.20 `int rand_r (unsigned long * __ctx)`

Variant of `rand()` that stores the context in the user-supplied variable located at `ctx` instead of a static library variable so the function becomes re-entrant.

23.10.4.21 `long random (void)`

The `random()` function computes a sequence of pseudo-random integers in the range of 0 to `RANDOM_MAX` (as defined by the header file <stdlib.h>).

The `srandom()` function sets its argument `seed` as the seed for a new sequence of pseudo-random numbers to be returned by `random()`. These sequences are repeatable by calling `srandom()` with the same seed value.

If no seed value is provided, the functions are automatically seeded with a value of 1.

23.10.4.22 `long random_r (unsigned long * __ctx)`

Variant of `random()` that stores the context in the user-supplied variable located at `ctx` instead of a static library variable so the function becomes re-entrant.

23.10.4.23 `void* realloc (void * __ptr, size_t __size)`

The `realloc()` function tries to change the size of the region allocated at `ptr` to the new `size` value. It returns a pointer to the new region. The returned pointer might be the same as the old pointer, or a pointer to a completely different region.

The contents of the returned region up to either the old or the new size value (whatever is less) will be identical to the contents of the old region, even in case a new region had to be allocated.

It is acceptable to pass `ptr` as `NULL`, in which case `realloc()` will behave identical to `malloc()`.

If the new memory cannot be allocated, `realloc()` returns `NULL`, and the region at `ptr` will not be changed.

23.10.4.24 `void srand (unsigned int __seed)`

Pseudo-random number generator seeding; see `rand()`.

23.10.4.25 `void srandom (unsigned long __seed)`

Pseudo-random number generator seeding; see `random()`.

23.10.4.26 `double strtod (const char * nptr, char ** endptr)`

The `strtod()` function converts the initial portion of the string pointed to by *nptr* to double representation.

The expected form of the string is an optional plus ('+') or minus sign ('-') followed by a sequence of digits optionally containing a decimal-point character, optionally followed by an exponent. An exponent consists of an 'E' or 'e', followed by an optional plus or minus sign, followed by a sequence of digits.

Leading white-space characters in the string are skipped.

The `strtod()` function returns the converted value, if any.

If *endptr* is not `NULL`, a pointer to the character after the last character used in the conversion is stored in the location referenced by *endptr*.

If no conversion is performed, zero is returned and the value of *nptr* is stored in the location referenced by *endptr*.

If the correct value would cause overflow, plus or minus `INFINITY` is returned (according to the sign of the value), and `ERANGE` is stored in `errno`. If the correct value would cause underflow, zero is returned and `ERANGE` is stored in `errno`.

23.10.4.27 `long strtol (const char * nptr, char ** endptr, int base)`

The `strtol()` function converts the string in *nptr* to a long value. The conversion is done according to the given base, which must be between 2 and 36 inclusive, or be the special value 0.

The string may begin with an arbitrary amount of white space (as determined by `isspace()`) followed by a single optional '+' or '-' sign. If *base* is zero or 16, the string may then include a "0x" prefix, and the number will be read in base 16; otherwise, a

zero base is taken as 10 (decimal) unless the next character is '0', in which case it is taken as 8 (octal).

The remainder of the string is converted to a long value in the obvious manner, stopping at the first character which is not a valid digit in the given base. (In bases above 10, the letter 'A' in either upper or lower case represents 10, 'B' represents 11, and so forth, with 'Z' representing 35.)

If `endptr` is not NULL, `strtol()` stores the address of the first invalid character in `*endptr`. If there were no digits at all, however, `strtol()` stores the original value of `nptr` in `endptr`. (Thus, if `*nptr` is not '\0' but `**endptr` is '\0' on return, the entire string was valid.)

The `strtol()` function returns the result of the conversion, unless the value would underflow or overflow. If no conversion could be performed, 0 is returned. If an overflow or underflow occurs, `errno` is set to `ERANGE` and the function return value is clamped to `LONG_MIN` or `LONG_MAX`, respectively.

23.10.4.28 unsigned long strtoul (const char * __nptr, char ** __endptr, int __base)

The `strtoul()` function converts the string in `nptr` to an unsigned long value. The conversion is done according to the given base, which must be between 2 and 36 inclusive, or be the special value 0.

The string may begin with an arbitrary amount of white space (as determined by `isspace()`) followed by a single optional '+' or '-' sign. If `base` is zero or 16, the string may then include a "0x" prefix, and the number will be read in base 16; otherwise, a zero base is taken as 10 (decimal) unless the next character is '0', in which case it is taken as 8 (octal).

The remainder of the string is converted to an unsigned long value in the obvious manner, stopping at the first character which is not a valid digit in the given base. (In bases above 10, the letter 'A' in either upper or lower case represents 10, 'B' represents 11, and so forth, with 'Z' representing 35.)

If `endptr` is not NULL, `strtoul()` stores the address of the first invalid character in `*endptr`. If there were no digits at all, however, `strtoul()` stores the original value of `nptr` in `endptr`. (Thus, if `*nptr` is not '\0' but `**endptr` is '\0' on return, the entire string was valid.)

The `strtoul()` function return either the result of the conversion or, if there was a leading minus sign, the negation of the result of the conversion, unless the original (non-negated) value would overflow; in the latter case, `strtoul()` returns `ULONG_MAX`, and `errno` is set to `ERANGE`. If no conversion could be performed, 0 is returned.

23.10.4.29 char* ultoa (unsigned long val, char * s, int radix)

Convert an unsigned long integer to a string.

The function `ultoa()` converts the unsigned long integer value from `val` into an ASCII

representation that will be stored under *s*. The caller is responsible for providing sufficient storage in *s*.

Note

The minimal size of the buffer *s* depends on the choice of radix. For example, if the radix is 2 (binary), you need to supply a buffer with a minimal length of $8 * \text{sizeof}(\text{unsigned long int}) + 1$ characters, i.e. one character for each bit plus one for the string terminator. Using a larger radix will require a smaller minimal buffer size.

Warning

If the buffer is too small, you risk a buffer overflow.

Conversion is done using the *radix* as base, which may be a number between 2 (binary conversion) and up to 36. If *radix* is greater than 10, the next digit after '9' will be the letter 'a'.

The `ultoa()` function returns the pointer passed as *s*.

23.10.4.30 char* utoa (unsigned int val, char *s, int radix)

Convert an unsigned integer to a string.

The function `utoa()` converts the unsigned integer value from *val* into an ASCII representation that will be stored under *s*. The caller is responsible for providing sufficient storage in *s*.

Note

The minimal size of the buffer *s* depends on the choice of radix. For example, if the radix is 2 (binary), you need to supply a buffer with a minimal length of $8 * \text{sizeof}(\text{unsigned int}) + 1$ characters, i.e. one character for each bit plus one for the string terminator. Using a larger radix will require a smaller minimal buffer size.

Warning

If the buffer is too small, you risk a buffer overflow.

Conversion is done using the *radix* as base, which may be a number between 2 (binary conversion) and up to 36. If *radix* is greater than 10, the next digit after '9' will be the letter 'a'.

The `utoa()` function returns the pointer passed as *s*.

23.10.5 Variable Documentation**23.10.5.1 char* __malloc_heap_end**

`malloc()` tunable.

23.10.5.2 char* __malloc_heap_start

`malloc()` tunable.

23.10.5.3 size_t __malloc_margin

`malloc()` tunable.

23.11 <string.h>: Strings

Defines

- `#define _FFS(x)`

Functions

- `int ffs (int __val)`
- `int ffsll (long __val)`
- `__extension__ int ffsll (long long __val)`
- `void * memccpy (void *, const void *, int, size_t)`
- `void * memchr (const void *, int, size_t) __ATTR_PURE__`
- `int memcmp (const void *, const void *, size_t) __ATTR_PURE__`
- `void * memcpy (void *, const void *, size_t)`
- `void * memmem (const void *, size_t, const void *, size_t) __ATTR_PURE__`
- `void * memmove (void *, const void *, size_t)`
- `void * memrchr (const void *, int, size_t) __ATTR_PURE__`
- `void * memset (void *, int, size_t)`
- `char * strcat (char *, const char *)`
- `char * strchr (const char *, int) __ATTR_PURE__`
- `char * strchrnul (const char *, int) __ATTR_PURE__`
- `int strcmp (const char *, const char *) __ATTR_PURE__`
- `char * strcpy (char *, const char *)`
- `int strcasecmp (const char *, const char *) __ATTR_PURE__`
- `char * strcasestr (const char *, const char *) __ATTR_PURE__`
- `size_t strcspn (const char *__s, const char *__reject) __ATTR_PURE__`
- `char * strdup (const char *s1)`
- `size_t strlcat (char *, const char *, size_t)`
- `size_t strlcpy (char *, const char *, size_t)`
- `size_t strlen (const char *) __ATTR_PURE__`
- `char * strlwr (char *)`
- `char * strncat (char *, const char *, size_t)`
- `int strncmp (const char *, const char *, size_t) __ATTR_PURE__`
- `char * strncpy (char *, const char *, size_t)`
- `int strncasecmp (const char *, const char *, size_t) __ATTR_PURE__`
- `size_t strnlen (const char *, size_t) __ATTR_PURE__`

- char * [strpbrk](#) (const char *__s, const char *__accept) __ATTR_PURE__
- char * [strchr](#) (const char *, int) __ATTR_PURE__
- char * [strrev](#) (char *)
- char * [strsep](#) (char **, const char *)
- size_t [strspn](#) (const char *__s, const char *__accept) __ATTR_PURE__
- char * [strstr](#) (const char *, const char *) __ATTR_PURE__
- char * [strtok](#) (char *, const char *)
- char * [strtok_r](#) (char *, const char *, char **)
- char * [strupr](#) (char *)

23.11.1 Detailed Description

```
#include <string.h>
```

The string functions perform string operations on NULL terminated strings.

Note

If the strings you are working on resident in program space (flash), you will need to use the string functions described in [<avr/pgmspace.h>: Program Space Utilities](#).

23.11.2 Define Documentation

23.11.2.1 #define _FFS(x)

This macro finds the first (least significant) bit set in the input value.

This macro is very similar to the function [ffs\(\)](#) except that it evaluates its argument at compile-time, so it should only be applied to compile-time constant expressions where it will reduce to a constant itself. Application of this macro to expressions that are not constant at compile-time is not recommended, and might result in a huge amount of code generated.

Returns

The [_FFS\(\)](#) macro returns the position of the first (least significant) bit set in the word val, or 0 if no bits are set. The least significant bit is position 1. Only 16 bits of argument are evaluated.

23.11.3 Function Documentation

23.11.3.1 int ffs (int val)

This function finds the first (least significant) bit set in the input value.

Returns

The [ffs\(\)](#) function returns the position of the first (least significant) bit set in the word val, or 0 if no bits are set. The least significant bit is position 1.

Note

For expressions that are constant at compile time, consider using the `_FFS` macro instead.

23.11.3.2 int ffs1 (long __val)

Same as `ffs()`, for an argument of type `long`.

23.11.3.3 int ffs11 (long long __val)

Same as `ffs()`, for an argument of type `long long`.

23.11.3.4 void * memccpy (void * *dest*, const void * *src*, int *val*, size_t *len*)

Copy memory area.

The `memccpy()` function copies no more than `len` bytes from memory area `src` to memory area `dest`, stopping when the character `val` is found.

Returns

The `memccpy()` function returns a pointer to the next character in `dest` after `val`, or `NULL` if `val` was not found in the first `len` characters of `src`.

23.11.3.5 void * memchr (const void * *src*, int *val*, size_t *len*)

Scan memory for a character.

The `memchr()` function scans the first `len` bytes of the memory area pointed to by `src` for the character `val`. The first byte to match `val` (interpreted as an unsigned character) stops the operation.

Returns

The `memchr()` function returns a pointer to the matching byte or `NULL` if the character does not occur in the given memory area.

23.11.3.6 `int memcmp (const void * s1, const void * s2, size_t len)`

Compare memory areas.

The `memcmp()` function compares the first `len` bytes of the memory areas `s1` and `s2`. The comparison is performed using unsigned char operations.

Returns

The `memcmp()` function returns an integer less than, equal to, or greater than zero if the first `len` bytes of `s1` is found, respectively, to be less than, to match, or be greater than the first `len` bytes of `s2`.

Note

Be sure to store the result in a 16 bit variable since you may get incorrect results if you use an unsigned char or char due to truncation.

Warning

This function is not -mint8 compatible, although if you only care about testing for equality, this function should be safe to use.

23.11.3.7 `void * memcpy (void * dest, const void * src, size_t len)`

Copy a memory area.

The `memcpy()` function copies `len` bytes from memory area `src` to memory area `dest`. The memory areas may not overlap. Use `memmove()` if the memory areas do overlap.

Returns

The `memcpy()` function returns a pointer to `dest`.

23.11.3.8 `void * memmem (const void * s1, size_t len1, const void * s2, size_t len2)`

The `memmem()` function finds the start of the first occurrence of the substring `s2` of length `len2` in the memory area `s1` of length `len1`.

Returns

The `memmem()` function returns a pointer to the beginning of the substring, or `NULL` if the substring is not found. If `len2` is zero, the function returns `s1`.

23.11.3.9 void * memmove (void * *dest*, const void * *src*, size_t *len*)

Copy memory area.

The [memmove\(\)](#) function copies *len* bytes from memory area *src* to memory area *dest*. The memory areas may overlap.

Returns

The [memmove\(\)](#) function returns a pointer to *dest*.

23.11.3.10 void * memrchr (const void * *src*, int *val*, size_t *len*)

The [memrchr\(\)](#) function is like the [memchr\(\)](#) function, except that it searches backwards from the end of the *len* bytes pointed to by *src* instead of forwards from the front. (Glibc, GNU extension.)

Returns

The [memrchr\(\)](#) function returns a pointer to the matching byte or NULL if the character does not occur in the given memory area.

23.11.3.11 void * memset (void * *dest*, int *val*, size_t *len*)

Fill memory with a constant byte.

The [memset\(\)](#) function fills the first *len* bytes of the memory area pointed to by *dest* with the constant byte *val*.

Returns

The [memset\(\)](#) function returns a pointer to the memory area *dest*.

23.11.3.12 int strcasecmp (const char * *s1*, const char * *s2*)

Compare two strings ignoring case.

The [strcasecmp\(\)](#) function compares the two strings *s1* and *s2*, ignoring the case of the characters.

Returns

The [strcasecmp\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* is found, respectively, to be less than, to match, or be greater than *s2*. A consequence of the ordering used by [strcasecmp\(\)](#) is that if *s1* is an initial substring of *s2*, then *s1* is considered to be "less than" *s2*.

23.11.3.13 char * strcasestr (const char * s1, const char * s2)

The [strcasestr\(\)](#) function finds the first occurrence of the substring *s2* in the string *s1*. This is like [strstr\(\)](#), except that it ignores case of alphabetic symbols in searching for the substring. (Glibc, GNU extension.)

Returns

The [strcasestr\(\)](#) function returns a pointer to the beginning of the substring, or `NULL` if the substring is not found. If *s2* points to a string of zero length, the function returns *s1*.

23.11.3.14 char * strcat (char * dest, const char * src)

Concatenate two strings.

The [strcat\(\)](#) function appends the *src* string to the *dest* string overwriting the `'\0'` character at the end of *dest*, and then adds a terminating `'\0'` character. The strings may not overlap, and the *dest* string must have enough space for the result.

Returns

The [strcat\(\)](#) function returns a pointer to the resulting string *dest*.

23.11.3.15 char * strchr (const char * src, int val)

Locate character in string.

The [strchr\(\)](#) function returns a pointer to the first occurrence of the character *val* in the string *src*.

Here "character" means "byte" - these functions do not work with wide or multi-byte characters.

Returns

The [strchr\(\)](#) function returns a pointer to the matched character or `NULL` if the character is not found.

23.11.3.16 char * strchrnul (const char * s, int c)

The [strchrnul\(\)](#) function is like [strchr\(\)](#) except that if *c* is not found in *s*, then it returns a pointer to the null byte at the end of *s*, rather than `NULL`. (Glibc, GNU extension.)

Returns

The `strchrnul()` function returns a pointer to the matched character, or a pointer to the null byte at the end of `s` (i.e., `s+strlen(s)`) if the character is not found.

23.11.3.17 int strcmp (const char * s1, const char * s2)

Compare two strings.

The `strcmp()` function compares the two strings `s1` and `s2`.

Returns

The `strcmp()` function returns an integer less than, equal to, or greater than zero if `s1` is found, respectively, to be less than, to match, or be greater than `s2`. A consequence of the ordering used by `strcmp()` is that if `s1` is an initial substring of `s2`, then `s1` is considered to be "less than" `s2`.

23.11.3.18 char * strcpy (char * dest, const char * src)

Copy a string.

The `strcpy()` function copies the string pointed to by `src` (including the terminating `'\0'` character) to the array pointed to by `dest`. The strings may not overlap, and the destination string `dest` must be large enough to receive the copy.

Returns

The `strcpy()` function returns a pointer to the destination string `dest`.

Note

If the destination string of a `strcpy()` is not large enough (that is, if the programmer was stupid/lazy, and failed to check the size before copying) then anything might happen. Overflowing fixed length strings is a favourite cracker technique.

23.11.3.19 size_t strcspn (const char * s, const char * reject)

The `strcspn()` function calculates the length of the initial segment of `s` which consists entirely of characters not in `reject`.

Returns

The `strcspn()` function returns the number of characters in the initial segment of `s` which are not in the string `reject`. The terminating zero is not considered as a part of string.

23.11.3.20 char * strdup (const char * s1)

Duplicate a string.

The `strdup()` function allocates memory and copies into it the string addressed by `s1`, including the terminating null character.

Warning

The `strdup()` function calls `malloc()` to allocate the memory for the duplicated string! The user is responsible for freeing the memory by calling `free()`.

Returns

The `strdup()` function returns a pointer to the resulting string `dest`. If `malloc()` cannot allocate enough storage for the string, `strdup()` will return `NULL`.

Warning

Be sure to check the return value of the `strdup()` function to make sure that the function has succeeded in allocating the memory!

23.11.3.21 size_t strlcat (char * dst, const char * src, size_t siz)

Concatenate two strings.

Appends `src` to string `dst` of size `siz` (unlike `strncat()`, `siz` is the full size of `dst`, not space left). At most `siz-1` characters will be copied. Always `NULL` terminates (unless `siz <= strlen(dst)`).

Returns

The `strlcat()` function returns `strlen(src) + MIN(siz, strlen(initial dst))`. If `retval >= siz`, truncation occurred.

Appends `src` to string `dst` of size `siz` (unlike `strncat()`, `siz` is the full size of `dst`, not space left). At most `siz-1` characters will be copied. Always `NULL` terminates (unless `siz <= strlen(dst)`).

Returns

The `strlcat()` function returns `strlen(src) + MIN(siz, strlen(initial dst))`. If `retval >= siz`, truncation occurred.

23.11.3.22 size_t strlcpy (char * dst, const char * src, size_t siz)

Copy a string.

Copy `src` to string `dst` of size `siz`. At most `siz-1` characters will be copied. Always NULL terminates (unless `siz == 0`).

Returns

The `strcpy()` function returns `strlen(src)`. If `retval >= siz`, truncation occurred.

Copy `src` to string `dst` of size `siz`. At most `siz-1` characters will be copied. Always NULL terminates (unless `siz == 0`).

Returns

The `strncpy()` function returns `strlen(src)`. If `retval >= siz`, truncation occurred.

23.11.3.23 `size_t strlen (const char * src)`

Calculate the length of a string.

The `strlen()` function calculates the length of the string `src`, not including the terminating `'\0'` character.

Returns

The `strlen()` function returns the number of characters in `src`.

23.11.3.24 `char * strlwr (char * s)`

Convert a string to lower case.

The `strlwr()` function will convert a string to lower case. Only the upper case alphabetic characters [A .. Z] are converted. Non-alphabetic characters will not be changed.

Returns

The `strlwr()` function returns a pointer to the converted string.

23.11.3.25 `int strncasecmp (const char * s1, const char * s2, size_t len)`

Compare two strings ignoring case.

The `strncasecmp()` function is similar to `strcasecmp()`, except it only compares the first `len` characters of `s1`.

Returns

The [strncasecmp\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* (or the first *len* bytes thereof) is found, respectively, to be less than, to match, or be greater than *s2*. A consequence of the ordering used by [strncasecmp\(\)](#) is that if *s1* is an initial substring of *s2*, then *s1* is considered to be "less than" *s2*.

23.11.3.26 char * strncat (char * *dest*, const char * *src*, size_t *len*)

Concatenate two strings.

The [strncat\(\)](#) function is similar to [strcat\(\)](#), except that only the first *n* characters of *src* are appended to *dest*.

Returns

The [strncat\(\)](#) function returns a pointer to the resulting string *dest*.

23.11.3.27 int strncmp (const char * *s1*, const char * *s2*, size_t *len*)

Compare two strings.

The [strncmp\(\)](#) function is similar to [strcmp\(\)](#), except it only compares the first (at most) *n* characters of *s1* and *s2*.

Returns

The [strncmp\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* (or the first *n* bytes thereof) is found, respectively, to be less than, to match, or be greater than *s2*.

23.11.3.28 char * strncpy (char * *dest*, const char * *src*, size_t *len*)

Copy a string.

The [strncpy\(\)](#) function is similar to [strcpy\(\)](#), except that not more than *n* bytes of *src* are copied. Thus, if there is no null byte among the first *n* bytes of *src*, the result will not be null-terminated.

In the case where the length of *src* is less than that of *n*, the remainder of *dest* will be padded with nulls.

Returns

The [strncpy\(\)](#) function returns a pointer to the destination string *dest*.

23.11.3.29 `size_t strlen (const char * src, size_t len)`

Determine the length of a fixed-size string.

The `strlen` function returns the number of characters in the string pointed to by `src`, not including the terminating `'\0'` character, but at most `len`. In doing this, `strlen` looks only at the first `len` characters at `src` and never beyond `src+len`.

Returns

The `strlen` function returns `strlen(src)`, if that is less than `len`, or `len` if there is no `'\0'` character among the first `len` characters pointed to by `src`.

23.11.3.30 `char * strpbrk (const char * s, const char * accept)`

The `strpbrk()` function locates the first occurrence in the string `s` of any of the characters in the string `accept`.

Returns

The `strpbrk()` function returns a pointer to the character in `s` that matches one of the characters in `accept`, or `NULL` if no such character is found. The terminating zero is not considered as a part of string: if one or both args are empty, the result will be `NULL`.

23.11.3.31 `char * strrchr (const char * src, int val)`

Locate character in string.

The `strrchr()` function returns a pointer to the last occurrence of the character `val` in the string `src`.

Here "character" means "byte" - these functions do not work with wide or multi-byte characters.

Returns

The `strrchr()` function returns a pointer to the matched character or `NULL` if the character is not found.

23.11.3.32 `char * strrev (char * s)`

Reverse a string.

The `strrev()` function reverses the order of the string.

Returns

The `strrev()` function returns a pointer to the beginning of the reversed string.

23.11.3.33 `char * strsep (char ** sp, const char * delim)`

Parse a string into tokens.

The `strsep()` function locates, in the string referenced by `*sp`, the first occurrence of any character in the string `delim` (or the terminating `'\0'` character) and replaces it with a `'\0'`. The location of the next character after the delimiter character (or `NULL`, if the end of the string was reached) is stored in `*sp`. An “empty” field, i.e. one caused by two adjacent delimiter characters, can be detected by comparing the location referenced by the pointer returned in `*sp` to `'\0'`.

Returns

The `strsep()` function returns a pointer to the original value of `*sp`. If `*sp` is initially `NULL`, `strsep()` returns `NULL`.

23.11.3.34 `size_t strspn (const char * s, const char * accept)`

The `strspn()` function calculates the length of the initial segment of `s` which consists entirely of characters in `accept`.

Returns

The `strspn()` function returns the number of characters in the initial segment of `s` which consist only of characters from `accept`. The terminating zero is not considered as a part of string.

23.11.3.35 `char * strstr (const char * s1, const char * s2)`

Locate a substring.

The `strstr()` function finds the first occurrence of the substring `s2` in the string `s1`. The terminating `'\0'` characters are not compared.

Returns

The `strstr()` function returns a pointer to the beginning of the substring, or `NULL` if the substring is not found. If `s2` points to a string of zero length, the function returns `s1`.

23.11.3.36 `char * strtok (char * s, const char * delim)`

Parses the string *s* into tokens.

`strtok` parses the string *s* into tokens. The first call to `strtok` should have *s* as its first argument. Subsequent calls should have the first argument set to `NULL`. If a token ends with a delimiter, this delimiting character is overwritten with a `'\0'` and a pointer to the next character is saved for the next call to `strtok`. The delimiter string *delim* may be different for each call.

Returns

The `strtok()` function returns a pointer to the next token or `NULL` when no more tokens are found.

Note

`strtok()` is NOT reentrant. For a reentrant version of this function see `strtok_r()`.

23.11.3.37 `char * strtok_r (char * string, const char * delim, char ** last)`

Parses *string* into tokens.

`strtok_r` parses *string* into tokens. The first call to `strtok_r` should have *string* as its first argument. Subsequent calls should have the first argument set to `NULL`. If a token ends with a delimiter, this delimiting character is overwritten with a `'\0'` and a pointer to the next character is saved for the next call to `strtok_r`. The delimiter string *delim* may be different for each call. *last* is a user allocated `char*` pointer. It must be the same while parsing the same string. `strtok_r` is a reentrant version of `strtok()`.

Returns

The `strtok_r()` function returns a pointer to the next token or `NULL` when no more tokens are found.

23.11.3.38 `char * strupr (char * s)`

Convert a string to upper case.

The `strupr()` function will convert a string to upper case. Only the lower case alphabetic characters [*a .. z*] are converted. Non-alphabetic characters will not be changed.

Returns

The `strupr()` function returns a pointer to the converted string. The pointer is the same as that passed in since the operation is performed in place.

23.12 <time.h>: Time

Data Structures

- struct [tm](#)
- struct [week_date](#)

Defines

- #define [ONE_HOUR](#) 3600
- #define [ONE_DEGREE](#) 3600
- #define [ONE_DAY](#) 86400
- #define [UNIX_OFFSET](#) 946684800
- #define [NTP_OFFSET](#) 3155673600

Typedefs

- typedef [uint32_t](#) [time_t](#)

Enumerations

- enum [_WEEK_DAYS_](#) {
 SUNDAY, MONDAY, TUESDAY, WEDNESDAY,
 THURSDAY, FRIDAY, SATURDAY }
- enum [_MONTHS_](#) {
 JANUARY, FEBRUARY, MARCH, APRIL,
 MAY, JUNE, JULY, AUGUST,
 SEPTEMBER, OCTOBER, NOVEMBER, DECEMBER }

Functions

- [time_t](#) [time](#) ([time_t](#) *timer)
- [int32_t](#) [difftime](#) ([time_t](#) time1, [time_t](#) time0)
- [time_t](#) [mktime](#) (struct [tm](#) *timeptr)
- [time_t](#) [mk_gmtime](#) (const struct [tm](#) *timeptr)
- struct [tm](#) * [gmtime](#) (const [time_t](#) *timer)
- void [gmtime_r](#) (const [time_t](#) *timer, struct [tm](#) *timeptr)
- struct [tm](#) * [localtime](#) (const [time_t](#) *timer)
- void [localtime_r](#) (const [time_t](#) *timer, struct [tm](#) *timeptr)
- char * [asctime](#) (const struct [tm](#) *timeptr)
- void [asctime_r](#) (const struct [tm](#) *timeptr, char *buf)
- char * [ctime](#) (const [time_t](#) *timer)
- void [ctime_r](#) (const [time_t](#) *timer, char *buf)
- char * [isotime](#) (const struct [tm](#) *tmptr)

- void `isotime_r` (const struct `tm` *, char *)
- size_t `strftime` (char *s, size_t maxsize, const char *format, const struct `tm` *timeptr)
- void `set_dst` (int(*) (const `time_t` *, `int32_t` *))
- void `set_zone` (`int32_t`)
- void `set_system_time` (`time_t` timestamp)
- void `system_tick` (void)
- `uint8_t` `is_leap_year` (`int16_t` year)
- `uint8_t` `month_length` (`int16_t` year, `uint8_t` month)
- `uint8_t` `week_of_year` (const struct `tm` *timeptr, `uint8_t` start)
- `uint8_t` `week_of_month` (const struct `tm` *timeptr, `uint8_t` start)
- struct `week_date` * `iso_week_date` (int year, int yday)
- void `iso_week_date_r` (int year, int yday, struct `week_date` *)
- `uint32_t` `fatfs_time` (const struct `tm` *timeptr)
- void `set_position` (`int32_t` latitude, `int32_t` longitude)
- `int16_t` `equation_of_time` (const `time_t` *timer)
- `int32_t` `daylight_seconds` (const `time_t` *timer)
- `time_t` `solar_noon` (const `time_t` *timer)
- `time_t` `sun_rise` (const `time_t` *timer)
- `time_t` `sun_set` (const `time_t` *timer)
- double `solar_declination` (const `time_t` *timer)
- `int8_t` `moon_phase` (const `time_t` *timer)
- unsigned long `gm_sidereal` (const `time_t` *timer)
- unsigned long `lm_sidereal` (const `time_t` *timer)

23.12.1 Detailed Description

```
#include <time.h>
```

Introduction to the Time functions This file declares the time functions implemented in `avr-libc`.

The implementation aspires to conform with ISO/IEC 9899 (C90). However, due to limitations of the target processor and the nature of its development environment, a practical implementation must of necessity deviate from the standard.

Section 7.23.2.1 `clock()` The type `clock_t`, the macro `CLOCKS_PER_SEC`, and the function `clock()` are not implemented. We consider these items belong to operating system code, or to application code when no operating system is present.

Section 7.23.2.3 `mktime()` The standard specifies that `mktime()` should return (`time_t`) -1, if the time cannot be represented. This implementation always returns a 'best effort' representation.

Section 7.23.2.4 `time()` The standard specifies that `time()` should return (`time_t`) -1, if the time is not available. Since the application must initialize the time system, this functionality is not implemented.

Section 7.23.2.2, [difftime\(\)](#) Due to the lack of a 64 bit double, the function [difftime\(\)](#) returns a long integer. In most cases this change will be invisible to the user, handled automatically by the compiler.

Section 7.23.1.4 struct tm Per the standard, struct tm->tm_isdst is greater than zero when Daylight Saving time is in effect. This implementation further specifies that, when positive, the value of tm_isdst represents the amount time is advanced during Daylight Saving time.

Section 7.23.3.5 [strftime\(\)](#) Only the 'C' locale is supported, therefore the modifiers 'E' and 'O' are ignored. The 'Z' conversion is also ignored, due to the lack of time zone name.

In addition to the above departures from the standard, there are some behaviors which are different from what is often expected, though allowed under the standard.

There is no 'platform standard' method to obtain the current time, time zone, or daylight savings 'rules' in the AVR environment. Therefore the application must initialize the time system with this information. The functions [set_zone\(\)](#), [set_dst\(\)](#), and [set_system_time\(\)](#) are provided for initialization. Once initialized, system time is maintained by calling the function [system_tick\(\)](#) at one second intervals.

Though not specified in the standard, it is often expected that time_t is a signed integer representing an offset in seconds from Midnight Jan 1 1970... i.e. 'Unix time'. This implementation uses an unsigned 32 bit integer offset from Midnight Jan 1 2000. The use of this 'epoch' helps to simplify the conversion functions, while the 32 bit value allows time to be properly represented until Tue Feb 7 06:28:15 2136 UTC. The macros UNIX_OFFSET and NTP_OFFSET are defined to assist in converting to and from Unix and NTP time stamps.

Unlike desktop counterparts, it is impractical to implement or maintain the 'zoneinfo' database. Therefore no attempt is made to account for time zone, daylight saving, or leap seconds in past dates. All calculations are made according to the currently configured time zone and daylight saving 'rule'.

In addition to C standard functions, re-entrant versions of [ctime\(\)](#), [asctime\(\)](#), [gmtime\(\)](#) and [localtime\(\)](#) are provided which, in addition to being re-entrant, have the property of claiming less permanent storage in RAM. An additional time conversion, [isotime\(\)](#) and its re-entrant version, uses far less storage than either [ctime\(\)](#) or [asctime\(\)](#).

Along with the usual smattering of utility functions, such as [is_leap_year\(\)](#), this library includes a set of functions related the sun and moon, as well as sidereal time functions.

23.12.2 Define Documentation

23.12.2.1 #define NTP_OFFSET 3155673600

Difference between the Y2K and the NTP epochs, in seconds. To convert a Y2K timestamp to NTP...

```
unsigned long ntp;
time_t y2k;

y2k = time(NULL);
ntp = y2k + NTP_OFFSET;
```

23.12.2.2 #define ONE_DAY 86400

One day, expressed in seconds

23.12.2.3 #define ONE_DEGREE 3600

Angular degree, expressed in arc seconds

23.12.2.4 #define ONE_HOUR 3600

One hour, expressed in seconds

23.12.2.5 #define UNIX_OFFSET 946684800

Difference between the Y2K and the UNIX epochs, in seconds. To convert a Y2K timestamp to UNIX...

```
long unix;  
time_t y2k;  
  
y2k = time(NULL);  
unix = y2k + UNIX_OFFSET;
```

23.12.3 Typedef Documentation**23.12.3.1 typedef uint32_t time_t**

time_t represents seconds elapsed from Midnight, Jan 1 2000 UTC (the Y2K 'epoch'). Its range allows this implementation to represent time up to Tue Feb 7 06:28:15 2136 UTC.

23.12.4 Enumeration Type Documentation**23.12.4.1 enum _MONTHS_**

Enumerated labels for the months.

23.12.4.2 enum _WEEK_DAYS_

Enumerated labels for the days of the week.

23.12.5 Function Documentation

23.12.5.1 `char* asctime (const struct tm * timeptr)`

The `asctime` function converts the broken-down time of `timeptr`, into an ascii string in the form

Sun Mar 23 01:03:52 2013

23.12.5.2 `void asctime_r (const struct tm * timeptr, char * buf)`

Re entrant version of [asctime\(\)](#).

23.12.5.3 `char* ctime (const time_t * timer)`

The `ctime` function is equivalent to `asctime(localtime(timer))`

23.12.5.4 `void ctime_r (const time_t * timer, char * buf)`

Re entrant version of [ctime\(\)](#).

23.12.5.5 `int32_t daylight_seconds (const time_t * timer)`

Computes the amount of time the sun is above the horizon, at the location of the observer.

NOTE: At observer locations inside a polar circle, this value can be zero during the winter, and can exceed `ONE_DAY` during the summer.

The returned value is in seconds.

23.12.5.6 `int32_t difftime (time_t time1, time_t time0)`

The `difftime` function returns the difference between two binary time stamps, `time1` - `time0`.

23.12.5.7 `int16_t equation_of_time (const time_t * timer)`

Computes the difference between apparent solar time and mean solar time. The returned value is in seconds.

23.12.5.8 `uint32_t fatfs_time (const struct tm * timeptr)`

Convert a Y2K time stamp into a FAT file system time stamp.

23.12.5.9 unsigned long gm_sidereal (const time_t * timer)

Returns Greenwich Mean Sidereal Time, as seconds into the sidereal day. The returned value will range from 0 through 86399 seconds.

23.12.5.10 struct tm* gmtime (const time_t * timer) [read]

The gmtime function converts the time stamp pointed to by timer into broken-down time, expressed as UTC.

23.12.5.11 void gmtime_r (const time_t * timer, struct tm * timeptr)

Re entrant version of [gmtime\(\)](#).

23.12.5.12 uint8_t is_leap_year (int16_t year)

Return 1 if year is a leap year, zero if it is not.

23.12.5.13 struct week_date* iso_week_date (int year, int yday) [read]

Return a [week_date](#) structure with the ISO_8601 week based date corresponding to the given year and day of year. See http://en.wikipedia.org/wiki/ISO_week_date for more information.

23.12.5.14 void iso_week_date_r (int year, int yday, struct week_date *)

Re-entrant version of iso-week_date.

23.12.5.15 char* isotime (const struct tm * tmptr)

The isotime function constructs an ascii string in the form

```
2013-03-23 01:03:52
```

23.12.5.16 void isotime_r (const struct tm *, char *)

Re entrant version of [isotime\(\)](#)

23.12.5.17 unsigned long lm_sidereal (const time_t * timer)

Returns Local Mean Sidereal Time, as seconds into the sidereal day. The returned value will range from 0 through 86399 seconds.

23.12.5.18 struct tm* localtime (const time_t * *timer*) [read]

The localtime function converts the time stamp pointed to by *timer* into broken-down time, expressed as Local time.

23.12.5.19 void localtime_r (const time_t * *timer*, struct tm * *timeptr*)

Re entrant version of [localtime\(\)](#).

23.12.5.20 time_t mktime (const struct tm * *timeptr*)

This function 'compiles' the elements of a broken-down time structure, returning a binary time stamp. The elements of *timeptr* are interpreted as representing UTC.

The original values of the *tm_wday* and *tm_yday* elements of the structure are ignored, and the original values of the other elements are not restricted to the ranges stated for struct tm.

Unlike [mktime\(\)](#), this function DOES NOT modify the elements of *timeptr*.

23.12.5.21 time_t mktime (struct tm * *timeptr*)

This function 'compiles' the elements of a broken-down time structure, returning a binary time stamp. The elements of *timeptr* are interpreted as representing Local Time.

The original values of the *tm_wday* and *tm_yday* elements of the structure are ignored, and the original values of the other elements are not restricted to the ranges stated for struct tm.

On successful completion, the values of all elements of *timeptr* are set to the appropriate range.

23.12.5.22 uint8_t month_length (int16_t *year*, uint8_t *month*)

Return the length of month, given the year and month, where month is in the range 1 to 12.

23.12.5.23 int8_t moon_phase (const time_t * *timer*)

Returns an approximation to the phase of the moon. The sign of the returned value indicates a waning or waxing phase. The magnitude of the returned value indicates the percentage illumination.

23.12.5.24 void set_dst (int(*) (const time_t *, int32_t *))

Specify the Daylight Saving function.

The Daylight Saving function should examine its parameters to determine whether Daylight Saving is in effect, and return a value appropriate for `tm_isdst`.

Working examples for the USA and the EU are available..

```
#include <util/eu_dst.h>
```

for the European Union, and

```
#include <util/usa_dst.h>
```

for the United States

If a Daylight Saving function is not specified, the system will ignore Daylight Saving.

23.12.5.25 void set_position (int32_t latitude, int32_t longitude)

Set the geographic coordinates of the 'observer', for use with several of the following functions. Parameters are passed as seconds of North Latitude, and seconds of East Longitude.

For New York City...

```
set_position( 40.7142 * ONE_DEGREE, -74.0064 * ONE_DEGREE);
```

23.12.5.26 void set_system_time (time_t timestamp)

Initialize the system time. Examples are...

From a Clock / Calendar type RTC:

```
struct tm rtc_time;

read_rtc(&rtc_time);
rtc_time.tm_isdst = 0;
set_system_time( mktime(&rtc_time) );
```

From a Network Time Protocol time stamp:

```
set_system_time( ntp_timestamp - NTP_OFFSET );
```

From a UNIX time stamp:

```
set_system_time( unix_timestamp - UNIX_OFFSET );
```

23.12.5.27 void set_zone (int32_t)

Set the 'time zone'. The parameter is given in seconds East of the Prime Meridian. Example for New York City:

```
set_zone( -5 * ONE_HOUR );
```

If the time zone is not set, the time system will operate in UTC only.

23.12.5.28 double solar_declination (const time_t * timer)

Returns the declination of the sun in radians.

23.12.5.29 time_t solar_noon (const time_t * timer)

Computes the time of solar noon, at the location of the observer.

23.12.5.30 size_t strftime (char * s, size_t maxsize, const char * format, const struct tm * timeptr)

A complete description of [strftime\(\)](#) is beyond the pale of this document. Refer to ISO/IEC document 9899 for details.

All conversions are made using the 'C Locale', ignoring the E or O modifiers. Due to the lack of a time zone 'name', the 'Z' conversion is also ignored.

23.12.5.31 time_t sun_rise (const time_t * timer)

Return the time of sunrise, at the location of the observer. See the note about [daylight_seconds\(\)](#).

23.12.5.32 time_t sun_set (const time_t * timer)

Return the time of sunset, at the location of the observer. See the note about [daylight_seconds\(\)](#).

23.12.5.33 void system_tick (void)

Maintain the system time by calling this function at a rate of 1 Hertz.

It is anticipated that this function will typically be called from within an Interrupt Service Routine, (though that is not required). It therefore includes code which makes it simple to use from within a 'Naked' ISR, avoiding the cost of saving and restoring all the cpu registers.

Such an ISR may resemble the following example...

```
ISR(RTC_OVF_vect, ISR_NAKED)
{
    system_tick();
    reti();
}
```

23.12.5.34 time_t time (time_t * timer)

The time function returns the systems current time stamp. If timer is not a null pointer, the return value is also assigned to the object it points to.

23.12.5.35 uint8_t week_of_month (const struct tm * *timeptr*, uint8_t *start*)

Return the calendar week of month, where the first week is considered to begin on the day of week specified by 'start'. The returned value may range from zero to 5.

23.12.5.36 uint8_t week_of_year (const struct tm * *timeptr*, uint8_t *start*)

Return the calendar week of year, where week 1 is considered to begin on the day of week specified by 'start'. The returned value may range from zero to 52.

23.13 <avr/boot.h>: Bootloader Support Utilities**Defines**

- #define `BOOTLOADER_SECTION` __attribute__((section(".bootloader")))
- #define `boot_spm_interrupt_enable()` (__SPM_REG |= (uint8_t)_BV(SPMIE))
- #define `boot_spm_interrupt_disable()` (__SPM_REG &= (uint8_t)~_BV(SPMIE))
- #define `boot_is_spm_interrupt()` (__SPM_REG & (uint8_t)_BV(SPMIE))
- #define `boot_rww_busy()` (__SPM_REG & (uint8_t)_BV(__COMMON_ASB))
- #define `boot_spm_busy()` (__SPM_REG & (uint8_t)_BV(__SPM_ENABLE))
- #define `boot_spm_busy_wait()` do{ }while(boot_spm_busy())
- #define `GET_LOW_FUSE_BITS` (0x0000)
- #define `GET_LOCK_BITS` (0x0001)
- #define `GET_EXTENDED_FUSE_BITS` (0x0002)
- #define `GET_HIGH_FUSE_BITS` (0x0003)
- #define `boot_lock_fuse_bits_get(address)`
- #define `boot_signature_byte_get(addr)`
- #define `boot_page_fill(address, data)` __boot_page_fill_normal(address, data)
- #define `boot_page_erase(address)` __boot_page_erase_normal(address)
- #define `boot_page_write(address)` __boot_page_write_normal(address)
- #define `boot_rww_enable()` __boot_rww_enable()
- #define `boot_lock_bits_set(lock_bits)` __boot_lock_bits_set(lock_bits)
- #define `boot_page_fill_safe(address, data)`
- #define `boot_page_erase_safe(address)`
- #define `boot_page_write_safe(address)`
- #define `boot_rww_enable_safe()`
- #define `boot_lock_bits_set_safe(lock_bits)`

23.13.1 Detailed Description

```
#include <avr/io.h>
#include <avr/boot.h>
```

The macros in this module provide a C language interface to the bootloader support functionality of certain AVR processors. These macros are designed to work with all sizes of flash memory.

Global interrupts are not automatically disabled for these macros. It is left up to the programmer to do this. See the code example below. Also see the processor datasheet for caveats on having global interrupts enabled during writing of the Flash.

Note

Not all AVR processors provide bootloader support. See your processor datasheet to see if it provides bootloader support.

Todo

From email with Marek: On smaller devices (all except ATmega64/128), __SPM__REG is in the I/O space, accessible with the shorter "in" and "out" instructions - since the boot loader has a limited size, this could be an important optimization.

API Usage Example

The following code shows typical usage of the boot API.

```
#include <inttypes.h>
#include <avr/interrupt.h>
#include <avr/pgmspace.h>

void boot_program_page (uint32_t page, uint8_t *buf)
{
    uint16_t i;
    uint8_t sreg;

    // Disable interrupts.

    sreg = SREG;
    cli();

    eeprom_busy_wait ();

    boot_page_erase (page);
    boot_spm_busy_wait ();      // Wait until the memory is erased.

    for (i=0; i<SPM_PAGESIZE; i+=2)
    {
        // Set up little-endian word.

        uint16_t w = *buf++;
        w += (*buf++) << 8;

        boot_page_fill (page + i, w);
    }

    boot_page_write (page);      // Store buffer in flash page.
    boot_spm_busy_wait();      // Wait until the memory is written.

    // Reenable RWW-section again. We need this if we want to jump back
    // to the application after bootloading.

    boot_rww_enable ();
}
```

```

        // Re-enable interrupts (if they were ever enabled).

        SREG = sreg;
    }

```

23.13.2 Define Documentation

23.13.2.1 `#define boot_is_spm_interrupt() (__SPM_REG & (uint8_t)_BV(SPMIE))`

Check if the SPM interrupt is enabled.

23.13.2.2 `#define boot_lock_bits_set(lock_bits) __boot_lock_bits_set(lock_bits)`

Set the bootloader lock bits.

Parameters

lock_bits A mask of which Boot Loader Lock Bits to set.

Note

In this context, a 'set bit' will be written to a zero value. Note also that only BLBxx bits can be programmed by this command.

For example, to disallow the SPM instruction from writing to the Boot Loader memory section of flash, you would use this macro as such:

```
boot_lock_bits_set (_BV (BLB11));
```

Note

Like any lock bits, the Boot Loader Lock Bits, once set, cannot be cleared again except by a chip erase which will in turn also erase the boot loader itself.

23.13.2.3 `#define boot_lock_bits_set_safe(lock_bits)`

Value:

```

do { \
    boot_spm_busy_wait(); \
    eeprom_busy_wait(); \
    boot_lock_bits_set (lock_bits); \
} while (0)

```

Same as [boot_lock_bits_set\(\)](#) except waits for eeprom and spm operations to complete before setting the lock bits.

23.13.2.4 #define boot_lock_fuse_bits_get(address)**Value:**

```

(__extension__({
    uint8_t __result;
    __asm__ __volatile__
    (
        "sts %1, %2\n\t"
        "lpm %0, Z\n\t"
        : "=r" (__result)
        : "i" (_SFR_MEM_ADDR(__SPM_REG)),
          "r" ((uint8_t)(__BOOT_LOCK_BITS_SET)),
          "z" ((uint16_t)(address))
        );
    __result;
}))

```

Read the lock or fuse bits at address.

Parameter `address` can be any of `GET_LOW_FUSE_BITS`, `GET_LOCK_BITS`, `GET_EXTENDED_FUSE_BITS`, or `GET_HIGH_FUSE_BITS`.

Note

The lock and fuse bits returned are the physical values, i.e. a bit returned as 0 means the corresponding fuse or lock bit is programmed.

23.13.2.5 #define boot_page_erase(address) __boot_page_erase_normal(address)

Erase the flash page that contains address.

Note

`address` is a byte address in flash, not a word address.

23.13.2.6 #define boot_page_erase_safe(address)**Value:**

```

do { \
    boot_spm_busy_wait(); \
    eeprom_busy_wait(); \
    boot_page_erase (address); \
} while (0)

```

Same as `boot_page_erase()` except it waits for eeprom and spm operations to complete before erasing the page.

23.13.2.7 #define boot_page_fill(address, data) __boot_page_fill_normal(address, data)

Fill the bootloader temporary page buffer for flash address with data word.

Note

The address is a byte address. The data is a word. The AVR writes data to the buffer a word at a time, but addresses the buffer per byte! So, increment your address by 2 between calls, and send 2 data bytes in a word format! The LSB of the data is written to the lower address; the MSB of the data is written to the higher address.

23.13.2.8 #define boot_page_fill_safe(address, data)**Value:**

```
do { \
    boot_spm_busy_wait();           \
    eeprom_busy_wait();             \
    boot_page_fill(address, data);   \
} while (0)
```

Same as [boot_page_fill\(\)](#) except it waits for eeprom and spm operations to complete before filling the page.

23.13.2.9 #define boot_page_write(address) __boot_page_write_normal(address)

Write the bootloader temporary page buffer to flash page that contains address.

Note

address is a byte address in flash, not a word address.

23.13.2.10 #define boot_page_write_safe(address)**Value:**

```
do { \
    boot_spm_busy_wait();           \
    eeprom_busy_wait();             \
    boot_page_write (address);       \
} while (0)
```

Same as [boot_page_write\(\)](#) except it waits for eeprom and spm operations to complete before writing the page.

23.13.2.11 `#define boot_rww_busy() (__SPM_REG & (uint8_t)_BV(__COMMON_ASB))`

Check if the RWW section is busy.

23.13.2.12 `#define boot_rww_enable() __boot_rww_enable()`

Enable the Read-While-Write memory section.

23.13.2.13 `#define boot_rww_enable_safe()`

Value:

```
do { \
    boot_spm_busy_wait();           \
    eeprom_busy_wait();             \
    boot_rww_enable();              \
} while (0)
```

Same as `boot_rww_enable()` except waits for eeprom and spm operations to complete before enabling the RWW mameory.

23.13.2.14 `#define boot_signature_byte_get(addr)`

Value:

```
(__extension__({ \
    uint8_t __result; \
    __asm__ __volatile__ \
    ( \
        "sts %1, %2\n\t" \
        "lpm %0, Z" "\n\t" \
        : "=r" (__result) \
        : "i" (_SFR_MEM_ADDR(__SPM_REG)), \
          "r" ((uint8_t)(__BOOT_SIGROW_READ)), \
          "z" ((uint16_t)(addr)) \
        ); \
    __result; \
}))
```

Read the Signature Row byte at `address`. For some MCU types, this function can also retrieve the factory-stored oscillator calibration bytes.

Parameter `address` can be 0-0x1f as documented by the datasheet.

Note

The values are MCU type dependent.

23.13.2.15 `#define boot_spm_busy() (__SPM_REG & (uint8_t)_BV(__SPM_ENABLE))`

Check if the SPM instruction is busy.

23.13.2.16 `#define boot_spm_busy_wait() do{while(boot_spm_busy())`

Wait while the SPM instruction is busy.

23.13.2.17 `#define boot_spm_interrupt_disable() (__SPM_REG &=`
`(uint8_t)~_BV(SPMIE))`

Disable the SPM interrupt.

23.13.2.18 `#define boot_spm_interrupt_enable() (__SPM_REG |=`
`(uint8_t)_BV(SPMIE))`

Enable the SPM interrupt.

23.13.2.19 `#define BOOTLOADER_SECTION __attribute__((section`
`(".bootloader")))`

Used to declare a function or variable to be placed into a new section called .bootloader. This section and its contents can then be relocated to any address (such as the bootloader NRWW area) at link-time.

23.13.2.20 `#define GET_EXTENDED_FUSE_BITS (0x0002)`

address to read the extended fuse bits, using boot_lock_fuse_bits_get

23.13.2.21 `#define GET_HIGH_FUSE_BITS (0x0003)`

address to read the high fuse bits, using boot_lock_fuse_bits_get

23.13.2.22 `#define GET_LOCK_BITS (0x0001)`

address to read the lock bits, using boot_lock_fuse_bits_get

23.13.2.23 `#define GET_LOW_FUSE_BITS (0x0000)`

address to read the low fuse bits, using boot_lock_fuse_bits_get

23.14 <avr/cpufunc.h>: Special AVR CPU functions

Defines

- `#define _NOP()`
- `#define _MemoryBarrier()`

Functions

- void `ccp_write_io` (`uint8_t * __ioaddr`, `uint8_t __value`)

23.14.1 Detailed Description

```
#include <avr/cpufunc.h>
```

This header file contains macros that access special functions of the AVR CPU which do not fit into any of the other header files.

23.14.2 Define Documentation

23.14.2.1 #define _MemoryBarrier()

Implement a read/write *memory barrier*. A memory barrier instructs the compiler to not cache any memory data in registers beyond the barrier. This can sometimes be more effective than blocking certain optimizations by declaring some object with a `volatile` qualifier.

See [Problems with reordering code](#) for things to be taken into account with respect to compiler optimizations.

23.14.2.2 #define _NOP()

Execute a *no operation* (NOP) CPU instruction. This should not be used to implement delays, better use the functions from [<util/delay_basic.h>](#) or [<util/delay.h>](#) for this. For debugging purposes, a NOP can be useful to have an instruction that is guaranteed to be not optimized away by the compiler, so it can always become a breakpoint in the debugger.

23.14.3 Function Documentation

23.14.3.1 void ccp_write_io (uint8_t * __ioaddr, uint8_t __value)

Write `__value` to Configuration Change Protected (CCP) IO register at `__ioaddr`.

23.15 <avr/eeprom.h>: EEPROM handling

Defines

- #define `EEMEM` `__attribute__((section(".eeprom")))`
- #define `eeprom_is_ready`()
- #define `eeprom_busy_wait`() `do { } while (!eeprom_is_ready())`

Functions

- `uint8_t eeprom_read_byte` (const `uint8_t *`__p) `__ATTR_PURE__`
- `uint16_t eeprom_read_word` (const `uint16_t *`__p) `__ATTR_PURE__`
- `uint32_t eeprom_read_dword` (const `uint32_t *`__p) `__ATTR_PURE__`
- `float eeprom_read_float` (const `float *`__p) `__ATTR_PURE__`
- `void eeprom_read_block` (void *__dst, const void *__src, `size_t` __n)
- `void eeprom_write_byte` (`uint8_t *`__p, `uint8_t` __value)
- `void eeprom_write_word` (`uint16_t *`__p, `uint16_t` __value)
- `void eeprom_write_dword` (`uint32_t *`__p, `uint32_t` __value)
- `void eeprom_write_float` (`float *`__p, `float` __value)
- `void eeprom_write_block` (const void *__src, void *__dst, `size_t` __n)
- `void eeprom_update_byte` (`uint8_t *`__p, `uint8_t` __value)
- `void eeprom_update_word` (`uint16_t *`__p, `uint16_t` __value)
- `void eeprom_update_dword` (`uint32_t *`__p, `uint32_t` __value)
- `void eeprom_update_float` (`float *`__p, `float` __value)
- `void eeprom_update_block` (const void *__src, void *__dst, `size_t` __n)

IAR C compatibility defines

- `#define _EPUT(addr, val) eeprom_write_byte ((uint8_t *) (addr), (uint8_t) (val))`
- `#define __EPUT(addr, val) eeprom_write_byte ((uint8_t *) (addr), (uint8_t) (val))`
- `#define _EGET(var, addr) (var) = eeprom_read_byte ((const uint8_t *) (addr))`
- `#define __EGET(var, addr) (var) = eeprom_read_byte ((const uint8_t *) (addr))`

23.15.1 Detailed Description

```
#include <avr/eeprom.h>
```

This header file declares the interface to some simple library routines suitable for handling the data EEPROM contained in the AVR microcontrollers. The implementation uses a simple polled mode interface. Applications that require interrupt-controlled EEPROM access to ensure that no time will be wasted in spinloops will have to deploy their own implementation.

Notes:

- In addition to the write functions there is a set of update ones. These functions read each byte first and skip the burning if the old value is the same with new. The scanning direction is from high address to low, to obtain quick return in common cases.

- All of the read/write functions first make sure the EEPROM is ready to be accessed. Since this may cause long delays if a write operation is still pending, time-critical applications should first poll the EEPROM e. g. using `eeprom_is_ready()` before attempting any actual I/O. But this functions are not wait until SELFPRGEN in SPMCSR becomes zero. Do this manually, if your software contains the Flash burning.
- As these functions modify IO registers, they are known to be non-reentrant. If any of these functions are used from both, standard and interrupt context, the applications must ensure proper protection (e.g. by disabling interrupts before accessing them).
- All write functions force erase_and_write programming mode.
- For Xmega the EEPROM start address is 0, like other architectures. The reading functions add the 0x2000 value to use EEPROM mapping into data space.

23.15.2 Define Documentation

23.15.2.1 `#define __EEGET(var, addr) (var) = eeprom_read_byte ((const uint8_t *) (addr))`

Read a byte from EEPROM. Compatibility define for IAR C.

23.15.2.2 `#define __EEPUT(addr, val) eeprom_write_byte ((uint8_t *) (addr), (uint8_t) (val))`

Write a byte to EEPROM. Compatibility define for IAR C.

23.15.2.3 `#define _EEGET(var, addr) (var) = eeprom_read_byte ((const uint8_t *) (addr))`

Read a byte from EEPROM. Compatibility define for IAR C.

23.15.2.4 `#define _EEPUT(addr, val) eeprom_write_byte ((uint8_t *) (addr), (uint8_t) (val))`

Write a byte to EEPROM. Compatibility define for IAR C.

23.15.2.5 `#define EEMEM __attribute__((section(".eeprom")))`

Attribute expression causing a variable to be allocated within the .eeprom section.

23.15.2.6 #define eeprom_busy_wait() do {} while (!eeprom_is_ready())

Loops until the eeprom is no longer busy.

Returns

Nothing.

23.15.2.7 #define eeprom_is_ready()**Returns**

1 if EEPROM is ready for a new read/write operation, 0 if not.

23.15.3 Function Documentation**23.15.3.1 void eeprom_read_block (void * __dst, const void * __src, size_t __n)**

Read a block of __n bytes from EEPROM address __src to SRAM __dst.

23.15.3.2 uint8_t eeprom_read_byte (const uint8_t * __p)

Read one byte from EEPROM address __p.

23.15.3.3 uint32_t eeprom_read_dword (const uint32_t * __p)

Read one 32-bit double word (little endian) from EEPROM address __p.

23.15.3.4 float eeprom_read_float (const float * __p)

Read one float value (little endian) from EEPROM address __p.

23.15.3.5 uint16_t eeprom_read_word (const uint16_t * __p)

Read one 16-bit word (little endian) from EEPROM address __p.

23.15.3.6 void eeprom_update_block (const void * __src, void * __dst, size_t __n)

Update a block of __n bytes to EEPROM address __dst from __src.

Note

The argument order is mismatch with common functions like [strcpy\(\)](#).

23.15.3.7 void eeprom_update_byte (uint8_t * __p, uint8_t __value)

Update a byte __value to EEPROM address __p.

23.15.3.8 void eeprom_update_dword (uint32_t * __p, uint32_t __value)

Update a 32-bit double word __value to EEPROM address __p.

23.15.3.9 void eeprom_update_float (float * __p, float __value)

Update a float __value to EEPROM address __p.

23.15.3.10 void eeprom_update_word (uint16_t * __p, uint16_t __value)

Update a word __value to EEPROM address __p.

23.15.3.11 void eeprom_write_block (const void * __src, void * __dst, size_t __n)

Write a block of __n bytes to EEPROM address __dst from __src.

Note

The argument order is mismatch with common functions like [strcpy\(\)](#).

23.15.3.12 void eeprom_write_byte (uint8_t * __p, uint8_t __value)

Write a byte __value to EEPROM address __p.

23.15.3.13 void eeprom_write_dword (uint32_t * __p, uint32_t __value)

Write a 32-bit double word __value to EEPROM address __p.

23.15.3.14 void eeprom_write_float (float * __p, float __value)

Write a float __value to EEPROM address __p.

23.15.3.15 void eeprom_write_word (uint16_t * __p, uint16_t __value)

Write a word __value to EEPROM address __p.

23.16 <avr/fuse.h>: Fuse Support

Introduction

The Fuse API allows a user to specify the fuse settings for the specific AVR device they are compiling for. These fuse settings will be placed in a special section in the ELF output file, after linking.

Programming tools can take advantage of the fuse information embedded in the ELF file, by extracting this information and determining if the fuses need to be programmed before programming the Flash and EEPROM memories. This also allows a single ELF file to contain all the information needed to program an AVR.

To use the Fuse API, include the <avr/io.h> header file, which in turn automatically includes the individual I/O header file and the <avr/fuse.h> file. These other two files provides everything necessary to set the AVR fuses.

Fuse API

Each I/O header file must define the FUSE_MEMORY_SIZE macro which is defined to the number of fuse bytes that exist in the AVR device.

A new type, `__fuse_t`, is defined as a structure. The number of fields in this structure are determined by the number of fuse bytes in the FUSE_MEMORY_SIZE macro.

If FUSE_MEMORY_SIZE == 1, there is only a single field: byte, of type unsigned char.

If FUSE_MEMORY_SIZE == 2, there are two fields: low, and high, of type unsigned char.

If FUSE_MEMORY_SIZE == 3, there are three fields: low, high, and extended, of type unsigned char.

If FUSE_MEMORY_SIZE > 3, there is a single field: byte, which is an array of unsigned char with the size of the array being FUSE_MEMORY_SIZE.

A convenience macro, FUSEMEM, is defined as a GCC attribute for a custom-named section of ".fuse".

A convenience macro, FUSES, is defined that declares a variable, `__fuse`, of type `__fuse_t` with the attribute defined by FUSEMEM. This variable allows the end user to easily set the fuse data.

Note

If a device-specific I/O header file has previously defined FUSEMEM, then FUSEMEM is not redefined. If a device-specific I/O header file has previously defined FUSES, then FUSES is not redefined.

Each AVR device I/O header file has a set of defined macros which specify the actual fuse bits available on that device. The AVR fuses have inverted values, logical 1 for

an unprogrammed (disabled) bit and logical 0 for a programmed (enabled) bit. The defined macros for each individual fuse bit represent this in their definition by a bitwise inversion of a mask. For example, the FUSE_EESAVE fuse in the ATmega128 is defined as:

```
#define FUSE_EESAVE      ~_BV(3)
```

Note

The `_BV` macro creates a bit mask from a bit number. It is then inverted to represent logical values for a fuse memory byte.

To combine the fuse bits macros together to represent a whole fuse byte, use the bitwise AND operator, like so:

```
(FUSE_BOOTSZ0 & FUSE_BOOTSZ1 & FUSE_EESAVE & FUSE_SPIEN & FUSE_JTAGEN)
```

Each device I/O header file also defines macros that provide default values for each fuse byte that is available. `LFUSE_DEFAULT` is defined for a Low Fuse byte. `HFUSE_DEFAULT` is defined for a High Fuse byte. `EFUSE_DEFAULT` is defined for an Extended Fuse byte.

If `FUSE_MEMORY_SIZE > 3`, then the I/O header file defines macros that provide default values for each fuse byte like so: `FUSE0_DEFAULT` `FUSE1_DEFAULT` `FUSE2_DEFAULT` `FUSE3_DEFAULT` `FUSE4_DEFAULT`

API Usage Example

Putting all of this together is easy. Using C99's designated initializers:

```
#include <avr/io.h>

FUSES =
{
    .low = LFUSE_DEFAULT,
    .high = (FUSE_BOOTSZ0 & FUSE_BOOTSZ1 & FUSE_EESAVE & FUSE_SPIEN & FUSE_JTAGEN),
    .extended = EFUSE_DEFAULT,
};

int main(void)
{
    return 0;
}
```

Or, using the variable directly instead of the FUSES macro,

```
#include <avr/io.h>

__fuse_t __fuse __attribute__((section (".fuse"))) =
{
    .low = LFUSE_DEFAULT,
    .high = (FUSE_BOOTSZ0 & FUSE_BOOTSZ1 & FUSE_EESAVE & FUSE_SPIEN & FUSE_JTAGEN)
}
```

```

    AGEN),
    .extended = EFUSE_DEFAULT,
};

int main(void)
{
    return 0;
}

```

If you are compiling in C++, you cannot use the designated initializers so you must do:

```

#include <avr/io.h>

FUSES =
{
    LFUSE_DEFAULT, // .low
    (FUSE_BOOTSZ0 & FUSE_BOOTSZ1 & FUSE_EESAVE & FUSE_SPIEN & FUSE_JTAGEN), /
    / .high
    EFUSE_DEFAULT, // .extended
};

int main(void)
{
    return 0;
}

```

However there are a number of caveats that you need to be aware of to use this API properly.

Be sure to include <avr/io.h> to get all of the definitions for the API. The FUSES macro defines a global variable to store the fuse data. This variable is assigned to its own linker section. Assign the desired fuse values immediately in the variable initialization.

The .fuse section in the ELF file will get its values from the initial variable assignment ONLY. This means that you can NOT assign values to this variable in functions and the new values will not be put into the ELF .fuse section.

The global variable is declared in the FUSES macro has two leading underscores, which means that it is reserved for the "implementation", meaning the library, so it will not conflict with a user-named variable.

You must initialize ALL fields in the __fuse_t structure. This is because the fuse bits in all bytes default to a logical 1, meaning unprogrammed. Normal uninitialized data defaults to all logical zeros. So it is vital that all fuse bytes are initialized, even with default data. If they are not, then the fuse bits may not be programmed to the desired settings.

Be sure to have the -mmcu=*device* flag in your compile command line and your linker command line to have the correct device selected and to have the correct I/O header file included when you include <avr/io.h>.

You can print out the contents of the .fuse section in the ELF file by using this command line:

```
avr-objdump -s -j .fuse <ELF file>
```

The section contents shows the address on the left, then the data going from lower address to a higher address, left to right.

23.17 <avr/interrupt.h>: Interrupts

Global manipulation of the interrupt flag

The global interrupt flag is maintained in the I bit of the status register (SREG).

Handling interrupts frequently requires attention regarding atomic access to objects that could be altered by code running within an interrupt context, see <[util/atomic.h](#)>.

Frequently, interrupts are being disabled for periods of time in order to perform certain operations without being disturbed; see [Problems with reordering code](#) for things to be taken into account with respect to compiler optimizations.

- #define [sei\(\)](#)
- #define [cli\(\)](#)

Macros for writing interrupt handler functions

- #define [ISR](#)(vector, attributes)
- #define [SIGNAL](#)(vector)
- #define [EMPTY_INTERRUPT](#)(vector)
- #define [ISR_ALIAS](#)(vector, target_vector)
- #define [reti](#)()
- #define [BADISR_vect](#)

ISR attributes

- #define [ISR_BLOCK](#)
- #define [ISR_NOBLOCK](#)
- #define [ISR_NAKED](#)
- #define [ISR_ALIASOF](#)(target_vector)

23.17.1 Detailed Description

Note

This discussion of interrupts was originally taken from Rich Neswold's document. See [Acknowledgments](#).

Introduction to avr-libc's interrupt handling It's nearly impossible to find compilers that agree on how to handle interrupt code. Since the C language tries to stay away from machine dependent details, each compiler writer is forced to design their method of support.

In the AVR-GCC environment, the vector table is predefined to point to interrupt routines with predetermined names. By using the appropriate name, your routine will be called when the corresponding interrupt occurs. The device library provides a set of default interrupt routines, which will get used if you don't define your own.

Patching into the vector table is only one part of the problem. The compiler uses, by convention, a set of registers when it's normally executing compiler-generated code. It's important that these registers, as well as the status register, get saved and restored. The extra code needed to do this is enabled by tagging the interrupt function with `__attribute__((signal))`.

These details seem to make interrupt routines a little messy, but all these details are handled by the Interrupt API. An interrupt routine is defined with `ISR()`. This macro registers and marks the routine as an interrupt handler for the specified peripheral. The following is an example definition of a handler for the ADC interrupt.

```
#include <avr/interrupt.h>

ISR(ADC_vect)
{
    // user code here
}
```

Refer to the chapter explaining [assembler programming](#) for an explanation about interrupt routines written solely in assembler language.

Catch-all interrupt vector If an unexpected interrupt occurs (interrupt is enabled and no handler is installed, which usually indicates a bug), then the default action is to reset the device by jumping to the reset vector. You can override this by supplying a function named `BADISR_vect` which should be defined with `ISR()` as such. (The name `BADISR_vect` is actually an alias for `__vector_default`. The latter must be used inside assembly code in case `<avr/interrupt.h>` is not included.)

```
#include <avr/interrupt.h>

ISR(BADISR_vect)
{
    // user code here
}
```

Nested interrupts The AVR hardware clears the global interrupt flag in SREG before entering an interrupt vector. Thus, normally interrupts will remain disabled inside the handler until the handler exits, where the `RETI` instruction (that is emitted by the compiler as part of the normal function epilogue for an interrupt handler) will eventually re-enable further interrupts. For that reason, interrupt handlers normally do not nest. For most interrupt handlers, this is the desired behaviour, for some it is even required in order to prevent infinitely recursive interrupts (like UART interrupts, or level-triggered external interrupts). In rare circumstances though it might be desired to re-enable the global interrupt flag as early as possible in the interrupt handler, in order to not defer any other interrupt more than absolutely needed. This could be done using an `sei()` instruction right at the beginning of the interrupt handler, but this still leaves

few instructions inside the compiler-generated function prologue to run with global interrupts disabled. The compiler can be instructed to insert an SEI instruction right at the beginning of an interrupt handler by declaring the handler the following way:

```
ISR(XXX_vect, ISR_NOBLOCK)
{
    ...
}
```

where `XXX_vect` is the name of a valid interrupt vector for the MCU type in question, as explained below.

Two vectors sharing the same code In some circumstances, the actions to be taken upon two different interrupts might be completely identical so a single implementation for the ISR would suffice. For example, pin-change interrupts arriving from two different ports could logically signal an event that is independent from the actual port (and thus interrupt vector) where it happened. Sharing interrupt vector code can be accomplished using the [ISR_ALIASOF\(\)](#) attribute to the ISR macro:

```
ISR(PCINT0_vect)
{
    ...
    // Code to handle the event.
}

ISR(PCINT1_vect, ISR_ALIASOF(PCINT0_vect));
```

Note

There is no body to the aliased ISR.

Note that the [ISR_ALIASOF\(\)](#) feature requires GCC 4.2 or above (or a patched version of GCC 4.1.x). See the documentation of the [ISR_ALIAS\(\)](#) macro for an implementation which is less elegant but could be applied to all compiler versions.

Empty interrupt service routines In rare circumstances, an interrupt vector does not need any code to be implemented at all. The vector must be declared anyway, so when the interrupt triggers it won't execute the `BADISR_vect` code (which by default restarts the application).

This could for example be the case for interrupts that are solely enabled for the purpose of getting the controller out of [sleep_mode\(\)](#).

A handler for such an interrupt vector can be declared using the [EMPTY_INTERRUPT\(\)](#) macro:

```
EMPTY_INTERRUPT(ADC_vect);
```

Note

There is no body to this macro.

Manually defined ISRs In some circumstances, the compiler-generated prologue and epilogue of the ISR might not be optimal for the job, and a manually defined ISR could be considered particularly to speedup the interrupt handling.

One solution to this could be to implement the entire ISR as manual assembly code in a separate (assembly) file. See [Combining C and assembly source files](#) for an example of how to implement it that way.

Another solution is to still implement the ISR in C language but take over the compiler's job of generating the prologue and epilogue. This can be done using the `ISR_NAKED` attribute to the `ISR()` macro. Note that the compiler does not generate *anything* as prologue or epilogue, so the final `reti()` must be provided by the actual implementation. `SREG` must be manually saved if the ISR code modifies it, and the compiler-implied assumption of `__zero_reg__` always being 0 could be wrong (e.g. when interrupting right after of a `MUL` instruction).

```
ISR(TIMER1_OVF_vect, ISR_NAKED)
{
    PORTB |= _BV(0); // results in SBI which does not affect SREG
    reti();
}
```

Choosing the vector: Interrupt vector names The interrupt is chosen by supplying one of the symbols in following table.

There are currently two different styles present for naming the vectors. One form uses names starting with `SIG_`, followed by a relatively verbose but arbitrarily chosen name describing the interrupt vector. This has been the only available style in `avr-libc` up to version 1.2.x.

Starting with `avr-libc` version 1.4.0, a second style of interrupt vector names has been added, where a short phrase for the vector description is followed by `_vect`. The short phrase matches the vector name as described in the datasheet of the respective device (and in Atmel's XML files), with spaces replaced by an underscore and other non-alphanumeric characters dropped. Using the suffix `_vect` is intended to improve portability to other C compilers available for the AVR that use a similar naming convention.

The historical naming style might become deprecated in a future release, so it is not recommended for new projects.

Note

The `ISR()` macro cannot really spell-check the argument passed to them. Thus, by misspelling one of the names below in a call to `ISR()`, a function will be created that, while possibly being usable as an interrupt function, is not actually wired into the interrupt vector table. The compiler will generate a warning if it detects a suspiciously looking name of a `ISR()` function (i.e. one that after macro replacement does not start with `__vector_`).

Vector name	Old vector name	Description	Applicable for device
-------------	-----------------	-------------	-----------------------

ADC_vect	SIG_ADC	ADC Conversion Complete		AT90S2333, AT90S4433, AT90S4434, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega8, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATtiny13, ATtiny15, ATtiny26, ATtiny43U, ATtiny48, ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85, ATtiny261, ATtiny461, ATtiny861, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
ANALOG_COMP_0_vect	SIG_COMPARATOR0	Analog Comparator 0	Com-	AT90PWM3, AT90PWM2, AT90PWM1
ANALOG_COMP_1_vect	SIG_COMPARATOR1	Analog Comparator 1	Com-	AT90PWM3, AT90PWM2, AT90PWM1
ANALOG_COMP_2_vect	SIG_COMPARATOR2	Analog Comparator 2	Com-	AT90PWM3, AT90PWM2, AT90PWM1
ANALOG_COMP_vect	SIG_COMPARATOR	Analog Comparator	Com-	AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
ANA_COMP_vect	SIG_COMPARATOR	Analog Comparator	Com-	AT90S1200, AT90S2313, AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, ATmega16, ATmega161, ATmega162, ATmega163, ATmega32, ATmega323, ATmega8, ATmega8515, ATmega8535, ATtiny11, ATtiny12, ATtiny13, ATtiny15, ATtiny2313, ATtiny26, ATtiny28, ATtiny43U, ATtiny48, ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85, ATtiny261, ATtiny461, ATtiny861
CANIT_vect	SIG_CAN_INTERRUPT1	CAN Transfer Complete or Error		AT90CAN128, AT90CAN32, AT90CAN64

EEPROM_- READY_vect	SIG_- EEPROM_- READY, SIG_EE_- READY		ATtiny2313
EE_RDY_vect	SIG_- EEPROM_- READY	EEPROM Ready	AT90S2333, AT90S4433, AT90S4434, AT90S8535, ATmega16, ATmega161, ATmega162, ATmega163, ATmega32, ATmega323, ATmega8, ATmega8515, ATmega8535, ATtiny12, ATtiny13, ATtiny15, ATtiny26, ATtiny43U, ATtiny48, ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85, ATtiny261, ATtiny461, ATtiny861
EE_READY_- vect	SIG_- EEPROM_- READY	EEPROM Ready	AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, ATmega406, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
EXT_INT0_- vect	SIG_- INTERRUPT0	External Interrupt Request 0	ATtiny24, ATtiny44, ATtiny84
INT0_vect	SIG_- INTERRUPT0	External Interrupt 0	AT90S1200, AT90S2313, AT90S2323, AT90S2333, AT90S2343, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, ATmega406, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega8, ATmega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny11, ATtiny12, ATtiny13, ATtiny15, ATtiny22, ATtiny2313, ATtiny26, ATtiny28, ATtiny43U, ATtiny48, ATtiny45, ATtiny25, ATtiny85, ATtiny261, ATtiny461, ATtiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646

INT1_vect	SIG_ INTERRUPT1	External Interrupt Request 1	AT90S2313, AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega168P, ATmega32, ATmega323, ATmega328P, ATmega32HVB, ATmega406, AT- mega48P, ATmega64, ATmega8, AT- mega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, AT- mega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny2313, ATtiny28, ATtiny48, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
INT2_vect	SIG_ INTERRUPT2	External Interrupt Request 2	AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega32, ATmega323, ATmega32HVB, ATmega406, ATmega64, ATmega8515, AT- mega8535, ATmega640, ATmega1280, ATmega1281, ATmega2560, AT- mega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, AT- mega16HVA, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
INT3_vect	SIG_ INTERRUPT3	External Interrupt Request 3	AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega32HVB, ATmega406, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
INT4_vect	SIG_ INTERRUPT4	External Interrupt Request 4	AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, AT- mega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, AT- mega2561, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
INT5_vect	SIG_ INTERRUPT5	External Interrupt Request 5	AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, AT- mega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, AT- mega2561, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646

INT6_vect	SIG_- INTERRUPT6	External Interrupt Request 6	AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, AT- mega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, AT- mega2561, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
INT7_vect	SIG_- INTERRUPT7	External Interrupt Request 7	AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, AT- mega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, AT- mega2561, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
IO_PINS_vect	SIG_PIN, SIG_PIN_- CHANGE	External Interrupt Request 0	ATtiny11, ATtiny12, ATtiny15, ATtiny26
LCD_vect	SIG_LCD	LCD Start of Frame	ATmega169, ATmega169P, ATmega329, ATmega3290, ATmega3290P, ATmega649, ATmega6490
LOWLEVEL_- IO_PINS_vect	SIG_PIN	Low-level Input on Port B	ATtiny28
OVRIT_vect	SIG_CAN_- OVERFLOW1	CAN Timer Overrun	AT90CAN128, AT90CAN32, AT90CAN64
PCINT0_vect	SIG_PIN_- CHANGE0	Pin Change Inter- rupt Request 0	ATmega162, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, AT- mega406, ATmega48P, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, AT- mega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATtiny13, ATtiny43U, ATtiny48, ATtiny24, AT- tiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
PCINT1_vect	SIG_PIN_- CHANGE1	Pin Change Inter- rupt Request 1	ATmega162, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, AT- mega406, ATmega48P, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATtiny43U, ATtiny48, AT- tiny24, ATtiny44, ATtiny84, AT90USB162, AT90USB82
PCINT2_vect	SIG_PIN_- CHANGE2	Pin Change Inter- rupt Request 2	ATmega3250, ATmega3250P, ATmega328P, ATmega3290, ATmega3290P, ATmega48P, ATmega6450, ATmega6490, ATmega88P, ATmega168, ATmega48, ATmega88, AT- mega640, ATmega1280, ATmega1281, AT- mega2560, ATmega2561, ATmega324P, AT- mega164P, ATmega644P, ATmega644, AT- tiny48

PCINT3_vect	SIG_PIN_- CHANGE3	Pin Change Inter- rupt Request 3	ATmega3250, ATmega3250P, ATmega3290, ATmega3290P, ATmega6450, ATmega6490, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATtiny48
PCINT_vect	SIG_PIN_- CHANGE, SIG_PCINT		ATtiny2313, ATtiny261, ATtiny461, AT- tiny861
PSC0_- CAPT_vect	SIG_PSC0_- CAPTURE	PSC0 Capture Event	AT90PWM3, AT90PWM2, AT90PWM1
PSC0_EC_- vect	SIG_PSC0_- END_CYCLE	PSC0 End Cycle	AT90PWM3, AT90PWM2, AT90PWM1
PSC1_- CAPT_vect	SIG_PSC1_- CAPTURE	PSC1 Capture Event	AT90PWM3, AT90PWM2, AT90PWM1
PSC1_EC_- vect	SIG_PSC1_- END_CYCLE	PSC1 End Cycle	AT90PWM3, AT90PWM2, AT90PWM1
PSC2_- CAPT_vect	SIG_PSC2_- CAPTURE	PSC2 Capture Event	AT90PWM3, AT90PWM2, AT90PWM1
PSC2_EC_- vect	SIG_PSC2_- END_CYCLE	PSC2 End Cycle	AT90PWM3, AT90PWM2, AT90PWM1
SPI_STC_vect	SIG_SPI	Serial Transfer Complete	AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega8, ATmega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, AT- mega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny48, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
SPM_RDY_- vect	SIG_SPM_- READY	Store Program Memory Ready	ATmega16, ATmega162, ATmega32, AT- mega323, ATmega8, ATmega8515, AT- mega8535
SPM_- READY_vect	SIG_SPM_- READY	Store Program Memory Read	AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega406, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646

TIM0_- COMPA_vect	SIG_- OUTPUT_- COMPARE0A	Timer/Counter Compare Match A	ATtiny13, ATtiny43U, ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85
TIM0_- COMPB_vect	SIG_- OUTPUT_- COMPARE0B	Timer/Counter Compare Match B	ATtiny13, ATtiny43U, ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85
TIM0_OVF_- vect	SIG_- OVERFLOW0	Timer/Counter0 Overflow	ATtiny13, ATtiny43U, ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85
TIM1_- CAPT_vect	SIG_INPUT_- CAPTURE1	Timer/Counter1 Capture Event	ATtiny24, ATtiny44, ATtiny84
TIM1_- COMPA_vect	SIG_- OUTPUT_- COMPARE1A	Timer/Counter1 Compare Match A	ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85
TIM1_- COMPB_vect	SIG_- OUTPUT_- COMPARE1B	Timer/Counter1 Compare Match B	ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85
TIM1_OVF_- vect	SIG_- OVERFLOW1	Timer/Counter1 Overflow	ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85
TIMER0_- CAPT_vect	SIG_INPUT_- CAPTURE0	ADC Conversion Complete	ATtiny261, ATtiny461, ATtiny861
TIMER0_- COMPA_vect	SIG_- OUTPUT_- COMPARE0A	TimerCounter0 Compare Match A	ATmega168, ATmega48, ATmega88, AT- mega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT- mega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny2313, ATtiny48, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER0_- COMPB_vect	SIG_- OUTPUT_- COMPARE0B, SIG_- OUTPUT_- COMPARE0_- B	Timer Counter 0 Compare Match B	AT90PWM3, AT90PWM2, AT90PWM1, ATmega1284P, ATmega168P, ATmega328P, ATmega32HVB, ATmega48P, AT- mega88P, ATmega168, ATmega48, AT- mega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny2313, ATtiny48, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER0_- COMP_A_- vect	SIG_- OUTPUT_- COMPARE0A, SIG_- OUTPUT_- COMPARE0_- A	Timer/Counter0 Compare Match A	AT90PWM3, AT90PWM2, AT90PWM1
TIMER0_- COMP_vect	SIG_- OUTPUT_- COMPARE0	Timer/Counter0 Compare Match	AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega16, AT- mega161, ATmega162, ATmega165, AT- mega165P, ATmega169, ATmega169P, AT- mega32, ATmega323, ATmega325, AT- mega3250, ATmega3250P, ATmega329, AT- mega3290, ATmega3290P, ATmega64, AT- mega645, ATmega6450, ATmega649, AT- mega6490, ATmega8515, ATmega8535
TIMER0_- OVF0_vect	SIG_- OVERFLOW0	Timer/Counter0 Overflow	AT90S2313, AT90S2323, AT90S2343, AT- tiny22, ATtiny26

TIMER0_ OVF_vect	SIG_ OVERFLOW0	Timer/Counter0 Overflow	AT90S1200, AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, AT- mega48P, ATmega64, ATmega645, AT- mega6450, ATmega649, ATmega6490, ATmega8, ATmega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny11, ATtiny12, ATtiny15, ATtiny2313, ATtiny28, ATtiny48, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER1_ CAPT1_vect	SIG_INPUT_ CAPTURE1	Timer/Counter1 Capture Event	AT90S2313
TIMER1_ CAPT_vect	SIG_INPUT_ CAPTURE1	Timer/Counter Capture Event	AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, AT- mega328P, ATmega329, ATmega3290, ATmega3290P, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega8, ATmega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, AT- mega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATtiny2313, ATtiny48, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER1_ CMPA_vect	SIG_ OUTPUT_ COMPARE1A	Timer/Counter1 Compare Match 1A	ATtiny26
TIMER1_ CMPB_vect	SIG_ OUTPUT_ COMPARE1B	Timer/Counter1 Compare Match 1B	ATtiny26
TIMER1_ COMP1_vect	SIG_ OUTPUT_ COMPARE1A	Timer/Counter1 Compare Match	AT90S2313

TIMER1_- COMPA_vect	SIG_- OUTPUT_- COMPARE1A	Timer/Counter1 Compare Match A	AT90S4414, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, AT- mega48P, ATmega64, ATmega645, AT- mega6450, ATmega649, ATmega6490, ATmega8, ATmega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny2313, ATtiny48, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER1_- COMPB_vect	SIG_- OUTPUT_- COMPARE1B	Timer/Counter1 Compare MatchB	AT90S4414, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, AT- mega48P, ATmega64, ATmega645, AT- mega6450, ATmega649, ATmega6490, ATmega8, ATmega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny2313, ATtiny48, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER1_- COMPC_vect	SIG_- OUTPUT_- COMPARE1C	Timer/Counter1 Compare Match C	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER1_- COMPD_vect	SIG_- OUTPUT_- COMPARE0D	Timer/Counter1 Compare Match D	ATtiny261, ATtiny461, ATtiny861
TIMER1_- COMP_vect	SIG_- OUTPUT_- COMPARE1A	Timer/Counter1 Compare Match	AT90S2333, AT90S4433, ATtiny15

TIMER1_ OVF1_vect	SIG_ OVERFLOW1	Timer/Counter1 Overflow	AT90S2313, ATtiny26
TIMER1_ OVF_vect	SIG_ OVERFLOW1	Timer/Counter1 Overflow	AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, AT90PWM216, AT90PWM2B, AT90PWM316, AT90PWM3B, AT90PWM3, AT90PWM2, AT90PWM1, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega32HVB, ATmega48P, AT- mega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega8, ATmega8515, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, AT- mega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny15, ATtiny2313, ATtiny48, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER2_ COMPA_vect	SIG_ OUTPUT_ COMPARE2A	Timer/Counter2 Compare Match A	ATmega168, ATmega48, ATmega88, AT- mega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT- mega324P, ATmega164P, ATmega644P, AT- mega644, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER2_ COMPB_vect	SIG_ OUTPUT_ COMPARE2B	Timer/Counter2 Compare Match A	ATmega168, ATmega48, ATmega88, AT- mega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT- mega324P, ATmega164P, ATmega644P, AT- mega644, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER2_ COMP_vect	SIG_ OUTPUT_ COMPARE2	Timer/Counter2 Compare Match	AT90S4434, AT90S8535, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega16, ATmega161, AT- mega162, ATmega163, ATmega165, AT- mega165P, ATmega169, ATmega169P, AT- mega32, ATmega323, ATmega325, AT- mega3250, ATmega3250P, ATmega329, AT- mega3290, ATmega3290P, ATmega64, AT- mega645, ATmega6450, ATmega649, AT- mega6490, ATmega8, ATmega8535

TIMER2_- OVF_vect	SIG_- OVERFLOW2	Timer/Counter2 Overflow	AT90S4434, AT90S8535, AT90CAN128, AT90CAN32, AT90CAN64, ATmega103, ATmega128, ATmega1284P, ATmega16, ATmega161, ATmega162, ATmega163, ATmega165, ATmega165P, ATmega168P, ATmega169, ATmega169P, ATmega32, ATmega323, ATmega325, ATmega3250, ATmega3250P, ATmega328P, ATmega329, ATmega3290, ATmega3290P, ATmega48P, ATmega64, ATmega645, ATmega6450, ATmega649, ATmega6490, ATmega8, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER3_- CAPT_vect	SIG_INPUT_- CAPTURE3	Timer/Counter3 Capture Event	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega162, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER3_- COMPA_vect	SIG_- OUTPUT_- COMPARE3A	Timer/Counter3 Compare Match A	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega162, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER3_- COMPB_vect	SIG_- OUTPUT_- COMPARE3B	Timer/Counter3 Compare Match B	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega162, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER3_- COMPC_vect	SIG_- OUTPUT_- COMPARE3C	Timer/Counter3 Compare Match C	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER3_- OVF_vect	SIG_- OVERFLOW3	Timer/Counter3 Overflow	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega162, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TIMER4_- CAPT_vect	SIG_INPUT_- CAPTURE4	Timer/Counter4 Capture Event	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER4_- COMPA_vect	SIG_- OUTPUT_- COMPARE4A	Timer/Counter4 Compare Match A	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER4_- COMPB_vect	SIG_- OUTPUT_- COMPARE4B	Timer/Counter4 Compare Match B	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER4_- COMPC_vect	SIG_- OUTPUT_- COMPARE4C	Timer/Counter4 Compare Match C	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER4_- OVF_vect	SIG_- OVERFLOW4	Timer/Counter4 Overflow	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER5_- CAPT_vect	SIG_INPUT_- CAPTURE5	Timer/Counter5 Capture Event	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561

TIMER5_- COMPA_vect	SIG_- OUTPUT_- COMPARE5A	Timer/Counter5 Compare Match A	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER5_- COMPB_vect	SIG_- OUTPUT_- COMPARE5B	Timer/Counter5 Compare Match B	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER5_- COMPC_vect	SIG_- OUTPUT_- COMPARE5C	Timer/Counter5 Compare Match C	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TIMER5_- OVF_vect	SIG_- OVERFLOW5	Timer/Counter5 Overflow	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
TWI_vect	SIG_2WIRE_- SERIAL	2-wire Serial In- terface	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega16, ATmega163, ATmega168P, ATmega32, AT- mega323, ATmega328P, ATmega32HVB, ATmega406, ATmega48P, ATmega64, ATmega8, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, AT- mega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATtiny48, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
TXDONE_- vect	SIG_- TXDONE	Transmission Done, Bit Timer Flag 2 Interrupt	AT86RF401
TXEMPTY_- vect	SIG_TXBE	Transmit Buffer Empty, Bit Timer Flag 0 Interrupt	AT86RF401
UART0_RX_- vect	SIG_- UART0_- RECV	UART0, Rx Complete	ATmega161
UART0_TX_- vect	SIG_- UART0_- TRANS	UART0, Tx Complete	ATmega161
UART0_- UDRE_vect	SIG_- UART0_- DATA	UART0 Data Register Empty	ATmega161
UART1_RX_- vect	SIG_- UART1_- RECV	UART1, Rx Complete	ATmega161
UART1_TX_- vect	SIG_- UART1_- TRANS	UART1, Tx Complete	ATmega161
UART1_- UDRE_vect	SIG_- UART1_- DATA	UART1 Data Register Empty	ATmega161
UART_RX_- vect	SIG_UART_- RECV	UART, Rx Com- plete	AT90S2313, AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, ATmega103, ATmega163, ATmega8515
UART_TX_- vect	SIG_UART_- TRANS	UART, Tx Com- plete	AT90S2313, AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, ATmega103, ATmega163, ATmega8515
UART_- UDRE_vect	SIG_UART_- DATA	UART Data Reg- ister Empty	AT90S2313, AT90S2333, AT90S4414, AT90S4433, AT90S4434, AT90S8515, AT90S8535, ATmega103, ATmega163, ATmega8515
USART0_- RXC_vect	SIG_- USART0_- RECV	USART0, Rx Complete	ATmega162

USART0_ RX_vect	SIG_ UART0_ - RECV	USART0, Complete	Rx	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega165, ATmega165P, ATmega169, ATmega169P, ATmega325, ATmega329, ATmega64, AT- mega645, ATmega649, ATmega640, AT- mega1280, ATmega1281, ATmega2560, AT- mega2561, ATmega324P, ATmega164P, AT- mega644P, ATmega644
USART0_ TXC_vect	SIG_ UART0_ - TRANS	USART0, Complete	Tx	ATmega162
USART0_ TX_vect	SIG_ UART0_ - TRANS	USART0, Complete	Tx	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega165, ATmega165P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega329, ATmega3290, ATmega3290P, ATmega64, ATmega645, ATmega6450, AT- mega649, ATmega6490, ATmega640, AT- mega1280, ATmega1281, ATmega2560, AT- mega2561, ATmega324P, ATmega164P, AT- mega644P, ATmega644
USART0_ UDRE_vect	SIG_ UART0_ - DATA	USART0 Data Register Empty		AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega162, ATmega165, ATmega165P, ATmega169, ATmega169P, ATmega325, ATmega329, ATmega64, ATmega645, ATmega649, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644
USART1_ RXC_vect	SIG_ USART1_ - RECV	USART1, Complete	Rx	ATmega162
USART1_ RX_vect	SIG_ UART1_ - RECV	USART1, Complete	Rx	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT- mega324P, ATmega164P, ATmega644P, ATmega644, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
USART1_ TXC_vect	SIG_ USART1_ - TRANS	USART1, Complete	Tx	ATmega162
USART1_ TX_vect	SIG_ UART1_ - TRANS	USART1, Complete	Tx	AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, AT- mega324P, ATmega164P, ATmega644P, ATmega644, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
USART1_ UDRE_vect	SIG_ UART1_ - DATA	USART1, Data Register Empty		AT90CAN128, AT90CAN32, AT90CAN64, ATmega128, ATmega1284P, ATmega162, ATmega64, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646
USART2_ RX_vect	SIG_ USART2_ - RECV	USART2, Complete	Rx	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561

USART2_ TX_vect	SIG_ USART2_ TRANS	USART2, Tx Complete	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
USART2_ UDRE_vect	SIG_ USART2_ DATA	USART2 Data register Empty	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
USART3_ RX_vect	SIG_ USART3_ RECV	USART3, Rx Complete	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
USART3_ TX_vect	SIG_ USART3_ TRANS	USART3, Tx Complete	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
USART3_ UDRE_vect	SIG_ USART3_ DATA	USART3 Data register Empty	ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561
USART_ RXC_vect	SIG_ USART_ RECV, SIG_ UART_RECV	USART, Complete Rx	ATmega16, ATmega32, ATmega323, AT- mega8
USART_RX_ vect	SIG_ USART_ RECV, SIG_ UART_RECV	USART, Complete Rx	AT90PWM3, AT90PWM2, AT90PWM1, ATmega168P, ATmega3250, ATmega3250P, ATmega328P, ATmega3290, ATmega3290P, ATmega48P, ATmega6450, ATmega6490, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATtiny2313
USART_ TXC_vect	SIG_ USART_ TRANS, SIG_UART_ TRANS	USART, Tx Complete	ATmega16, ATmega32, ATmega323, AT- mega8
USART_TX_ vect	SIG_ USART_ TRANS, SIG_UART_ TRANS	USART, Tx Complete	AT90PWM3, AT90PWM2, AT90PWM1, ATmega168P, ATmega328P, ATmega48P, ATmega8535, ATmega88P, ATmega168, ATmega48, ATmega88, ATtiny2313
USART_ UDRE_vect	SIG_ USART_ DATA, SIG_ UART_DATA	USART Data Register Empty	AT90PWM3, AT90PWM2, AT90PWM1, ATmega16, ATmega168P, ATmega32, AT- mega323, ATmega3250, ATmega3250P, AT- mega328P, ATmega3290, ATmega3290P, ATmega48P, ATmega6450, ATmega6490, ATmega8, ATmega8535, ATmega88P, AT- mega168, ATmega48, ATmega88, AT- tiny2313
USI_ OVERFLOW_ vect	SIG_USI_ OVERFLOW	USI Overflow	ATmega165, ATmega165P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega329, ATmega3290, ATmega3290P, ATmega645, ATmega6450, ATmega649, ATmega6490, ATtiny2313
USI_OVF_ vect	SIG_USI_ OVERFLOW	USI Overflow	ATtiny26, ATtiny43U, ATtiny24, ATtiny44, ATtiny84, ATtiny45, ATtiny25, ATtiny85, ATtiny261, ATtiny461, ATtiny861
USI_START_ vect	SIG_USI_ START	USI Start Condi- tion	ATmega165, ATmega165P, ATmega169, ATmega169P, ATmega325, ATmega3250, ATmega3250P, ATmega329, ATmega3290, ATmega3290P, ATmega645, ATmega6450, ATmega649, ATmega6490, ATtiny2313, ATtiny43U, ATtiny45, ATtiny25, ATtiny85, ATtiny261, ATtiny461, ATtiny861
USI_STRT_ vect	SIG_USI_ START	USI Start	ATtiny26
USI_STR_ vect	SIG_USI_ START	USI START	ATtiny24, ATtiny44, ATtiny84

WATCHDOG_ vect	SIG_ WATCHDOG_ TIMEOUT	Watchdog Time- out	ATtiny24, ATtiny44, ATtiny84
WDT_ OVERFLOW_ vect	SIG_ WATCHDOG_ TIMEOUT, SIG_WDT_ OVERFLOW	Watchdog Timer Overflow	ATtiny2313
WDT_vect	SIG_WDT, SIG_ WATCHDOG_ TIMEOUT	Watchdog Time- out Interrupt	AT90PWM3, AT90PWM2, AT90PWM1, ATmega1284P, ATmega168P, ATmega328P, ATmega32HVB, ATmega406, ATmega48P, ATmega88P, ATmega168, ATmega48, ATmega88, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega324P, ATmega164P, ATmega644P, ATmega644, ATmega16HVA, ATtiny13, ATtiny43U, ATtiny48, ATtiny45, ATtiny25, ATtiny85, ATtiny261, ATtiny461, AT- tiny861, AT90USB162, AT90USB82, AT90USB1287, AT90USB1286, AT90USB647, AT90USB646

23.17.2 Define Documentation

23.17.2.1 #define BADISR_vect

```
#include <avr/interrupt.h>
```

This is a vector which is aliased to `__vector_default`, the vector executed when an ISR fires with no accompanying ISR handler. This may be used along with the [ISR\(\)](#) macro to create a catch-all for undefined but used ISRs for debugging purposes.

23.17.2.2 #define cli()

Disables all interrupts by clearing the global interrupt mask. This function actually compiles into a single line of assembly, so there is no function call overhead. However, the macro also implies a *memory barrier* which can cause additional loss of optimization.

In order to implement atomic access to multi-byte objects, consider using the macros from [<util/atomic.h>](#), rather than implementing them manually with [cli\(\)](#) and [sei\(\)](#).

23.17.2.3 #define EMPTY_INTERRUPT(vector)

Defines an empty interrupt handler function. This will not generate any prolog or epilog code and will only return from the ISR. Do not define a function body as this will define it for you. Example:

```
EMPTY_INTERRUPT(ADC_vect);
```

23.17.2.4 #define ISR(vector, attributes)

Introduces an interrupt handler function (interrupt service routine) that runs with global interrupts initially disabled by default with no attributes specified.

The attributes are optional and alter the behaviour and resultant generated code of the interrupt routine. Multiple attributes may be used for a single function, with a space separating each attribute.

Valid attributes are `ISR_BLOCK`, `ISR_NOBLOCK`, `ISR_NAKED` and [ISR_ALIASOF\(vect\)](#).

`vector` must be one of the interrupt vector names that are valid for the particular MCU type.

23.17.2.5 #define ISR_ALIAS(vector, target_vector)

Aliases a given vector to another one in the same manner as the `ISR_ALIASOF` attribute for the [ISR\(\)](#) macro. Unlike the `ISR_ALIASOF` attribute macro however, this is compatible for all versions of GCC rather than just GCC version 4.2 onwards.

Note

This macro creates a trampoline function for the aliased macro. This will result in a two cycle penalty for the aliased vector compared to the ISR the vector is aliased to, due to the `JMP/RJMP` opcode used.

Deprecated

For new code, the use of `ISR(..., ISR_ALIASOF(...))` is recommended.

Example:

```
ISR(INT0_vect)
{
    PORTB = 42;
}

ISR_ALIAS(INT1_vect, INT0_vect);
```

23.17.2.6 #define ISR_ALIASOF(target_vector)

The ISR is linked to another ISR, specified by the `vect` parameter. This is compatible with GCC 4.2 and greater only.

Use this attribute in the attributes parameter of the `ISR` macro.

23.17.2.7 #define ISR_BLOCK

Identical to an ISR with no attributes specified. Global interrupts are initially disabled by the AVR hardware when entering the ISR, without the compiler modifying this state.

Use this attribute in the attributes parameter of the ISR macro.

23.17.2.8 #define ISR_NAKED

ISR is created with no prologue or epilogue code. The user code is responsible for preservation of the machine state including the SREG register, as well as placing a [reti\(\)](#) at the end of the interrupt routine.

Use this attribute in the attributes parameter of the ISR macro.

23.17.2.9 #define ISR_NOBLOCK

ISR runs with global interrupts initially enabled. The interrupt enable flag is activated by the compiler as early as possible within the ISR to ensure minimal processing delay for nested interrupts.

This may be used to create nested ISRs, however care should be taken to avoid stack overflows, or to avoid infinitely entering the ISR for those cases where the AVR hardware does not clear the respective interrupt flag before entering the ISR.

Use this attribute in the attributes parameter of the ISR macro.

23.17.2.10 #define reti()

Returns from an interrupt routine, enabling global interrupts. This should be the last command executed before leaving an ISR defined with the `ISR_NAKED` attribute.

This macro actually compiles into a single line of assembly, so there is no function call overhead.

23.17.2.11 #define sei()

Enables interrupts by setting the global interrupt mask. This function actually compiles into a single line of assembly, so there is no function call overhead. However, the macro also implies a *memory barrier* which can cause additional loss of optimization.

In order to implement atomic access to multi-byte objects, consider using the macros from <util/atomic.h>, rather than implementing them manually with [cli\(\)](#) and [sei\(\)](#).

23.17.2.12 #define SIGNAL(vector)

Introduces an interrupt handler function that runs with global interrupts initially disabled.

This is the same as the ISR macro without optional attributes.

Deprecated

Do not use [SIGNAL\(\)](#) in new code. Use [ISR\(\)](#) instead.

23.18 <avr/io.h>: AVR device-specific IO definitions

Defines

- `#define _PROTECTED_WRITE(reg, value)`
- `#define _PROTECTED_WRITE_SPM(reg, value)`

23.18.1 Detailed Description

```
#include <avr/io.h>
```

This header file includes the appropriate IO definitions for the device that has been specified by the `-mmcu=` compiler command-line switch. This is done by diverting to the appropriate file `<avr/ioXXXX.h>` which should never be included directly. Some register names common to all AVR devices are defined directly within `<avr/common.h>`, which is included in `<avr/io.h>`, but most of the details come from the respective include file.

Note that this file always includes the following files:

```
#include <avr/sfr_defs.h>
#include <avr/portpins.h>
#include <avr/common.h>
#include <avr/version.h>
```

See [<avr/sfr_defs.h>: Special function registers](#) for more details about that header file.

Included are definitions of the IO register set and their respective bit values as specified in the Atmel documentation. Note that inconsistencies in naming conventions, so even identical functions sometimes get different names on different devices.

Also included are the specific names useable for interrupt function definitions as documented [here](#).

Finally, the following macros are defined:

- **RAMEND**
The last on-chip RAM address.
- **XRAMEND**
The last possible RAM location that is addressable. This is equal to RAMEND for devices that do not allow for external RAM. For devices that allow external RAM, this will be larger than RAMEND.
- **E2END**
The last EEPROM address.
- **FLASHEND**
The last byte address in the Flash program space.
- **SPM_PAGESIZE**
For devices with bootloader support, the flash pagesize (in bytes) to be used for the SPM instruction.

- **E2PAGESIZE**

The size of the EEPROM page.

23.18.2 Define Documentation

23.18.2.1 #define _PROTECTED_WRITE(reg, value)

Write value `value` to IO register `reg` that is protected through the Xmega configuration change protection (CCP) mechanism. This implements the timed sequence that is required for CCP.

Example to modify the CPU clock:

```
#include <avr/io.h>

_PROTECTED_WRITE(CLK_PSCTRL, CLK_PSADIV0_bm);
_PROTECTED_WRITE(CLK_CTRL, CLK_SCLKSEL0_bm);
```

23.18.2.2 #define _PROTECTED_WRITE_SPM(reg, value)

Write value `value` to register `reg` that is protected through the Xmega configuration change protection (CCP) key for self programming (SPM). This implements the timed sequence that is required for CCP.

Example to modify the CPU clock:

```
#include <avr/io.h>

_PROTECTED_WRITE_SPM(NVMCTRL_CTRLA, NVMCTRL_CMD_PAGEERASEWRITE_gc);
```

23.19 <avr/lock.h>: Lockbit Support

Introduction

The Lockbit API allows a user to specify the lockbit settings for the specific AVR device they are compiling for. These lockbit settings will be placed in a special section in the ELF output file, after linking.

Programming tools can take advantage of the lockbit information embedded in the ELF file, by extracting this information and determining if the lockbits need to be programmed after programming the Flash and EEPROM memories. This also allows a single ELF file to contain all the information needed to program an AVR.

To use the Lockbit API, include the <avr/io.h> header file, which in turn automatically includes the individual I/O header file and the <avr/lock.h> file. These other two files provides everything necessary to set the AVR lockbits.

Lockbit API

Each I/O header file may define up to 3 macros that controls what kinds of lockbits are available to the user.

If `__LOCK_BITS_EXIST` is defined, then two lock bits are available to the user and 3 mode settings are defined for these two bits.

If `__BOOT_LOCK_BITS_0_EXIST` is defined, then the two BLB0 lock bits are available to the user and 4 mode settings are defined for these two bits.

If `__BOOT_LOCK_BITS_1_EXIST` is defined, then the two BLB1 lock bits are available to the user and 4 mode settings are defined for these two bits.

If `__BOOT_LOCK_APPLICATION_TABLE_BITS_EXIST` is defined then two lock bits are available to set the locking mode for the Application Table Section (which is used in the XMEGA family).

If `__BOOT_LOCK_APPLICATION_BITS_EXIST` is defined then two lock bits are available to set the locking mode for the Application Section (which is used in the XMEGA family).

If `__BOOT_LOCK_BOOT_BITS_EXIST` is defined then two lock bits are available to set the locking mode for the Boot Loader Section (which is used in the XMEGA family).

The AVR lockbit modes have inverted values, logical 1 for an unprogrammed (disabled) bit and logical 0 for a programmed (enabled) bit. The defined macros for each individual lock bit represent this in their definition by a bit-wise inversion of a mask. For example, the `LB_MODE_3` macro is defined as:

```
#define LB_MODE_3    (0xFC)
,
```

To combine the lockbit mode macros together to represent a whole byte, use the bitwise AND operator, like so:

```
(LB_MODE_3 & BLB0_MODE_2)
```

<avr/lock.h> also defines a macro that provides a default lockbit value: `LOCKBITS_DEFAULT` which is defined to be `0xFF`.

See the AVR device specific datasheet for more details about these lock bits and the available mode settings.

A convenience macro, `LOCKMEM`, is defined as a GCC attribute for a custom-named section of ".lock".

A convenience macro, `LOCKBITS`, is defined that declares a variable, `__lock`, of type unsigned char with the attribute defined by `LOCKMEM`. This variable allows the end user to easily set the lockbit data.

Note

If a device-specific I/O header file has previously defined `LOCKMEM`, then `LOCKMEM` is not redefined. If a device-specific I/O header file has previously defined `LOCKBITS`, then `LOCKBITS` is not redefined. `LOCKBITS` is currently known to be defined in the I/O header files for the XMEGA devices.

API Usage Example

Putting all of this together is easy:

```
#include <avr/io.h>

LOCKBITS = (LB_MODE_1 & BLB0_MODE_3 & BLB1_MODE_4);

int main(void)
{
    return 0;
}
```

Or:

```
#include <avr/io.h>

unsigned char __lock __attribute__((section (".lock"))) =
    (LB_MODE_1 & BLB0_MODE_3 & BLB1_MODE_4);

int main(void)
{
    return 0;
}
```

However there are a number of caveats that you need to be aware of to use this API properly.

Be sure to include <avr/io.h> to get all of the definitions for the API. The LOCKBITS macro defines a global variable to store the lockbit data. This variable is assigned to its own linker section. Assign the desired lockbit values immediately in the variable initialization.

The .lock section in the ELF file will get its values from the initial variable assignment ONLY. This means that you can NOT assign values to this variable in functions and the new values will not be put into the ELF .lock section.

The global variable is declared in the LOCKBITS macro has two leading underscores, which means that it is reserved for the "implementation", meaning the library, so it will not conflict with a user-named variable.

You must initialize the lockbit variable to some meaningful value, even if it is the default value. This is because the lockbits default to a logical 1, meaning unprogrammed. Normal uninitialized data defaults to all logical zeros. So it is vital that all lockbits are initialized, even with default data. If they are not, then the lockbits may not be programmed to the desired settings and can possibly put your device into an unrecoverable state.

Be sure to have the -mmcu=*device* flag in your compile command line and your linker command line to have the correct device selected and to have the correct I/O header file included when you include <avr/io.h>.

You can print out the contents of the .lock section in the ELF file by using this command line:

```
avr-objdump -s -j .lock <ELF file>
```

23.20 <avr/pgmspace.h>: Program Space Utilities

Defines

- #define `PROGMEM` `__ATTR_PROGMEM__`
- #define `PGM_P` `const char *`
- #define `PGM_VOID_P` `const void *`
- #define `PSTR(s)` `((const PROGMEM char *) (s))`
- #define `pgm_read_byte_near(address_short)` `__LPM__((uint16_t)(address_short))`
- #define `pgm_read_word_near(address_short)` `__LPM_word__((uint16_t)(address_short))`
- #define `pgm_read_dword_near(address_short)` `__LPM_dword__((uint16_t)(address_short))`
- #define `pgm_read_float_near(address_short)` `__LPM_float__((uint16_t)(address_short))`
- #define `pgm_read_ptr_near(address_short)` `(void*)__LPM_word__((uint16_t)(address_short))`
- #define `pgm_read_byte_far(address_long)` `__ELPM__((uint32_t)(address_long))`
- #define `pgm_read_word_far(address_long)` `__ELPM_word__((uint32_t)(address_long))`
- #define `pgm_read_dword_far(address_long)` `__ELPM_dword__((uint32_t)(address_long))`
- #define `pgm_read_float_far(address_long)` `__ELPM_float__((uint32_t)(address_long))`
- #define `pgm_read_ptr_far(address_long)` `(void*)__ELPM_word__((uint32_t)(address_long))`
- #define `pgm_read_byte(address_short)` `pgm_read_byte_near(address_short)`
- #define `pgm_read_word(address_short)` `pgm_read_word_near(address_short)`
- #define `pgm_read_dword(address_short)` `pgm_read_dword_near(address_short)`
- #define `pgm_read_float(address_short)` `pgm_read_float_near(address_short)`
- #define `pgm_read_ptr(address_short)` `pgm_read_ptr_near(address_short)`
- #define `pgm_get_far_address(var)`

Typedefs

- typedef void PROGMEM `prog_void`
- typedef char PROGMEM `prog_char`
- typedef unsigned char PROGMEM `prog_uchar`
- typedef int8_t PROGMEM `prog_int8_t`
- typedef uint8_t PROGMEM `prog_uint8_t`
- typedef int16_t PROGMEM `prog_int16_t`
- typedef uint16_t PROGMEM `prog_uint16_t`
- typedef int32_t PROGMEM `prog_int32_t`
- typedef uint32_t PROGMEM `prog_uint32_t`
- typedef int64_t PROGMEM `prog_int64_t`
- typedef uint64_t PROGMEM `prog_uint64_t`

Functions

- const void * [memchr_P](#) (const void *, int __val, size_t __len)
- int [memcmp_P](#) (const void *, const void *, size_t) __ATTR_PURE__
- void * [memcpy_P](#) (void *, const void *, int __val, size_t)
- void * [memcpy_P](#) (void *, const void *, size_t)
- void * [memmem_P](#) (const void *, size_t, const void *, size_t) __ATTR_PURE_-
-
- const void * [memrchr_P](#) (const void *, int __val, size_t __len)
- char * [strcat_P](#) (char *, const char *)
- const char * [strchr_P](#) (const char *, int __val)
- const char * [strchrnul_P](#) (const char *, int __val)
- int [strcmp_P](#) (const char *, const char *) __ATTR_PURE__
- char * [strcpy_P](#) (char *, const char *)
- int [strcasecmp_P](#) (const char *, const char *) __ATTR_PURE__
- char * [strcasestr_P](#) (const char *, const char *) __ATTR_PURE__
- size_t [strcspn_P](#) (const char *__s, const char *__reject) __ATTR_PURE__
- size_t [strlcat_P](#) (char *, const char *, size_t)
- size_t [strncpy_P](#) (char *, const char *, size_t)
- size_t [strlen_P](#) (const char *, size_t)
- int [strncmp_P](#) (const char *, const char *, size_t) __ATTR_PURE__
- int [strncasecmp_P](#) (const char *, const char *, size_t) __ATTR_PURE__
- char * [strncat_P](#) (char *, const char *, size_t)
- char * [strncpy_P](#) (char *, const char *, size_t)
- char * [strpbrk_P](#) (const char *__s, const char *__accept) __ATTR_PURE__
- const char * [strrchr_P](#) (const char *, int __val)
- char * [strsep_P](#) (char **__sp, const char *__delim)
- size_t [strspn_P](#) (const char *__s, const char *__accept) __ATTR_PURE__
- char * [strstr_P](#) (const char *, const char *) __ATTR_PURE__
- char * [strtok_P](#) (char *__s, const char *__delim)
- char * [strtok_rP](#) (char *__s, const char *__delim, char **__last)
- size_t [strlen_PF](#) (uint_farptr_t src)
- size_t [strlen_PF](#) (uint_farptr_t src, size_t len)
- void * [memcpy_PF](#) (void *dest, uint_farptr_t src, size_t len)
- char * [strcpy_PF](#) (char *dest, uint_farptr_t src)
- char * [strncpy_PF](#) (char *dest, uint_farptr_t src, size_t len)
- char * [strcat_PF](#) (char *dest, uint_farptr_t src)
- size_t [strlcat_PF](#) (char *dst, uint_farptr_t src, size_t siz)
- char * [strncat_PF](#) (char *dest, uint_farptr_t src, size_t len)
- int [strcmp_PF](#) (const char *s1, uint_farptr_t s2) __ATTR_PURE__
- int [strncmp_PF](#) (const char *s1, uint_farptr_t s2, size_t n) __ATTR_PURE__
- int [strcasecmp_PF](#) (const char *s1, uint_farptr_t s2) __ATTR_PURE__
- int [strncasecmp_PF](#) (const char *s1, uint_farptr_t s2, size_t n) __ATTR_PURE_-
-
- char * [strstr_PF](#) (const char *s1, uint_farptr_t s2)
- size_t [strncpy_PF](#) (char *dst, uint_farptr_t src, size_t siz)
- int [memcmp_PF](#) (const void *, uint_farptr_t, size_t) __ATTR_PURE__
- static size_t [strlen_P](#) (const char *s)

23.20.1 Detailed Description

```
#include <avr/io.h>
#include <avr/pgmspace.h>
```

The functions in this module provide interfaces for a program to access data stored in program space (flash memory) of the device. In order to use these functions, the target device must support either the LPM or ELPM instructions.

Note

These functions are an attempt to provide some compatibility with header files that come with IAR C, to make porting applications between different compilers easier. This is not 100% compatibility though (GCC does not have full support for multiple address spaces yet).

If you are working with strings which are completely based in ram, use the standard string functions described in <string.h>: [Strings](#).

If possible, put your constant tables in the lower 64 KB and use [pgm_read_byte_near\(\)](#) or [pgm_read_word_near\(\)](#) instead of [pgm_read_byte_far\(\)](#) or [pgm_read_word_far\(\)](#) since it is more efficient that way, and you can still use the upper 64K for executable code. All functions that are suffixed with a `_P` *require* their arguments to be in the lower 64 KB of the flash ROM, as they do not use ELPM instructions. This is normally not a big concern as the linker setup arranges any program space constants declared using the macros from this header file so they are placed right after the interrupt vectors, and in front of any executable code. However, it can become a problem if there are too many of these constants, or for bootloaders on devices with more than 64 KB of ROM. *All these functions will not work in that situation.*

For **Xmega** devices, make sure the NVM controller command register (NVM.CMD or NVM_CMD) is set to 0x00 (NOP) before using any of these functions.

23.20.2 Define Documentation

23.20.2.1 #define pgm_get_far_address(var)

Value:

```
((
    uint_farp_ptr_t tmp;
    __asm__ __volatile__(
        "ldi    %A0, lo8(%1)"      "\n\t"
        "ldi    %B0, hi8(%1)"      "\n\t"
        "ldi    %C0, hh8(%1)"      "\n\t"
        "clr    %D0"               "\n\t"
        :
        "=>" (tmp)
        :
        "p"   (&(var))
    );
    tmp;
))
```

This macro facilitates the obtention of a 32 bit "far" pointer (only 24 bits used) to data even passed the 64KB limit for the 16 bit ordinary pointer. It is similar to the '&' operator, with some limitations.

Comments:

- The overhead is minimal and it's mainly due to the 32 bit size operation.
- 24 bit sizes guarantees the code compatibility for use in future devices.
- hh8() is an undocumented feature but seems to give the third significant byte of a 32 bit data and accepts symbols, complementing the functionality of hi8() and lo8(). There is not an equivalent assembler function to get the high significant byte.
- 'var' has to be resolved at linking time as an existing symbol, i.e, a simple type variable name, an array name (not an indexed element of the array, if the index is a constant the compiler does not complain but fails to get the address if optimization is enabled), a struct name or a struct field name, a function identifier, a linker defined identifier,...
- The returned value is the identifier's VMA (virtual memory address) determined by the linker and falls in the corresponding memory region. The AVR Harvard architecture requires non overlapping VMA areas for the multiple address spaces in the processor: Flash ROM, RAM, and EEPROM. Typical offset for this are 0x00000000, 0x00800xx0, and 0x00810000 respectively, derived from the linker script used and linker options. The value returned can be seen then as a universal pointer.

23.20.2.2 #define PGM_P const char *

Used to declare a variable that is a pointer to a string in program space.

23.20.2.3 #define pgm_read_byte(address_short) pgm_read_byte_near(address_short)

Read a byte from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.4 #define pgm_read_byte_far(address_long) __ELPM__((uint32_t)(address_long))

Read a byte from the program space with a 32-bit (far) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.5 `#define pgm_read_byte_near(address_short) __LPM__((uint16_t)(address_short))`

Read a byte from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.6 `#define pgm_read_dword(address_short) pgm_read_dword_near(address_short)`

Read a double word from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.7 `#define pgm_read_dword_far(address_long) __ELPM_dword__((uint32_t)(address_long))`

Read a double word from the program space with a 32-bit (far) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.8 `#define pgm_read_dword_near(address_short) __LPM_dword__((uint16_t)(address_short))`

Read a double word from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.9 `#define pgm_read_float(address_short) pgm_read_float_near(address_short)`

Read a float from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.10 `#define pgm_read_float_far(address_long) __ELPM_float((uint32_t)(address_long))`

Read a float from the program space with a 32-bit (far) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.11 `#define pgm_read_float_near(address_short) __LPM_float((uint16_t)(address_short))`

Read a float from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.12 `#define pgm_read_ptr(address_short) pgm_read_ptr_near(address_short)`

Read a pointer from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.13 `#define pgm_read_ptr_far(address_long) (void*)__ELPM_word((uint32_t)(address_long))`

Read a pointer from the program space with a 32-bit (far) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.14 `#define pgm_read_ptr_near(address_short) (void*)__LPM_word((uint16_t)(address_short))`

Read a pointer from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.15 `#define pgm_read_word(address_short) pgm_read_word_near(address_short)`

Read a word from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.16 `#define pgm_read_word_far(address_long) __ELPM_word((uint32_t)(address_long))`

Read a word from the program space with a 32-bit (far) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.17 `#define pgm_read_word_near(address_short) __LPM_word((uint16_t)(address_short))`

Read a word from the program space with a 16-bit (near) address.

Note

The address is a byte address. The address is in the program space.

23.20.2.18 `#define PGM_VOID_P const void *`

Used to declare a generic pointer to an object in program space.

23.20.2.19 `#define PROGMEM __ATTR_PROGMEM__`

Attribute to use in order to declare an object being located in flash ROM.

23.20.2.20 `#define PSTR(s) ((const PROGMEM char *)(s))`

Used to declare a static pointer to a string in program space.

23.20.3 Typedef Documentation

23.20.3.1 `prog_char`

Note

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of a "char" object located in flash ROM.

23.20.3.2 prog_int16_t**Note**

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "int16_t" object located in flash ROM.

23.20.3.3 prog_int32_t**Note**

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "int32_t" object located in flash ROM.

23.20.3.4 prog_int64_t

Note

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "int64_t" object located in flash ROM.

Note

This type is not available when the compiler option `-mint8` is in effect.

23.20.3.5 prog_int8_t**Note**

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "int8_t" object located in flash ROM.

23.20.3.6 prog_uchar**Note**

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "unsigned char" object located in flash ROM.

23.20.3.7 prog_uint16_t

Note

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "uint16_t" object located in flash ROM.

23.20.3.8 prog_uint32_t

Note

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "uint32_t" object located in flash ROM.

23.20.3.9 prog_uint64_t

Note

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including [<avr/pgmspace.h>](#) (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "uint64_t" object located in flash ROM.

Note

This type is not available when the compiler option `-mint8` is in effect.

23.20.3.10 prog_uint8_t**Note**

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including `<avr/pgmspace.h>` (either by a `#define` directive, or by a `-D` compiler option.)

Type of an "uint8_t" object located in flash ROM.

23.20.3.11 prog_void**Note**

DEPRECATED

This typedef is now deprecated because the usage of the `__progmem__` attribute on a type is not supported in GCC. However, the use of the `__progmem__` attribute on a variable declaration is supported, and this is now the recommended usage.

The typedef is only visible if the macro `__PROG_TYPES_COMPAT__` has been defined before including `<avr/pgmspace.h>` (either by a `#define` directive, or by a `-D` compiler option.)

Type of a "void" object located in flash ROM. Does not make much sense by itself, but can be used to declare a "void *" object in flash ROM.

23.20.4 Function Documentation**23.20.4.1 void * memcpy_P (void * dest, const void * src, int val, size_t len)**

This function is similar to `memcpy()` except that `src` is pointer to a string in program space.

23.20.4.2 const void * memchr_P (const void * s, int val, size_t len)

Scan flash memory for a character.

The `memchr_P()` function scans the first `len` bytes of the flash memory area pointed to by `s` for the character `val`. The first byte to match `val` (interpreted as an unsigned character) stops the operation.

Returns

The `memchr_P()` function returns a pointer to the matching byte or `NULL` if the character does not occur in the given memory area.

23.20.4.3 `int memcmp_P (const void * s1, const void * s2, size_t len)`

Compare memory areas.

The `memcmp_P()` function compares the first `len` bytes of the memory areas `s1` and flash `s2`. The comparison is performed using unsigned char operations.

Returns

The `memcmp_P()` function returns an integer less than, equal to, or greater than zero if the first `len` bytes of `s1` is found, respectively, to be less than, to match, or be greater than the first `len` bytes of `s2`.

23.20.4.4 `int memcmp_PF (const void * s1, uint_farptr_t s2, size_t len)`

Compare memory areas.

The `memcmp_PF()` function compares the first `len` bytes of the memory areas `s1` and flash `s2`. The comparison is performed using unsigned char operations. It is an equivalent of `memcmp_P()` function, except that it is capable working on all FLASH including the extended area above 64kB.

Returns

The `memcmp_PF()` function returns an integer less than, equal to, or greater than zero if the first `len` bytes of `s1` is found, respectively, to be less than, to match, or be greater than the first `len` bytes of `s2`.

23.20.4.5 `void * memcpy_P (void * dest, const void * src, size_t n)`

The `memcpy_P()` function is similar to `memcpy()`, except the `src` string resides in program space.

Returns

The `memcpy_P()` function returns a pointer to `dest`.

23.20.4.6 void * memcpy_P (void * *dest*, uint_farptr_t *src*, size_t *n*)

Copy a memory block from flash to SRAM.

The [memcpy_P\(\)](#) function is similar to [memcpy\(\)](#), except the data is copied from the program space and is addressed using a far pointer.

Parameters

- dest* A pointer to the destination buffer
- src* A far pointer to the origin of data in flash memory
- n* The number of bytes to be copied

Returns

The [memcpy_P\(\)](#) function returns a pointer to *dst*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.7 void * memmem_P (const void * *s1*, size_t *len1*, const void * *s2*, size_t *len2*)

The [memmem_P\(\)](#) function is similar to [memmem\(\)](#) except that *s2* is pointer to a string in program space.

23.20.4.8 const void * memrchr_P (const void * *src*, int *val*, size_t *len*)

The [memrchr_P\(\)](#) function is like the [memchr_P\(\)](#) function, except that it searches backwards from the end of the *len* bytes pointed to by *src* instead of forwards from the front. (Glibc, GNU extension.)

Returns

The [memrchr_P\(\)](#) function returns a pointer to the matching byte or NULL if the character does not occur in the given memory area.

23.20.4.9 int strcasecmp_P (const char * *s1*, const char * *s2*)

Compare two strings ignoring case.

The [strcasecmp_P\(\)](#) function compares the two strings *s1* and *s2*, ignoring the case of the characters.

Parameters

- s1* A pointer to a string in the devices SRAM.

s2 A pointer to a string in the devices Flash.

Returns

The [strcasecmp_P\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* is found, respectively, to be less than, to match, or be greater than *s2*. A consequence of the ordering used by [strcasecmp_P\(\)](#) is that if *s1* is an initial substring of *s2*, then *s1* is considered to be "less than" *s2*.

23.20.4.10 int strcasecmp_PF (const char * *s1*, uint_farptr_t *s2*)

Compare two strings ignoring case.

The [strcasecmp_PF\(\)](#) function compares the two strings *s1* and *s2*, ignoring the case of the characters.

Parameters

s1 A pointer to the first string in SRAM

s2 A far pointer to the second string in Flash

Returns

The [strcasecmp_PF\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* is found, respectively, to be less than, to match, or be greater than *s2*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.11 char * strcasestr_P (const char * *s1*, const char * *s2*)

This function is similar to [strcasestr\(\)](#) except that *s2* is pointer to a string in program space.

23.20.4.12 char * strcat_P (char * *dest*, const char * *src*)

The [strcat_P\(\)](#) function is similar to [strcat\(\)](#) except that the *src* string must be located in program space (flash).

Returns

The [strcat\(\)](#) function returns a pointer to the resulting string *dest*.

23.20.4.13 char * strcat_PF (char * *dst*, uint_farptr_t *src*)

Concatenates two strings.

The `strcat_PF()` function is similar to `strcat()` except that the *src* string must be located in program space (flash) and is addressed using a far pointer

Parameters

dst A pointer to the destination string in SRAM

src A far pointer to the string to be appended in Flash

Returns

The `strcat_PF()` function returns a pointer to the resulting string *dst*. The contents of RAMPZ SFR are undefined when the function returns

23.20.4.14 `const char * strchr_P (const char * s, int val)`

Locate character in program space string.

The `strchr_P()` function locates the first occurrence of *val* (converted to a char) in the string pointed to by *s* in program space. The terminating null character is considered to be part of the string.

The `strchr_P()` function is similar to `strchr()` except that *s* is pointer to a string in program space.

Returns

The `strchr_P()` function returns a pointer to the matched character or `NULL` if the character is not found.

23.20.4.15 `const char * strchrnul_P (const char * s, int c)`

The `strchrnul_P()` function is like `strchr_P()` except that if *c* is not found in *s*, then it returns a pointer to the null byte at the end of *s*, rather than `NULL`. (Glibc, GNU extension.)

Returns

The `strchrnul_P()` function returns a pointer to the matched character, or a pointer to the null byte at the end of *s* (i.e., `s+strlen(s)`) if the character is not found.

23.20.4.16 `int strcmp_P (const char * s1, const char * s2)`

The `strcmp_P()` function is similar to `strcmp()` except that *s2* is pointer to a string in program space.

Returns

The [strcmp_P\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* is found, respectively, to be less than, to match, or be greater than *s2*. A consequence of the ordering used by [strcmp_P\(\)](#) is that if *s1* is an initial substring of *s2*, then *s1* is considered to be "less than" *s2*.

23.20.4.17 int strcmp_PF (const char * *s1*, uint_farptr_t *s2*)

Compares two strings.

The [strcmp_PF\(\)](#) function is similar to [strcmp\(\)](#) except that *s2* is a far pointer to a string in program space.

Parameters

- s1* A pointer to the first string in SRAM
- s2* A far pointer to the second string in Flash

Returns

The [strcmp_PF\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* is found, respectively, to be less than, to match, or be greater than *s2*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.18 char * strcpy_P (char * *dest*, const char * *src*)

The [strcpy_P\(\)](#) function is similar to [strcpy\(\)](#) except that *src* is a pointer to a string in program space.

Returns

The [strcpy_P\(\)](#) function returns a pointer to the destination string *dest*.

23.20.4.19 char * strcpy_PF (char * *dst*, uint_farptr_t *src*)

Duplicate a string.

The [strcpy_PF\(\)](#) function is similar to [strcpy\(\)](#) except that *src* is a far pointer to a string in program space.

Parameters

- dst* A pointer to the destination string in SRAM
- src* A far pointer to the source string in Flash

Returns

The [strcpy_P\(\)](#) function returns a pointer to the destination string *dst*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.20 size_t strcspn_P (const char * s, const char * reject)

The [strcspn_P\(\)](#) function calculates the length of the initial segment of *s* which consists entirely of characters not in *reject*. This function is similar to [strcspn\(\)](#) except that *reject* is a pointer to a string in program space.

Returns

The [strcspn_P\(\)](#) function returns the number of characters in the initial segment of *s* which are not in the string *reject*. The terminating zero is not considered as a part of string.

23.20.4.21 size_t strlcat_P (char * dst, const char * src, size_t siz)

Concatenate two strings.

The [strlcat_P\(\)](#) function is similar to [strlcat\(\)](#), except that the *src* string must be located in program space (flash).

Appends *src* to string *dst* of size *siz* (unlike [strncat\(\)](#), *siz* is the full size of *dst*, not space left). At most *siz*-1 characters will be copied. Always NULL terminates (unless *siz* <= *strlen(dst)*).

Returns

The [strlcat_P\(\)](#) function returns *strlen(src)* + MIN(*siz*, *strlen(initial dst)*). If *retval* >= *siz*, truncation occurred.

23.20.4.22 size_t strlcat_P (char * dst, uint_farptr_t src, size_t n)

Concatenate two strings.

The [strlcat_P\(\)](#) function is similar to [strlcat\(\)](#), except that the *src* string must be located in program space (flash) and is addressed using a far pointer.

Appends *src* to string *dst* of size *n* (unlike [strncat\(\)](#), *n* is the full size of *dst*, not space left). At most *n*-1 characters will be copied. Always NULL terminates (unless *n* <= *strlen(dst)*).

Parameters

dst A pointer to the destination string in SRAM

src A far pointer to the source string in Flash

n The total number of bytes allocated to the destination string

Returns

The [strlcat_PF\(\)](#) function returns `strlen(src) + MIN(n, strlen(initial dst))`. If `retval >= n`, truncation occurred. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.23 `size_t strlcpy_P(char *dst, const char *src, size_t siz)`

Copy a string from progmem to RAM.

Copy `src` to string `dst` of size `siz`. At most `siz-1` characters will be copied. Always NULL terminates (unless `siz == 0`). The [strlcpy_P\(\)](#) function is similar to [strlcpy\(\)](#) except that the `src` is pointer to a string in memory space.

Returns

The [strlcpy_P\(\)](#) function returns `strlen(src)`. If `retval >= siz`, truncation occurred.

23.20.4.24 `size_t strlcpy_PF(char *dst, uint_farptr_t src, size_t siz)`

Copy a string from progmem to RAM.

Copy `src` to string `dst` of size `siz`. At most `siz-1` characters will be copied. Always NULL terminates (unless `siz == 0`).

Returns

The [strlcpy_PF\(\)](#) function returns `strlen(src)`. If `retval >= siz`, truncation occurred. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.25 `size_t strlen_P(const char *src) [static]`

The [strlen_P\(\)](#) function is similar to [strlen\(\)](#), except that `src` is a pointer to a string in program space.

Returns

The [strlen_P\(\)](#) function returns the number of characters in `src`.

Note

[strlen_P\(\)](#) is implemented as an inline function in the [avr/pgmspace.h](#) header file, which will check if the length of the string is a constant and known at compile

time. If it is not known at compile time, the macro will issue a call to `__strlen_P()` which will then calculate the length of the string as normal.

23.20.4.26 `size_t strlen_PF (uint_farptr_t s)`

Obtain the length of a string.

The `strlen_PF()` function is similar to `strlen()`, except that `s` is a far pointer to a string in program space.

Parameters

`s` A far pointer to the string in flash

Returns

The `strlen_PF()` function returns the number of characters in `s`. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.27 `int strncasecmp_P (const char * s1, const char * s2, size_t n)`

Compare two strings ignoring case.

The `strncasecmp_P()` function is similar to `strncasecmp_P()`, except it only compares the first `n` characters of `s1`.

Parameters

`s1` A pointer to a string in the devices SRAM.

`s2` A pointer to a string in the devices Flash.

`n` The maximum number of bytes to compare.

Returns

The `strncasecmp_P()` function returns an integer less than, equal to, or greater than zero if `s1` (or the first `n` bytes thereof) is found, respectively, to be less than, to match, or be greater than `s2`. A consequence of the ordering used by `strncasecmp_P()` is that if `s1` is an initial substring of `s2`, then `s1` is considered to be "less than" `s2`.

23.20.4.28 `int strncasecmp_PF (const char * s1, uint_farptr_t s2, size_t n)`

Compare two strings ignoring case.

The `strncasecmp_PF()` function is similar to `strncasecmp_PF()`, except it only compares the first `n` characters of `s1` and the string in flash is addressed using a far pointer.

Parameters

- s1* A pointer to a string in SRAM
- s2* A far pointer to a string in Flash
- n* The maximum number of bytes to compare

Returns

The [strncasecmp_PF\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* (or the first *n* bytes thereof) is found, respectively, to be less than, to match, or be greater than *s2*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.29 char * strncat_P (char * *dest*, const char * *src*, size_t *len*)

Concatenate two strings.

The [strncat_P\(\)](#) function is similar to [strncat\(\)](#), except that the *src* string must be located in program space (flash).

Returns

The [strncat_P\(\)](#) function returns a pointer to the resulting string *dest*.

23.20.4.30 char * strncat_PF (char * *dst*, uint_farptr_t *src*, size_t *n*)

Concatenate two strings.

The [strncat_PF\(\)](#) function is similar to [strncat\(\)](#), except that the *src* string must be located in program space (flash) and is addressed using a far pointer.

Parameters

- dst* A pointer to the destination string in SRAM
- src* A far pointer to the source string in Flash
- n* The maximum number of bytes to append

Returns

The [strncat_PF\(\)](#) function returns a pointer to the resulting string *dst*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.31 int strncmp_P (const char * *s1*, const char * *s2*, size_t *n*)

The [strncmp_P\(\)](#) function is similar to [strcmp_P\(\)](#) except it only compares the first (at most) *n* characters of *s1* and *s2*.

Returns

The [strncmp_P\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* (or the first *n* bytes thereof) is found, respectively, to be less than, to match, or be greater than *s2*.

23.20.4.32 int strncmp_PF (const char * *s1*, uint_farptr_t *s2*, size_t *n*)

Compare two strings with limited length.

The [strncmp_PF\(\)](#) function is similar to [strcmp_PF\(\)](#) except it only compares the first (at most) *n* characters of *s1* and *s2*.

Parameters

- s1* A pointer to the first string in SRAM
- s2* A far pointer to the second string in Flash
- n* The maximum number of bytes to compare

Returns

The [strncmp_PF\(\)](#) function returns an integer less than, equal to, or greater than zero if *s1* (or the first *n* bytes thereof) is found, respectively, to be less than, to match, or be greater than *s2*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.33 char * strncpy_P (char * *dest*, const char * *src*, size_t *n*)

The [strncpy_P\(\)](#) function is similar to [strncpy\(\)](#) except that not more than *n* bytes of *src* are copied. Thus, if there is no null byte among the first *n* bytes of *src*, the result will not be null-terminated.

In the case where the length of *src* is less than that of *n*, the remainder of *dest* will be padded with nulls.

Returns

The [strncpy_P\(\)](#) function returns a pointer to the destination string *dest*.

23.20.4.34 `char * strncpy_PF (char * dst, uint_farptr_t src, size_t n)`

Duplicate a string until a limited length.

The `strncpy_PF()` function is similar to `strncpy()` except that not more than *n* bytes of *src* are copied. Thus, if there is no null byte among the first *n* bytes of *src*, the result will not be null-terminated.

In the case where the length of *src* is less than that of *n*, the remainder of *dst* will be padded with nulls.

Parameters

dst A pointer to the destination string in SRAM

src A far pointer to the source string in Flash

n The maximum number of bytes to copy

Returns

The `strncpy_PF()` function returns a pointer to the destination string *dst*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.35 `size_t strlen_P (const char * src, size_t len)`

Determine the length of a fixed-size string.

The `strlen_P()` function is similar to `strlen()`, except that *src* is a pointer to a string in program space.

Returns

The `strlen_P` function returns `strlen_P(src)`, if that is less than *len*, or *len* if there is no `'\0'` character among the first *len* characters pointed to by *src*.

23.20.4.36 `size_t strlen_PF (uint_farptr_t s, size_t len)`

Determine the length of a fixed-size string.

The `strlen_PF()` function is similar to `strlen()`, except that *s* is a far pointer to a string in program space.

Parameters

s A far pointer to the string in Flash

len The maximum number of length to return

Returns

The `strlen_PF` function returns `strlen_P(s)`, if that is less than *len*, or *len* if there is no `'\0'` character among the first *len* characters pointed to by *s*. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.37 `char * strpbrk_P (const char * s, const char * accept)`

The `strpbrk_P()` function locates the first occurrence in the string *s* of any of the characters in the flash string *accept*. This function is similar to `strpbrk()` except that *accept* is a pointer to a string in program space.

Returns

The `strpbrk_P()` function returns a pointer to the character in *s* that matches one of the characters in *accept*, or `NULL` if no such character is found. The terminating zero is not considered as a part of string: if one or both args are empty, the result will `NULL`.

23.20.4.38 `const char * strrchr_P (const char * s, int val)`

Locate character in string.

The `strrchr_P()` function returns a pointer to the last occurrence of the character *val* in the flash string *s*.

Returns

The `strrchr_P()` function returns a pointer to the matched character or `NULL` if the character is not found.

23.20.4.39 `char * strsep_P (char ** sp, const char * delim)`

Parse a string into tokens.

The `strsep_P()` function locates, in the string referenced by **sp*, the first occurrence of any character in the string *delim* (or the terminating `'\0'` character) and replaces it with a `'\0'`. The location of the next character after the delimiter character (or `NULL`, if the end of the string was reached) is stored in **sp*. An “empty” field, i.e. one caused by two adjacent delimiter characters, can be detected by comparing the location referenced by the pointer returned in **sp* to `'\0'`. This function is similar to `strsep()` except that *delim* is a pointer to a string in program space.

Returns

The `strsep_P()` function returns a pointer to the original value of **sp*. If **sp* is initially `NULL`, `strsep_P()` returns `NULL`.

23.20.4.40 `size_t strspn_P (const char * s, const char * accept)`

The `strspn_P()` function calculates the length of the initial segment of `s` which consists entirely of characters in `accept`. This function is similar to `strspn()` except that `accept` is a pointer to a string in program space.

Returns

The `strspn_P()` function returns the number of characters in the initial segment of `s` which consist only of characters from `accept`. The terminating zero is not considered as a part of string.

23.20.4.41 `char * strstr_P (const char * s1, const char * s2)`

Locate a substring.

The `strstr_P()` function finds the first occurrence of the substring `s2` in the string `s1`. The terminating `'\0'` characters are not compared. The `strstr_P()` function is similar to `strstr()` except that `s2` is pointer to a string in program space.

Returns

The `strstr_P()` function returns a pointer to the beginning of the substring, or NULL if the substring is not found. If `s2` points to a string of zero length, the function returns `s1`.

23.20.4.42 `char * strstr_PF (const char * s1, uint_farptr_t s2)`

Locate a substring.

The `strstr_PF()` function finds the first occurrence of the substring `s2` in the string `s1`. The terminating `'\0'` characters are not compared. The `strstr_PF()` function is similar to `strstr()` except that `s2` is a far pointer to a string in program space.

Returns

The `strstr_PF()` function returns a pointer to the beginning of the substring, or NULL if the substring is not found. If `s2` points to a string of zero length, the function returns `s1`. The contents of RAMPZ SFR are undefined when the function returns.

23.20.4.43 `char * strtok_P (char * s, const char * delim)`

Parses the string into tokens.

[strtok_P\(\)](#) parses the string *s* into tokens. The first call to [strtok_P\(\)](#) should have *s* as its first argument. Subsequent calls should have the first argument set to NULL. If a token ends with a delimiter, this delimiting character is overwritten with a `'\0'` and a pointer to the next character is saved for the next call to [strtok_P\(\)](#). The delimiter string *delim* may be different for each call.

The [strtok_P\(\)](#) function is similar to [strtok\(\)](#) except that *delim* is pointer to a string in program space.

Returns

The [strtok_P\(\)](#) function returns a pointer to the next token or NULL when no more tokens are found.

Note

[strtok_P\(\)](#) is NOT reentrant. For a reentrant version of this function see [strtok_rP\(\)](#).

23.20.4.44 `char * strtok_rP (char * string, const char * delim, char ** last)`

Parses string into tokens.

The [strtok_rP\(\)](#) function parses *string* into tokens. The first call to [strtok_rP\(\)](#) should have *string* as its first argument. Subsequent calls should have the first argument set to NULL. If a token ends with a delimiter, this delimiting character is overwritten with a `'\0'` and a pointer to the next character is saved for the next call to [strtok_rP\(\)](#). The delimiter string *delim* may be different for each call. *last* is a user allocated `char*` pointer. It must be the same while parsing the same string. [strtok_rP\(\)](#) is a reentrant version of [strtok_P\(\)](#).

The [strtok_rP\(\)](#) function is similar to [strtok_r\(\)](#) except that *delim* is pointer to a string in program space.

Returns

The [strtok_rP\(\)](#) function returns a pointer to the next token or NULL when no more tokens are found.

23.21 <avr/power.h>: Power Reduction Management

Functions

- void [clock_prescale_set](#) (clock_div_t __x)

23.21.1 Detailed Description

```
#include <avr/power.h>
```

Many AVR devices contain a Power Reduction Register (PRR) or Registers (PRRx) that allow you to reduce power consumption by disabling or enabling various on-board peripherals as needed. Some devices have the XTAL Divide Control Register (XDIV) which offer similar functionality as System Clock Prescale Register (CLKPR).

There are many macros in this header file that provide an easy interface to enable or disable on-board peripherals to reduce power. See the table below.

Note

Not all AVR devices have a Power Reduction Register (for example the ATmega8). On those devices without a Power Reduction Register, the power reduction macros are not available..

Not all AVR devices contain the same peripherals (for example, the LCD interface), or they will be named differently (for example, USART and USART0). Please consult your device's datasheet, or the header file, to find out which macros are applicable to your device.

For device using the XTAL Divide Control Register (XDIV), when prescaler is used, Timer/Counter0 can only be used in asynchronous mode. Keep in mind that Timer/Counter0 source shall be less than 1/4th of peripheral clock. Therefore, when using a typical 32.768 kHz crystal, one shall not scale the clock below 131.072 kHz.

Power Macro	Description
power_aca_disable()	Disable the Analog Comparator on PortA.
power_aca_enable()	Enable the Analog Comparator on PortA.
power_adc_enable()	Enable the Analog to Digital Converter module.
power_adc_disable()	Disable the Analog to Digital Converter module.
power_adca_disable()	Disable the Analog to Digital Converter module on PortA
power_adca_enable()	Enable the Analog to Digital Converter module on PortA
power_evsys_disable()	Disable the EVSYS module
power_evsys_enable()	Enable the EVSYS module
power_hiresc_disable()	Disable the HIRES module on PortC
power_hiresc_enable()	Enable the HIRES module on PortC
power_lcd_enable()	Enable the LCD module.
power_lcd_disable()	Disable the LCD module.
power_pga_enable()	Enable the Programmable Gain Amplifier module.
power_pga_disable()	Disable the Programmable Gain Amplifier module.
power_pscr_enable()	Enable the Reduced Power Stage Controller module.
power_pscr_disable()	Disable the Reduced Power Stage Controller module.
power_psc0_enable()	Enable the Power Stage Controller 0 module.
power_psc0_disable()	Disable the Power Stage Controller 0 module.
power_psc1_enable()	Enable the Power Stage Controller 1 module.
power_psc1_disable()	Disable the Power Stage Controller 1 module.
power_psc2_enable()	Enable the Power Stage Controller 2 module.
power_psc2_disable()	Disable the Power Stage Controller 2 module.
power_ram0_enable()	Enable the SRAM block 0 .
power_ram0_disable()	Disable the SRAM block 0.
power_ram1_enable()	Enable the SRAM block 1 .
power_ram1_disable()	Disable the SRAM block 1.
power_ram2_enable()	Enable the SRAM block 2 .
power_ram2_disable()	Disable the SRAM block 2.
power_ram3_enable()	Enable the SRAM block 3 .
power_ram3_disable()	Disable the SRAM block 3.
power_rtc_disable()	Disable the RTC module
power_rtc_enable()	Enable the RTC module
power_spi_enable()	Enable the Serial Peripheral Interface module.
power_spi_disable()	Disable the Serial Peripheral Interface module.

Generated on Wed May 11 11:27:14 2022 for avr-libc by Doxygen

Some of the newer AVR's contain a System Clock Prescale Register (CLKPR) that allows you to decrease the system clock frequency and the power consumption when the need for processing power is low. On some earlier AVR's (ATmega103, ATmega64, ATmega128), similar functionality can be achieved through the XTAL Divide Control Register. Below are two macros and an enumerated type that can be used to interface to the Clock Prescale Register or XTAL Divide Control Register.

Note

Not all AVR devices have a clock prescaler. On those devices without a Clock Prescale Register or XTAL Divide Control Register, these macros are not available.

```
typedef enum
{
    clock_div_1 = 0,
    clock_div_2 = 1,
    clock_div_4 = 2,
    clock_div_8 = 3,
    clock_div_16 = 4,
    clock_div_32 = 5,
    clock_div_64 = 6,
    clock_div_128 = 7,
    clock_div_256 = 8,
    clock_div_1_rc = 15, // ATmega128RFA1 only
} clock_div_t;
```

Clock prescaler setting enumerations for device using System Clock Prescale Register.

```
typedef enum
{
    clock_div_1 = 1,
    clock_div_2 = 2,
    clock_div_4 = 4,
    clock_div_8 = 8,
    clock_div_16 = 16,
    clock_div_32 = 32,
    clock_div_64 = 64,
    clock_div_128 = 128
} clock_div_t;
```

Clock prescaler setting enumerations for device using XTAL Divide Control Register.

23.21.2 Function Documentation**23.21.2.1 clock_prescale_set (clock_div_t x)**

Set the clock prescaler register select bits, selecting a system clock division setting. This function is inlined, even if compiler optimizations are disabled.

The type of *x* is `clock_div_t`.

Note

For device with XTAL Divide Control Register (XDIV), *x* can actually range from 1 to 129. Thus, one does not need to use `clock_div_t` type as argument.

23.22 Additional notes from <avr/sfr_defs.h>

The <avr/sfr_defs.h> file is included by all of the <avr/ioXXXX.h> files, which use macros defined here to make the special function register definitions look like C variables or simple constants, depending on the `_SFR_ASM_COMPAT` define. Some examples from <avr/iocanxx.h> to show how to define such macros:

```
#define PORTA    _SFR_IO8(0x02)
#define EEAR     _SFR_IO16(0x21)
#define UDR0     _SFR_MEM8(0xC6)
#define TCNT3    _SFR_MEM16(0x94)
#define CANIDT   _SFR_MEM32(0xF0)
```

If `_SFR_ASM_COMPAT` is not defined, C programs can use names like `PORTA` directly in C expressions (also on the left side of assignment operators) and GCC will do the right thing (use short I/O instructions if possible). The `__SFR_OFFSET` definition is not used in any way in this case.

Define `_SFR_ASM_COMPAT` as 1 to make these names work as simple constants (addresses of the I/O registers). This is necessary when included in preprocessed assembler (*.S) source files, so it is done automatically if `__ASSEMBLER__` is defined. By default, all addresses are defined as if they were memory addresses (used in `lds/sts` instructions). To use these addresses in `in/out` instructions, you must subtract 0x20 from them.

For more backwards compatibility, insert the following at the start of your old assembler source file:

```
#define __SFR_OFFSET 0
```

This automatically subtracts 0x20 from I/O space addresses, but it's a hack, so it is recommended to change your source: wrap such addresses in macros defined here, as shown below. After this is done, the `__SFR_OFFSET` definition is no longer necessary and can be removed.

Real example - this code could be used in a boot loader that is portable between devices with SPMCR at different addresses.

```
<avr/iom163.h>: #define SPMCR _SFR_IO8(0x37)
<avr/iom128.h>: #define SPMCR _SFR_MEM8(0x68)

#if _SFR_IO_REG_P(SPMCR)
    out _SFR_IO_ADDR(SPMCR), r24
#else
    sts _SFR_MEM_ADDR(SPMCR), r24
#endif
```

You can use the `in/out/cbi/sbi/sbic/sbis` instructions, without the `_SFR_IO_REG_P` test, if you know that the register is in the I/O space (as with `SREG`, for example). If it isn't, the assembler will complain (I/O address out of range 0...0x3f), so this should be fairly safe.

If you do not define `__SFR_OFFSET` (so it will be 0x20 by default), all special register addresses are defined as memory addresses (so `SREG` is 0x5f), and (if code size and

speed are not important, and you don't like the ugly #if above) you can always use lds/sts to access them. But, this will not work if `__SFR_OFFSET != 0x20`, so use a different macro (defined only if `__SFR_OFFSET == 0x20`) for safety:

```
sts _SFR_ADDR(SPMCR), r24
```

In C programs, all 3 combinations of `_SFR_ASM_COMPAT` and `__SFR_OFFSET` are supported - the `_SFR_ADDR(SPMCR)` macro can be used to get the address of the SPMCR register (0x57 or 0x68 depending on device).

23.23 <avr/sfr_defs.h>: Special function registers

Modules

- [Additional notes from <avr/sfr_defs.h>](#)

Bit manipulation

- `#define _BV(bit) (1 << (bit))`

IO register bit manipulation

- `#define bit_is_set(sfr, bit) (_SFR_BYTE(sfr) & _BV(bit))`
- `#define bit_is_clear(sfr, bit) (!_SFR_BYTE(sfr) & _BV(bit))`
- `#define loop_until_bit_is_set(sfr, bit) do { } while (bit_is_clear(sfr, bit))`
- `#define loop_until_bit_is_clear(sfr, bit) do { } while (bit_is_set(sfr, bit))`

23.23.1 Detailed Description

When working with microcontrollers, many tasks usually consist of controlling internal peripherals, or external peripherals that are connected to the device. The entire IO address space is made available as *memory-mapped IO*, i.e. it can be accessed using all the MCU instructions that are applicable to normal data memory. For most AVR devices, the IO register space is mapped into the data memory address space with an offset of 0x20 since the bottom of this space is reserved for direct access to the MCU registers. (Actual SRAM is available only behind the IO register area, starting at some specific address depending on the device.)

For example the user can access memory-mapped IO registers as if they were globally defined variables like this:

```
PORTA = 0x33;
unsigned char foo = PINA;
```

The compiler will choose the correct instruction sequence to generate based on the address of the register being accessed.

The advantage of using the memory-mapped registers in C programs is that it makes the programs more portable to other C compilers for the AVR platform.

Note that special care must be taken when accessing some of the 16-bit timer IO registers where access from both the main program and within an interrupt context can happen. See [Why do some 16-bit timer registers sometimes get trashed?](#).

Porting programs that use the deprecated sbi/cbi macros

Access to the AVR single bit set and clear instructions are provided via the standard C bit manipulation commands. The sbi and cbi macros are no longer directly supported. sbi (sfr,bit) can be replaced by sfr |= [_BV\(bit\)](#) .

i.e.: [sbi\(PORTB, PB1\)](#); is now `PORTB |= _BV\(PB1\)`;

This actually is more flexible than having sbi directly, as the optimizer will use a hardware sbi if appropriate, or a read/or/write operation if not appropriate. You do not need to keep track of which registers sbi/cbi will operate on.

Likewise, cbi (sfr,bit) is now `sfr &= ~(_BV\(bit\))`;

23.23.2 Define Documentation

23.23.2.1 #define [_BV\(bit\)](#) (1 << (bit))

```
#include <avr/io.h>
```

Converts a bit number into a byte value.

Note

The bit shift is performed by the compiler which then inserts the result into the code. Thus, there is no run-time overhead when using [_BV\(\)](#).

23.23.2.2 #define [bit_is_clear\(sfr, bit\)](#) (!(_SFR_BYTE(sfr) & [_BV\(bit\)](#)))

```
#include <avr/io.h>
```

Test whether bit `bit` in IO register `sfr` is clear. This will return non-zero if the bit is clear, and a 0 if the bit is set.

23.23.2.3 #define bit_is_set(sfr, bit) (_SFR_BYTE(sfr) & _BV(bit))

```
#include <avr/io.h>
```

Test whether bit `bit` in IO register `sfr` is set. This will return a 0 if the bit is clear, and non-zero if the bit is set.

23.23.2.4 #define loop_until_bit_is_clear(sfr, bit) do { } while (bit_is_set(sfr, bit))

```
#include <avr/io.h>
```

Wait until bit `bit` in IO register `sfr` is clear.

23.23.2.5 #define loop_until_bit_is_set(sfr, bit) do { } while (bit_is_clear(sfr, bit))

```
#include <avr/io.h>
```

Wait until bit `bit` in IO register `sfr` is set.

23.24 <avr/signature.h>: Signature Support**Introduction**

The [<avr/signature.h>](#) header file allows the user to automatically and easily include the device's signature data in a special section of the final linked ELF file.

This value can then be used by programming software to compare the on-device signature with the signature recorded in the ELF file to look for a match before programming the device.

API Usage Example

Usage is very simple; just include the header file:

```
#include <avr/signature.h>
```

This will declare a constant unsigned char array and it is initialized with the three signature bytes, MSB first, that are defined in the device I/O header file. This array is then placed in the .signature section in the resulting linked ELF file.

The three signature bytes that are used to initialize the array are these defined macros in the device I/O header file, from MSB to LSB: SIGNATURE_2, SIGNATURE_1, SIGNATURE_0.

This header file should only be included once in an application.

23.25 <avr/sleep.h>: Power Management and Sleep Modes

Functions

- void `sleep_enable` (void)
- void `sleep_disable` (void)
- void `sleep_cpu` (void)
- void `sleep_mode` (void)
- void `sleep_bod_disable` (void)

23.25.1 Detailed Description

```
#include <avr/sleep.h>
```

Use of the `SLEEP` instruction can allow an application to reduce its power consumption considerably. AVR devices can be put into different sleep modes. Refer to the datasheet for the details relating to the device you are using.

There are several macros provided in this header file to actually put the device into sleep mode. The simplest way is to optionally set the desired sleep mode using `set_sleep_mode()` (it usually defaults to idle mode where the CPU is put on sleep but all peripheral clocks are still running), and then call `sleep_mode()`. This macro automatically sets the sleep enable bit, goes to sleep, and clears the sleep enable bit.

Example:

```
#include <avr/sleep.h>

...
set_sleep_mode(<mode>);
sleep_mode();
```

Note that unless your purpose is to completely lock the CPU (until a hardware reset), interrupts need to be enabled before going to sleep.

As the `sleep_mode()` macro might cause race conditions in some situations, the individual steps of manipulating the sleep enable (SE) bit, and actually issuing the `SLEEP` instruction, are provided in the macros `sleep_enable()`, `sleep_disable()`, and `sleep_cpu()`. This also allows for test-and-sleep scenarios that take care of not missing the interrupt that will awake the device from sleep.

Example:

```
#include <avr/interrupt.h>
#include <avr/sleep.h>

...
set_sleep_mode(<mode>);
cli();
if (some_condition)
{
    sleep_enable();
    sei();
    sleep_cpu();
    sleep_disable();
}
sei();
```

This sequence ensures an atomic test of `some_condition` with interrupts being disabled. If the condition is met, sleep mode will be prepared, and the `SLEEP` instruction will be scheduled immediately after an `SEI` instruction. As the instruction right after the `SEI` is guaranteed to be executed before an interrupt could trigger, it is sure the device will really be put to sleep.

Some devices have the ability to disable the Brown Out Detector (BOD) before going to sleep. This will also reduce power while sleeping. If the specific AVR device has this ability then an additional macro is defined: `sleep_bod_disable()`. This macro generates inlined assembly code that will correctly implement the timed sequence for disabling the BOD before sleeping. However, there is a limited number of cycles after the BOD has been disabled that the device can be put into sleep mode, otherwise the BOD will not truly be disabled. Recommended practice is to disable the BOD (`sleep_bod_disable()`), set the interrupts (`sei()`), and then put the device to sleep (`sleep_cpu()`), like so:

```
#include <avr/interrupt.h>
#include <avr/sleep.h>

...
set_sleep_mode(<mode>);
cli();
if (some_condition)
{
    sleep_enable();
    sleep_bod_disable();
    sei();
    sleep_cpu();
    sleep_disable();
}
sei();
```

23.25.2 Function Documentation

23.25.2.1 void sleep_bod_disable(void)

Disable BOD before going to sleep. Not available on all devices.

23.25.2.2 void sleep_cpu (void)

Put the device into sleep mode. The SE bit must be set beforehand, and it is recommended to clear it afterwards.

23.25.2.3 void sleep_disable (void)

Clear the SE (sleep enable) bit.

23.25.2.4 void sleep_enable (void)

Put the device in sleep mode. How the device is brought out of sleep mode depends on the specific mode selected with the set_sleep_mode() function. See the data sheet for your device for more details.

Set the SE (sleep enable) bit.

23.25.2.5 void sleep_mode (void)

Put the device into sleep mode, taking care of setting the SE bit before, and clearing it afterwards.

23.26 <avr/version.h>: avr-libc version macros

Defines

- #define `__AVR_LIBC_VERSION_STRING__` "2.0.0"
- #define `__AVR_LIBC_VERSION__` 20000UL
- #define `__AVR_LIBC_DATE_STRING__` "20150208"
- #define `__AVR_LIBC_DATE__` 20150208UL
- #define `__AVR_LIBC_MAJOR__` 2
- #define `__AVR_LIBC_MINOR__` 0
- #define `__AVR_LIBC_REVISION__` 0

23.26.1 Detailed Description

```
#include <avr/version.h>
```

This header file defines macros that contain version numbers and strings describing the current version of avr-libc.

The version number itself basically consists of three pieces that are separated by a dot: the major number, the minor number, and the revision number. For development versions (which use an odd minor number), the string representation additionally gets the date code (YYYYMMDD) appended.

This file will also be included by <avr/io.h>. That way, portable tests can be implemented using <avr/io.h> that can be used in code that wants to remain backwards-compatible to library versions prior to the date when the library version API had been added, as referenced but undefined C preprocessor macros automatically evaluate to 0.

23.26.2 Define Documentation

23.26.2.1 #define __AVR_LIBC_DATE_ 20150208UL

Numerical representation of the release date.

23.26.2.2 #define __AVR_LIBC_DATE_STRING__ "20150208"

String literal representation of the release date.

23.26.2.3 #define __AVR_LIBC_MAJOR__ 2

Library major version number.

23.26.2.4 #define __AVR_LIBC_MINOR__ 0

Library minor version number.

23.26.2.5 #define __AVR_LIBC_REVISION__ 0

Library revision number.

23.26.2.6 #define __AVR_LIBC_VERSION__ 20000UL

Numerical representation of the current library version.

In the numerical representation, the major number is multiplied by 10000, the minor number by 100, and all three parts are then added. It is intended to provide a monotonically increasing numerical value that can easily be used in numerical checks.

23.26.2.7 #define __AVR_LIBC_VERSION_STRING__ "2.0.0"

String literal representation of the current library version.

23.27 <avr/wdt.h>: Watchdog timer handling

Defines

- #define `wdt_reset()` `__asm__ __volatile__ ("wdr")`

- #define WDTO_15MS 0
- #define WDTO_30MS 1
- #define WDTO_60MS 2
- #define WDTO_120MS 3
- #define WDTO_250MS 4
- #define WDTO_500MS 5
- #define WDTO_1S 6
- #define WDTO_2S 7
- #define WDTO_4S 8
- #define WDTO_8S 9

Functions

- static `__inline__ __attribute__((__always_inline__)) void wdt_enable(const uint8_t value)`

23.27.1 Detailed Description

```
#include <avr/wdt.h>
```

This header file declares the interface to some inline macros handling the watchdog timer present in many AVR devices. In order to prevent the watchdog timer configuration from being accidentally altered by a crashing application, a special timed sequence is required in order to change it. The macros within this header file handle the required sequence automatically before changing any value. Interrupts will be disabled during the manipulation.

Note

Depending on the fuse configuration of the particular device, further restrictions might apply, in particular it might be disallowed to turn off the watchdog timer.

Note that for newer devices (ATmega88 and newer, effectively any AVR that has the option to also generate interrupts), the watchdog timer remains active even after a system reset (except a power-on condition), using the fastest prescaler value (approximately 15 ms). It is therefore required to turn off the watchdog early during program startup, the datasheet recommends a sequence like the following:

```
#include <stdint.h>
#include <avr/wdt.h>

uint8_t mcusr_mirror __attribute__((section (".noinit")));

void get_mcusr(void) \
    __attribute__((naked)) \
    __attribute__((section (".init3")));
void get_mcusr(void)
{
    mcusr_mirror = MCUSR;
    MCUSR = 0;
    wdt_disable();
}
```


Saving the value of MCUSR in `mcusr_mirror` is only needed if the application later wants to examine the reset source, but in particular, clearing the watchdog reset flag before disabling the watchdog is required, according to the datasheet.

23.27.2 Define Documentation

23.27.2.1 `#define wdt_reset() __asm__ __volatile__ ("wdr")`

Reset the watchdog timer. When the watchdog timer is enabled, a call to this instruction is required before the timer expires, otherwise a watchdog-initiated device reset will occur.

23.27.2.2 `#define WDTO_120MS 3`

See `WDTO_15MS`

23.27.2.3 `#define WDTO_15MS 0`

Symbolic constants for the watchdog timeout. Since the watchdog timer is based on a free-running RC oscillator, the times are approximate only and apply to a supply voltage of 5 V. At lower supply voltages, the times will increase. For older devices, the times will be as large as three times when operating at $V_{cc} = 3$ V, while the newer devices (e. g. ATmega128, ATmega8) only experience a negligible change.

Possible timeout values are: 15 ms, 30 ms, 60 ms, 120 ms, 250 ms, 500 ms, 1 s, 2 s. (Some devices also allow for 4 s and 8 s.) Symbolic constants are formed by the prefix `WDTO_`, followed by the time.

Example that would select a watchdog timer expiry of approximately 500 ms:

```
wdt_enable(WDTO_500MS);
```

23.27.2.4 `#define WDTO_1S 6`

See `WDTO_15MS`

23.27.2.5 `#define WDTO_250MS 4`

See `WDTO_15MS`

23.27.2.6 `#define WDTO_2S 7`

See `WDTO_15MS`

23.27.2.7 #define WDTO_30MS 1

See WDTO_15MS

23.27.2.8 #define WDTO_4S 8

See WDTO_15MS Note: This is only available on the ATtiny2313, ATtiny24, ATtiny44, ATtiny84, ATtiny84A, ATtiny25, ATtiny45, ATtiny85, ATtiny261, ATtiny461, ATtiny861, ATmega48, ATmega88, ATmega168, ATmega48P, ATmega88P, ATmega168P, ATmega328P, ATmega164P, ATmega324P, ATmega644P, ATmega644, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega8HVA, ATmega16HVA, ATmega32HVB, ATmega406, ATmega1284P, AT90PWM1, AT90PWM2, AT90PWM2B, AT90PWM3, AT90PWM3B, AT90PWM216, AT90PWM316, AT90PWM81, AT90PWM161, AT90USB82, AT90USB162, AT90USB646, AT90USB647, AT90USB1286, AT90USB1287, ATtiny48, ATtiny88.

Note: This value does *not* match the bit pattern of the respective control register. It is solely meant to be used together with `wdt_enable()`.

23.27.2.9 #define WDTO_500MS 5

See WDTO_15MS

23.27.2.10 #define WDTO_60MS 2

See WDTO_15MS

23.27.2.11 #define WDTO_8S 9

See WDTO_15MS Note: This is only available on the ATtiny2313, ATtiny24, ATtiny44, ATtiny84, ATtiny84A, ATtiny25, ATtiny45, ATtiny85, ATtiny261, ATtiny461, ATtiny861, ATmega48, ATmega48A, ATmega48PA, ATmega88, ATmega168, ATmega48P, ATmega88P, ATmega168P, ATmega328P, ATmega164P, ATmega324P, ATmega644P, ATmega644, ATmega640, ATmega1280, ATmega1281, ATmega2560, ATmega2561, ATmega8HVA, ATmega16HVA, ATmega32HVB, ATmega406, ATmega1284P, ATmega2564RFR2, ATmega256RFR2, ATmega1284RFR2, ATmega128RFR2, ATmega644RFR2, ATmega64RFR2 AT90PWM1, AT90PWM2, AT90PWM2B, AT90PWM3, AT90PWM3B, AT90PWM216, AT90PWM316, AT90PWM81, AT90PWM161, AT90USB82, AT90USB162, AT90USB646, AT90USB647, AT90USB1286, AT90USB1287, ATtiny48, ATtiny88, ATxmega16a4u, ATxmega32a4u, ATxmega16c4, ATxmega32c4, ATxmega128c3, ATxmega192c3, ATxmega256c3.

Note: This value does *not* match the bit pattern of the respective control register. It is solely meant to be used together with `wdt_enable()`.

23.27.3 Function Documentation

23.27.3.1 `static __inline__ __attribute__((__always_inline__)) const [static]`

Enable the watchdog timer, configuring it for expiry after `timeout` (which is a combination of the `WDP0` through `WDP2` bits to write into the `WDTCR` register; For those devices that have a `WDTCR` register, it uses the combination of the `WDP0` through `WDP3` bits).

See also the symbolic constants `WDTO_15MS` et al.

23.28 <util/atomic.h> Atomically and Non-Atomically Executed Code Blocks

Defines

- `#define ATOMIC_BLOCK(type)`
- `#define NONATOMIC_BLOCK(type)`
- `#define ATOMIC_RESTORESTATE`
- `#define ATOMIC_FORCEON`
- `#define NONATOMIC_RESTORESTATE`
- `#define NONATOMIC_FORCEOFF`

23.28.1 Detailed Description

```
#include <util/atomic.h>
```

Note

The macros in this header file require the ISO/IEC 9899:1999 ("ISO C99") feature of for loop variables that are declared inside the for loop itself. For that reason, this header file can only be used if the standard level of the compiler (option `--std=`) is set to either `c99` or `gnu99`.

The macros in this header file deal with code blocks that are guaranteed to be executed Atomically or Non-Atomically. The term "Atomic" in this context refers to the inability of the respective code to be interrupted.

These macros operate via automatic manipulation of the Global Interrupt Status (I) bit of the `SREG` register. Exit paths from both block types are all managed automatically without the need for special considerations, i. e. the interrupt status will be restored to the same value it has been when entering the respective block.

A typical example that requires atomic access is a 16 (or more) bit variable that is shared between the main execution path and an ISR. While declaring such a variable as volatile ensures that the compiler will not optimize accesses to it away, it does not guarantee atomic access to it. Assuming the following example:

```
#include <inttypes.h>
#include <avr/interrupt.h>
#include <avr/io.h>

volatile uint16_t ctr;

ISR(TIMER1_OVF_vect)
{
    ctr--;
}

...
int
main(void)
{
    ...
    ctr = 0x200;
    start_timer();
    while (ctr != 0)
        // wait
        ;
    ...
}
```

There is a chance where the main context will exit its wait loop when the variable `ctr` just reached the value `0xFF`. This happens because the compiler cannot natively access a 16-bit variable atomically in an 8-bit CPU. So the variable is for example at `0x100`, the compiler then tests the low byte for 0, which succeeds. It then proceeds to test the high byte, but that moment the ISR triggers, and the main context is interrupted. The ISR will decrement the variable from `0x100` to `0xFF`, and the main context proceeds. It now tests the high byte of the variable which is (now) also 0, so it concludes the variable has reached 0, and terminates the loop.

Using the macros from this header file, the above code can be rewritten like:

```
#include <inttypes.h>
#include <avr/interrupt.h>
#include <avr/io.h>
#include <util/atomic.h>

volatile uint16_t ctr;

ISR(TIMER1_OVF_vect)
{
    ctr--;
}

...
int
main(void)
{
    ...
    ctr = 0x200;
    start_timer();
    sei();
    uint16_t ctr_copy;
    do
    {
        ATOMIC_BLOCK(ATOMIC_FORCEON)
        {
```

23.28 <util/atomic.h> Atomically and Non-Atomically Executed Code Blocks 305

```
        ctr_copy = ctr;
    }
}
while (ctr_copy != 0);
...
}
```

This will install the appropriate interrupt protection before accessing variable `ctr`, so it is guaranteed to be consistently tested. If the global interrupt state were uncertain before entering the `ATOMIC_BLOCK`, it should be executed with the parameter `ATOMIC_RESTORESTATE` rather than `ATOMIC_FORCEON`.

See [Problems with reordering code](#) for things to be taken into account with respect to compiler optimizations.

23.28.2 Define Documentation

23.28.2.1 `#define ATOMIC_BLOCK(type)`

Creates a block of code that is guaranteed to be executed atomically. Upon entering the block the Global Interrupt Status flag in `SREG` is disabled, and re-enabled upon exiting the block from any exit path.

Two possible macro parameters are permitted, `ATOMIC_RESTORESTATE` and `ATOMIC_FORCEON`.

23.28.2.2 `#define ATOMIC_FORCEON`

This is a possible parameter for `ATOMIC_BLOCK`. When used, it will cause the `ATOMIC_BLOCK` to force the state of the `SREG` register on exit, enabling the Global Interrupt Status flag bit. This saves on flash space as the previous value of the `SREG` register does not need to be saved at the start of the block.

Care should be taken that `ATOMIC_FORCEON` is only used when it is known that interrupts are enabled before the block's execution or when the side effects of enabling global interrupts at the block's completion are known and understood.

23.28.2.3 `#define ATOMIC_RESTORESTATE`

This is a possible parameter for `ATOMIC_BLOCK`. When used, it will cause the `ATOMIC_BLOCK` to restore the previous state of the `SREG` register, saved before the Global Interrupt Status flag bit was disabled. The net effect of this is to make the `ATOMIC_BLOCK`'s contents guaranteed atomic, without changing the state of the Global Interrupt Status flag when execution of the block completes.

23.28.2.4 `#define NONATOMIC_BLOCK(type)`

Creates a block of code that is executed non-atomically. Upon entering the block the Global Interrupt Status flag in `SREG` is

enabled, and disabled upon exiting the block from any exit path. This is useful when nested inside `ATOMIC_BLOCK` sections, allowing for non-atomic execution of small blocks of code while maintaining the atomic access of the other sections of the parent `ATOMIC_BLOCK`.

Two possible macro parameters are permitted, `NONATOMIC_RESTORESTATE` and `NONATOMIC_FORCEOFF`.

23.28.2.5 `#define NONATOMIC_FORCEOFF`

This is a possible parameter for `NONATOMIC_BLOCK`. When used, it will cause the `NONATOMIC_BLOCK` to force the state of the `SREG` register on exit, disabling the Global Interrupt Status flag bit. This saves on flash space as the previous value of the `SREG` register does not need to be saved at the start of the block.

Care should be taken that `NONATOMIC_FORCEOFF` is only used when it is known that interrupts are disabled before the block's execution or when the side effects of disabling global interrupts at the block's completion are known and understood.

23.28.2.6 `#define NONATOMIC_RESTORESTATE`

This is a possible parameter for `NONATOMIC_BLOCK`. When used, it will cause the `NONATOMIC_BLOCK` to restore the previous state of the `SREG` register, saved before the Global Interrupt Status flag bit was enabled. The net effect of this is to make the `NONATOMIC_BLOCK`'s contents guaranteed non-atomic, without changing the state of the Global Interrupt Status flag when execution of the block completes.

23.29 <util/crc16.h>: CRC Computations

Functions

- static `__inline__ uint16_t _crc16_update (uint16_t __crc, uint8_t __data)`
- static `__inline__ uint16_t _crc_xmodem_update (uint16_t __crc, uint8_t __data)`
- static `__inline__ uint16_t _crc_ccitt_update (uint16_t __crc, uint8_t __data)`
- static `__inline__ uint8_t _crc_ibutton_update (uint8_t __crc, uint8_t __data)`
- static `__inline__ uint8_t _crc8_ccitt_update (uint8_t __crc, uint8_t __data)`

23.29.1 Detailed Description

```
#include <util/crc16.h>
```

This header file provides a optimized inline functions for calculating cyclic redundancy checks (CRC) using common polynomials.

References:

See the Dallas Semiconductor app note 27 for 8051 assembler example and general CRC optimization suggestions. The table on the last page of the app note is the key to understanding these implementations.

Jack Crenshaw's "Implementing CRCs" article in the January 1992 issue of *Embedded Systems Programming*. This may be difficult to find, but it explains CRC's in very clear and concise terms. Well worth the effort to obtain a copy.

A typical application would look like:

```
// Dallas iButton test vector.
uint8_t serno[] = { 0x02, 0x1c, 0xb8, 0x01, 0, 0, 0, 0xa2 };

int
checkcrc(void)
{
    uint8_t crc = 0, i;

    for (i = 0; i < sizeof serno / sizeof serno[0]; i++)
        crc = _crc_ibutton_update(crc, serno[i]);

    return crc; // must be 0
}
```

23.29.2 Function Documentation

23.29.2.1 `static __inline__ uint16_t _crc16_update (uint16_t __crc, uint8_t __data) [static]`

Optimized CRC-16 calculation.

Polynomial: $x^{16} + x^{15} + x^2 + 1$ (0xa001)

Initial value: 0xffff

This CRC is normally used in disk-drive controllers.

The following is the equivalent functionality written in C.

```
uint16_t
crc16_update(uint16_t crc, uint8_t a)
{
    int i;

    crc ^= a;
    for (i = 0; i < 8; ++i)
    {
        if (crc & 1)
            crc = (crc >> 1) ^ 0xA001;
        else
            crc >>= 1;
    }
}
```

```

        crc = (crc >> 1);
    }

    return crc;
}

```

23.29.2.2 `static __inline__ uint8_t _crc8_ccitt_update (uint8_t __crc, uint8_t __data) [static]`

Optimized CRC-8-CCITT calculation.

Polynomial: $x^8 + x^2 + x + 1$ (0xE0)

For use with simple CRC-8

Initial value: 0x0

For use with CRC-8-ROHC

Initial value: 0xff

Reference: <http://tools.ietf.org/html/rfc3095#section-5.9.1>

For use with CRC-8-ATM/ITU

Initial value: 0xff

Final XOR value: 0x55

Reference: <http://www.itu.int/rec/T-REC-I.432.1-199902-I/en>

The C equivalent has been originally written by Dave Hylands. Assembly code is based on `_crc_ibutton_update` optimization.

The following is the equivalent functionality written in C.

```

uint8_t
_crc8_ccitt_update (uint8_t inCrc, uint8_t inData)
{
    uint8_t  i;
    uint8_t  data;

    data = inCrc ^ inData;

    for ( i = 0; i < 8; i++ )
    {
        if (( data & 0x80 ) != 0 )
        {
            data <<= 1;
            data ^= 0x07;
        }
        else
        {
            data <<= 1;
        }
    }
    return data;
}

```


23.29.2.3 `static __inline__ uint16_t _crc_ccitt_update (uint16_t __crc, uint8_t __data) [static]`

Optimized CRC-CCITT calculation.

Polynomial: $x^{16} + x^{12} + x^5 + 1$ (0x8408)

Initial value: 0xffff

This is the CRC used by PPP and IrDA.

See RFC1171 (PPP protocol) and IrDA IrLAP 1.1

Note

Although the CCITT polynomial is the same as that used by the Xmodem protocol, they are quite different. The difference is in how the bits are shifted through the algorithm. Xmodem shifts the MSB of the CRC and the input first, while CCITT shifts the LSB of the CRC and the input first.

The following is the equivalent functionality written in C.

```
uint16_t
_crc_ccitt_update (uint16_t crc, uint8_t data)
{
    data ^= lo8 (crc);
    data ^= data << 4;

    return (((uint16_t)data << 8) | hi8 (crc)) ^ (uint8_t)(data >> 4)
           ^ ((uint16_t)data << 3));
}
```

23.29.2.4 `static __inline__ uint8_t _crc_ibutton_update (uint8_t __crc, uint8_t __data) [static]`

Optimized Dallas (now Maxim) iButton 8-bit CRC calculation.

Polynomial: $x^8 + x^5 + x^4 + 1$ (0x8C)

Initial value: 0x0

See http://www.maxim-ic.com/appnotes.cfm/appnote_number/27

The following is the equivalent functionality written in C.

```
uint8_t
_crc_ibutton_update(uint8_t crc, uint8_t data)
{
    uint8_t i;

    crc = crc ^ data;
    for (i = 0; i < 8; i++)
    {
        if (crc & 0x01)
            crc = (crc >> 1) ^ 0x8C;
        else
            crc = (crc >> 1);
    }
}
```

```

        crc >>= 1;
    }

    return crc;
}

```

23.29.2.5 static __inline__ uint16_t _crc_xmodem_update (uint16_t __crc, uint8_t __data) [static]

Optimized CRC-XMODEM calculation.

Polynomial: $x^{16} + x^{12} + x^5 + 1$ (0x1021)

Initial value: 0x0

This is the CRC used by the Xmodem-CRC protocol.

The following is the equivalent functionality written in C.

```

uint16_t
_crc_xmodem_update (uint16_t crc, uint8_t data)
{
    int i;

    crc = crc ^ ((uint16_t)data << 8);
    for (i=0; i<8; i++)
    {
        if (crc & 0x8000)
            crc = (crc << 1) ^ 0x1021;
        else
            crc <<= 1;
    }

    return crc;
}

```

23.30 <util/delay.h>: Convenience functions for busy-wait delay loops

Defines

- #define [F_CPU](#) 1000000UL

Functions

- void [_delay_ms](#) (double __ms)
- void [_delay_us](#) (double __us)

23.30.1 Detailed Description

```

#define F_CPU 1000000UL // 1 MHz
// #define F_CPU 14.7456E6
#include <util/delay.h>

```

Note

As an alternative method, it is possible to pass the `F_CPU` macro down to the compiler from the Makefile. Obviously, in that case, no `#define` statement should be used.

The functions in this header file are wrappers around the basic busy-wait functions from `<util/delay_basic.h>`. They are meant as convenience functions where actual time values can be specified rather than a number of cycles to wait for. The idea behind is that compile-time constant expressions will be eliminated by compiler optimization so floating-point expressions can be used to calculate the number of delay cycles needed based on the CPU frequency passed by the macro `F_CPU`.

Note

In order for these functions to work as intended, compiler optimizations *must* be enabled, and the delay time *must* be an expression that is a known constant at compile-time. If these requirements are not met, the resulting delay will be much longer (and basically unpredictable), and applications that otherwise do not use floating-point calculations will experience severe code bloat by the floating-point library routines linked into the application.

The functions available allow the specification of microsecond, and millisecond delays directly, using the application-supplied macro `F_CPU` as the CPU clock frequency (in Hertz).

23.30.2 Define Documentation**23.30.2.1 `#define F_CPU 1000000UL`**

CPU frequency in Hz.

The macro `F_CPU` specifies the CPU frequency to be considered by the delay macros. This macro is normally supplied by the environment (e.g. from within a project header, or the project's Makefile). The value 1 MHz here is only provided as a "vanilla" fallback if no such user-provided definition could be found.

In terms of the delay functions, the CPU frequency can be given as a floating-point constant (e.g. `3.6864E6` for 3.6864 MHz). However, the macros in `<util/setbaud.h>` require it to be an integer value.

23.30.3 Function Documentation**23.30.3.1 `void _delay_ms (double __ms)`**

Perform a delay of `__ms` milliseconds, using `_delay_loop_2()`.

The macro `F_CPU` is supposed to be defined to a constant defining the CPU clock frequency (in Hertz).

The maximal possible delay is $262.14 \text{ ms} / F_{\text{CPU}}$ in MHz.

When the user request delay which exceed the maximum possible one, `_delay_ms()` provides a decreased resolution functionality. In this mode `_delay_ms()` will work with a resolution of 1/10 ms, providing delays up to 6.5535 seconds (independent from CPU frequency). The user will not be informed about decreased resolution.

If the avr-gcc toolchain has `__builtin_avr_delay_cycles()` support, maximal possible delay is $4294967.295 \text{ ms} / F_{\text{CPU}}$ in MHz. For values greater than the maximal possible delay, overflows results in no delay i.e., 0ms.

Conversion of `__ms` into clock cycles may not always result in integer. By default, the clock cycles rounded up to next integer. This ensures that the user gets at least `__ms` microseconds of delay.

Alternatively, by defining the macro `__DELAY_ROUND_DOWN__`, or `__DELAY_ROUND_CLOSEST__`, before including this header file, the algorithm can be made to round down, or round to closest integer, respectively.

Note

The implementation of `_delay_ms()` based on `__builtin_avr_delay_cycles()` is not backward compatible with older implementations. In order to get functionality backward compatible with previous versions, the macro "`__DELAY_BACKWARD_COMPATIBLE__`" must be defined before including this header file. Also, the backward compatible algorithm will be chosen if the code is compiled in a *freestanding environment* (GCC option `-ffreestanding`), as the math functions required for rounding are not available to the compiler then.

23.30.3.2 void _delay_us (double __us)

Perform a delay of `__us` microseconds, using `_delay_loop_1()`.

The macro `F_CPU` is supposed to be defined to a constant defining the CPU clock frequency (in Hertz).

The maximal possible delay is $768 \text{ us} / F_{\text{CPU}}$ in MHz.

If the user requests a delay greater than the maximal possible one, `_delay_us()` will automatically call `_delay_ms()` instead. The user will not be informed about this case.

If the avr-gcc toolchain has `__builtin_avr_delay_cycles()` support, maximal possible delay is $4294967.295 \text{ us} / F_{\text{CPU}}$ in MHz. For values greater than the maximal possible delay, overflow results in no delay i.e., 0us.

Conversion of `__us` into clock cycles may not always result in integer. By default, the clock cycles rounded up to next integer. This ensures that the user gets at least `__us` microseconds of delay.

Alternatively, by defining the macro `__DELAY_ROUND_DOWN__`, or `__DELAY_ROUND_CLOSEST__`, before including this header file, the algorithm can be made to round down, or round to closest integer, respectively.

Note

The implementation of `_delay_ms()` based on `__builtin_avr_delay_cycles()` is not backward compatible with older implementations. In order to get functionality backward compatible with previous versions, the macro `__DELAY_BACKWARD_COMPATIBLE__` must be defined before including this header file. Also, the backward compatible algorithm will be chosen if the code is compiled in a *freestanding environment* (GCC option `-ffreestanding`), as the math functions required for rounding are not available to the compiler then.

23.31 <util/delay_basic.h>: Basic busy-wait delay loops**Functions**

- void `_delay_loop_1 (uint8_t __count)`
- void `_delay_loop_2 (uint16_t __count)`

23.31.1 Detailed Description

```
#include <util/delay_basic.h>
```

The functions in this header file implement simple delay loops that perform a busy-waiting. They are typically used to facilitate short delays in the program execution. They are implemented as count-down loops with a well-known CPU cycle count per loop iteration. As such, no other processing can occur simultaneously. It should be kept in mind that the functions described here do not disable interrupts.

In general, for long delays, the use of hardware timers is much preferable, as they free the CPU, and allow for concurrent processing of other events while the timer is running. However, in particular for very short delays, the overhead of setting up a hardware timer is too much compared to the overall delay time.

Two inline functions are provided for the actual delay algorithms.

23.31.2 Function Documentation**23.31.2.1 void _delay_loop_1 (uint8_t __count)**

Delay loop using an 8-bit counter `__count`, so up to 256 iterations are possible. (The value 256 would have to be passed as 0.) The loop executes three CPU cycles per iteration, not including the overhead the compiler needs to setup the counter register.

Thus, at a CPU speed of 1 MHz, delays of up to 768 microseconds can be achieved.

23.31.2.2 void _delay_loop_2 (uint16_t __count)

Delay loop using a 16-bit counter `__count`, so up to 65536 iterations are possible. (The value 65536 would have to be passed as 0.) The loop executes four

CPU cycles per iteration, not including the overhead the compiler requires to setup the counter register pair.

Thus, at a CPU speed of 1 MHz, delays of up to about 262.1 milliseconds can be achieved.

23.32 <util/parity.h>: Parity bit generation

Defines

- #define [parity_even_bit](#)(val)

23.32.1 Detailed Description

```
#include <util/parity.h>
```

This header file contains optimized assembler code to calculate the parity bit for a byte.

23.32.2 Define Documentation

23.32.2.1 #define [parity_even_bit](#)(val)

Value:

```
(__extension__({
    unsigned char __t;
    __asm__ (
        "mov __tmp_reg__, %0" "\n\t" \
        "swap %0" "\n\t" \
        "eor %0, __tmp_reg__" "\n\t" \
        "mov __tmp_reg__, %0" "\n\t" \
        "lsr %0" "\n\t" \
        "lsr %0" "\n\t" \
        "eor %0, __tmp_reg__" \
        : "=r" (__t) \
        : "0" ((unsigned char)(val)) \
        : "r0" \
    );
    (((__t + 1) >> 1) & 1);
})))
```

Returns

1 if `val` has an odd number of bits set.

23.33 <util/setbaud.h>: Helper macros for baud rate calculations

Defines

- #define [BAUD_TOL](#) 2
- #define [UBRR_VALUE](#)

- `#define UBRRL_VALUE`
- `#define UBRRH_VALUE`
- `#define USE_2X 0`

23.33.1 Detailed Description

```
#define F_CPU 11059200
#define BAUD 38400
#include <util/setbaud.h>
```

This header file requires that on entry values are already defined for `F_CPU` and `BAUD`. In addition, the macro `BAUD_TOL` will define the baud rate tolerance (in percent) that is acceptable during the calculations. The value of `BAUD_TOL` will default to 2 %.

This header file defines macros suitable to setup the UART baud rate prescaler registers of an AVR. All calculations are done using the C preprocessor. Including this header file causes no other side effects so it is possible to include this file more than once (supposedly, with different values for the `BAUD` parameter), possibly even within the same function.

Assuming that the requested `BAUD` is valid for the given `F_CPU` then the macro `UBRR_VALUE` is set to the required prescaler value. Two additional macros are provided for the low and high bytes of the prescaler, respectively: `UBRRL_VALUE` is set to the lower byte of the `UBRR_VALUE` and `UBRRH_VALUE` is set to the upper byte. An additional macro `USE_2X` will be defined. Its value is set to 1 if the desired `BAUD` rate within the given tolerance could only be achieved by setting the `U2X` bit in the UART configuration. It will be defined to 0 if `U2X` is not needed.

Example usage:

```
#include <avr/io.h>

#define F_CPU 4000000

static void
uart_9600(void)
{
    #define BAUD 9600
    #include <util/setbaud.h>
    UBRRH = UBRRH_VALUE;
    UBRRL = UBRRL_VALUE;
    #if USE_2X
    UCSRA |= (1 << U2X);
    #else
    UCSRA &= ~(1 << U2X);
    #endif
}

static void
uart_38400(void)
{
    #undef BAUD // avoid compiler warning
    #define BAUD 38400
    #include <util/setbaud.h>
    UBRRH = UBRRH_VALUE;
    UBRRL = UBRRL_VALUE;
```

```
#if USE_2X
UCSRA |= (1 << U2X);
#else
UCSRA &= ~(1 << U2X);
#endif
}
```

In this example, two functions are defined to setup the UART to run at 9600 Bd, and 38400 Bd, respectively. Using a CPU clock of 4 MHz, 9600 Bd can be achieved with an acceptable tolerance without setting U2X (prescaler 25), while 38400 Bd require U2X to be set (prescaler 12).

23.33.2 Define Documentation

23.33.2.1 #define BAUD_TOL 2

Input and output macro from <[util/setbaud.h](#)>

Define the acceptable baud rate tolerance in percent. If not set on entry, it will be set to its default value of 2.

23.33.2.2 #define UBRR_VALUE

Output macro from <[util/setbaud.h](#)>

Contains the calculated baud rate prescaler value for the UBRR register.

23.33.2.3 #define UBRRH_VALUE

Output macro from <[util/setbaud.h](#)>

Contains the upper byte of the calculated prescaler value (UBRR_VALUE).

23.33.2.4 #define UBRRL_VALUE

Output macro from <[util/setbaud.h](#)>

Contains the lower byte of the calculated prescaler value (UBRR_VALUE).

23.33.2.5 #define USE_2X 0

Output macro from <[util/setbaud.h](#)>

Contains the value 1 if the desired baud rate tolerance could only be achieved by setting the U2X bit in the UART configuration. Contains 0 otherwise.

23.34 <util/twi.h>: TWI bit mask definitions**TWSR values**

Mnemonics:

TW_MT_XXX - master transmitter

TW_MR_XXX - master receiver

TW_ST_XXX - slave transmitter

TW_SR_XXX - slave receiver

- #define TW_START 0x08
- #define TW_REP_START 0x10
- #define TW_MT_SLA_ACK 0x18
- #define TW_MT_SLA_NACK 0x20
- #define TW_MT_DATA_ACK 0x28
- #define TW_MT_DATA_NACK 0x30
- #define TW_MT_ARB_LOST 0x38
- #define TW_MR_ARB_LOST 0x38
- #define TW_MR_SLA_ACK 0x40
- #define TW_MR_SLA_NACK 0x48
- #define TW_MR_DATA_ACK 0x50
- #define TW_MR_DATA_NACK 0x58
- #define TW_ST_SLA_ACK 0xA8
- #define TW_ST_ARB_LOST_SLA_ACK 0xB0
- #define TW_ST_DATA_ACK 0xB8
- #define TW_ST_DATA_NACK 0xC0
- #define TW_ST_LAST_DATA 0xC8
- #define TW_SR_SLA_ACK 0x60
- #define TW_SR_ARB_LOST_SLA_ACK 0x68
- #define TW_SR_GCALL_ACK 0x70
- #define TW_SR_ARB_LOST_GCALL_ACK 0x78
- #define TW_SR_DATA_ACK 0x80
- #define TW_SR_DATA_NACK 0x88
- #define TW_SR_GCALL_DATA_ACK 0x90
- #define TW_SR_GCALL_DATA_NACK 0x98
- #define TW_SR_STOP 0xA0
- #define TW_NO_INFO 0xF8
- #define TW_BUS_ERROR 0x00
- #define TW_STATUS_MASK
- #define TW_STATUS (TWSR & TW_STATUS_MASK)

R/~W bit in SLA+R/W address field.

- #define TW_READ 1
- #define TW_WRITE 0

23.34.1 Detailed Description

```
#include <util/twi.h>
```

This header file contains bit mask definitions for use with the AVR TWI interface.

23.34.2 Define Documentation

23.34.2.1 #define TW_BUS_ERROR 0x00

illegal start or stop condition

23.34.2.2 #define TW_MR_ARB_LOST 0x38

arbitration lost in SLA+R or NACK

23.34.2.3 #define TW_MR_DATA_ACK 0x50

data received, ACK returned

23.34.2.4 #define TW_MR_DATA_NACK 0x58

data received, NACK returned

23.34.2.5 #define TW_MR_SLA_ACK 0x40

SLA+R transmitted, ACK received

23.34.2.6 #define TW_MR_SLA_NACK 0x48

SLA+R transmitted, NACK received

23.34.2.7 #define TW_MT_ARB_LOST 0x38

arbitration lost in SLA+W or data

23.34.2.8 #define TW_MT_DATA_ACK 0x28

data transmitted, ACK received

23.34.2.9 #define TW_MT_DATA_NACK 0x30

data transmitted, NACK received

23.34.2.10 #define TW_MT_SLA_ACK 0x18

SLA+W transmitted, ACK received

23.34.2.11 #define TW_MT_SLA_NACK 0x20

SLA+W transmitted, NACK received

23.34.2.12 #define TW_NO_INFO 0xF8

no state information available

23.34.2.13 #define TW_READ 1

SLA+R address

23.34.2.14 #define TW_REP_START 0x10

repeated start condition transmitted

23.34.2.15 #define TW_SR_ARB_LOST_GCALL_ACK 0x78

arbitration lost in SLA+RW, general call received, ACK returned

23.34.2.16 #define TW_SR_ARB_LOST_SLA_ACK 0x68

arbitration lost in SLA+RW, SLA+W received, ACK returned

23.34.2.17 #define TW_SR_DATA_ACK 0x80

data received, ACK returned

23.34.2.18 #define TW_SR_DATA_NACK 0x88

data received, NACK returned

23.34.2.19 #define TW_SR_GCALL_ACK 0x70

general call received, ACK returned

23.34.2.20 #define TW_SR_GCALL_DATA_ACK 0x90

general call data received, ACK returned

23.34.2.21 #define TW_SR_GCALL_DATA_NACK 0x98

general call data received, NACK returned

23.34.2.22 #define TW_SR_SLA_ACK 0x60

SLA+W received, ACK returned

23.34.2.23 #define TW_SR_STOP 0xA0

stop or repeated start condition received while selected

23.34.2.24 #define TW_ST_ARB_LOST_SLA_ACK 0xB0

arbitration lost in SLA+RW, SLA+R received, ACK returned

23.34.2.25 #define TW_ST_DATA_ACK 0xB8

data transmitted, ACK received

23.34.2.26 #define TW_ST_DATA_NACK 0xC0

data transmitted, NACK received

23.34.2.27 #define TW_ST_LAST_DATA 0xC8

last data byte transmitted, ACK received

23.34.2.28 #define TW_ST_SLA_ACK 0xA8

SLA+R received, ACK returned

23.34.2.29 #define TW_START 0x08

start condition transmitted

23.34.2.30 #define TW_STATUS (TWSR & TW_STATUS_MASK)

TWSR, masked by TW_STATUS_MASK

23.34.2.31 #define TW_STATUS_MASK**Value:**

```
(_BV(TWS7) | _BV(TWS6) | _BV(TWS5) | _BV(TWS4) | \
 _BV(TWS3))
```

The lower 3 bits of TWSR are reserved on the ATmega163. The 2 LSB carry the prescaler bits on the newer ATmegs.

23.34.2.32 #define TW_WRITE 0

SLA+W address

23.35 <compat/deprecated.h>: Deprecated items**Allowing specific system-wide interrupts**

In addition to globally enabling interrupts, each device's particular interrupt needs to be enabled separately if interrupts for this device are desired. While some devices maintain their interrupt enable bit inside the device's register set, external and timer interrupts have system-wide configuration registers.

Example:

```
// Enable timer 1 overflow interrupts.
timer_enable_int(_BV(TOIE1));

// Do some work...

// Disable all timer interrupts.
timer_enable_int(0);
```

Note

Be careful when you use these functions. If you already have a different interrupt enabled, you could inadvertently disable it by enabling another interrupt.

- static `__inline__ void timer_enable_int (unsigned char ints)`
- `#define enable_external_int(mask) (__EICR = mask)`
- `#define INTERRUPT(signame)`
- `#define __INTR_ATTRS used`

Obsolete IO macros

Back in a time when AVR-GCC and avr-libc could not handle IO port access in the direct assignment form as they are handled now, all IO port access had to be done through specific macros that eventually resulted in inline assembly instructions performing the desired action.

These macros became obsolete, as reading and writing IO ports can be done by simply using the IO port name in an expression, and all bit manipulation (including those on IO ports) can be done using generic C bit manipulation operators.

The macros in this group simulate the historical behaviour. While they are supposed to be applied to IO ports, the emulation actually uses standard C methods, so they could be applied to arbitrary memory locations as well.

- `#define inp(port) (port)`
- `#define outp(val, port) (port) = (val)`
- `#define inb(port) (port)`
- `#define outb(port, val) (port) = (val)`
- `#define sbi(port, bit) (port) |= (1 << (bit))`
- `#define cbi(port, bit) (port) &= ~(1 << (bit))`

23.35.1 Detailed Description

This header file contains several items that used to be available in previous versions of this library, but have eventually been deprecated over time.

```
#include <compat/deprecated.h>
```

These items are supplied within that header file for backward compatibility reasons only, so old source code that has been written for previous library versions could easily be maintained until its end-of-life. Use of any of these items in new code is strongly discouraged.

23.35.2 Define Documentation

23.35.2.1 `#define cbi(port, bit) (port) &= ~(1 << (bit))`

Deprecated

Clear `bit` in IO port `port`.

23.35.2.2 `#define enable_external_int(mask) (__EICR = mask)`

Deprecated

This macro gives access to the `GIMSK` register (or `EIMSK` register if using an AVR Mega device or `GICR` register for others). Although this macro is essentially the same as assigning to the register, it does adapt slightly to the type of device being used. This macro is unavailable if none of the registers listed above are defined.

23.35.2.3 #define inb(port) (port)

Deprecated

Read a value from an IO port `port`.

23.35.2.4 #define inp(port) (port)

Deprecated

Read a value from an IO port `port`.

23.35.2.5 #define INTERRUPT(signame)

Value:

```
void signame (void) __attribute__((interrupt, __INTR_ATTRS)); \
void signame (void)
```

Deprecated

Introduces an interrupt handler function that runs with global interrupts initially enabled. This allows interrupt handlers to be interrupted.

As this macro has been used by too many unsuspecting people in the past, it has been deprecated, and will be removed in a future version of the library. Users who want to legitimately re-enable interrupts in their interrupt handlers as quickly as possible are encouraged to explicitly declare their handlers as described [above](#).

23.35.2.6 #define outb(port, val) (port) = (val)

Deprecated

Write `val` to IO port `port`.

23.35.2.7 `#define outp(val, port) (port) = (val)`

Deprecated

Write `val` to IO port `port`.

23.35.2.8 `#define sbi(port, bit) (port) |= (1 << (bit))`

Deprecated

Set `bit` in IO port `port`.

23.35.3 Function Documentation

23.35.3.1 `static __inline__ void timer_enable_int (unsigned char ints)`
`[static]`

Deprecated

This function modifies the `timsk` register. The value you pass via `ints` is device specific.

23.36 <compat/ina90.h>: Compatibility with IAR EWB 3.x

```
#include <compat/ina90.h>
```

This is an attempt to provide some compatibility with header files that come with IAR C, to make porting applications between different compilers easier. No 100% compatibility though.

Note

For actual documentation, please see the IAR manual.

23.37 Demo projects

Modules

- [Combining C and assembly source files](#)
- [A simple project](#)
- [A more sophisticated project](#)
- [Using the standard IO facilities](#)
- [Example using the two-wire interface \(TWI\)](#)

23.37.1 Detailed Description

Various small demo projects are provided to illustrate several aspects of using the open-source utilities for the AVR controller series. It should be kept in mind that these demos serve mainly educational purposes, and are normally not directly suitable for use in any production environment. Usually, they have been kept as simple as sufficient to demonstrate one particular feature.

The [simple project](#) is somewhat like the "Hello world!" application for a microcontroller, about the most simple project that can be done. It is explained in good detail, to allow the reader to understand the basic concepts behind using the tools on an AVR microcontroller.

The [more sophisticated demo project](#) builds on top of that simple project, and adds some controls to it. It touches a number of avr-libc's basic concepts on its way.

A [comprehensive example on using the standard IO facilities](#) intends to explain that complex topic, using a practical microcontroller peripheral setup with one RS-232 connection, and an HD44780-compatible industry-standard LCD display.

The [Example using the two-wire interface \(TWI\)](#) project explains the use of the two-wire hardware interface (also known as "I2C") that is present on many AVR controllers.

Finally, the [Combining C and assembly source files](#) demo shows how C and assembly language source files can collaborate within one project. While the overall project is managed by a C program part for easy maintenance, time-critical parts are written directly in manually optimized assembly language for shortest execution times possible. Naturally, this kind of project is very closely tied to the hardware design, thus it is custom-tailored to a particular controller type and peripheral setup. As an alternative to the assembly-language solution, this project also offers a C-only implementation (deploying the exact same peripheral setup) based on a more sophisticated (and thus more expensive) but pin-compatible controller.

While the simple demo is meant to run on about any AVR setup possible where a LED could be connected to the OCR1[A] output, the [large](#) and [stdio](#) demos are mainly targeted to the Atmel STK500 starter kit, and the [TWI](#) example requires a controller where some 24Cxx two-wire EEPROM can be connected to. For the STK500 demos, the default CPU (either an AT90S8515 or an ATmega8515) should be removed from its socket, and the ATmega16 that ships with the kit should be inserted into socket SCKT3100A3. The ATmega16 offers an on-board ADC that is used in the [large](#) demo, and all AVRs with an ADC feature a different pinout than the industry-standard compatible devices.

In order to fully utilize the [large](#) demo, a female 10-pin header with cable, connecting to a 10 kOhm potentiometer will be useful.

For the [stdio](#) demo, an industry-standard HD44780-compatible LCD display of at least 16x1 characters will be needed. Among other things, the [LCD4Linux](#) project page describes many things around these displays, including common pinouts.

23.38 Combining C and assembly source files

For time- or space-critical applications, it can often be desirable to combine C code (for easy maintenance) and assembly code (for maximal speed or minimal code size) together. This demo provides an example of how to do that.

The objective of the demo is to decode radio-controlled model PWM signals, and control an output PWM based on the current input signal's value. The incoming PWM pulses follow a standard encoding scheme where a pulse width of 920 microseconds denotes one end of the scale (represented as 0 % pulse width on output), and 2120 microseconds mark the other end (100 % output PWM). Normally, multiple channels would be encoded that way in subsequent pulses, followed by a larger gap, so the entire frame will repeat each 14 through 20 ms, but this is ignored for the purpose of the demo, so only a single input PWM channel is assumed.

The basic challenge is to use the cheapest controller available for the task, an ATtiny13 that has only a single timer channel. As this timer channel is required to run the outgoing PWM signal generation, the incoming PWM decoding had to be adjusted to the constraints set by the outgoing PWM.

As PWM generation toggles the counting direction of timer 0 between up and down after each 256 timer cycles, the current time cannot be deduced by reading TCNT0 only, but the current counting direction of the timer needs to be considered as well. This requires servicing interrupts whenever the timer hits *TOP* (255) and *BOTTOM* (0) to learn about each change of the counting direction. For PWM generation, it is usually desired to run it at the highest possible speed so filtering the PWM frequency from the modulated output signal is made easy. Thus, the PWM timer runs at full CPU speed. This causes the overflow and compare match interrupts to be triggered each 256 CPU clocks, so they must run with the minimal number of processor cycles possible in order to not impose a too high CPU load by these interrupt service routines. This is the main reason to implement the entire interrupt handling in fine-tuned assembly code rather than in C.

In order to verify parts of the algorithm, and the underlying hardware, the demo has been set up in a way so the pin-compatible but more expensive ATtiny45 (or its siblings ATtiny25 and ATtiny85) could be used as well. In that case, no separate assembly code is required, as two timer channels are available.

23.38.1 Hardware setup

The incoming PWM pulse train is fed into PB4. It will generate a pin change interrupt there on each edge of the incoming signal.

The outgoing PWM is generated through OC0B of timer channel 0 (PB1). For demon-

stration purposes, a LED should be connected to that pin (like, one of the LEDs of an STK500).

The controllers run on their internal calibrated RC oscillators, 1.2 MHz on the ATtiny13, and 1.0 MHz on the ATtiny45.

23.38.2 A code walkthrough

23.38.2.1 `asmdemo.c`

After the usual include files, two variables are defined. The first one, `pwm_incoming` is used to communicate the most recent pulse width detected by the incoming PWM decoder up to the main loop.

The second variable actually only constitutes of a single bit, `intbits.pwm_received`. This bit will be set whenever the incoming PWM decoder has updated `pwm_incoming`.

Both variables are marked *volatile* to ensure their readers will always pick up an updated value, as both variables will be set by interrupt service routines.

The function `ioinit()` initializes the microcontroller peripheral devices. In particular, it starts timer 0 to generate the outgoing PWM signal on OC0B. Setting OCR0A to 255 (which is the *TOP* value of timer 0) is used to generate a timer 0 overflow A interrupt on the ATtiny13. This interrupt is used to inform the incoming PWM decoder that the counting direction of channel 0 is just changing from up to down. Likewise, an overflow interrupt will be generated whenever the countdown reached *BOTTOM* (value 0), where the counter will again alter its counting direction to upwards. This information is needed in order to know whether the current counter value of TCNT0 is to be evaluated from bottom or top.

Further, `ioinit()` activates the pin-change interrupt PCINT0 on any edge of PB4. Finally, PB1 (OC0B) will be activated as an output pin, and global interrupts are being enabled.

In the ATtiny45 setup, the C code contains an ISR for PCINT0. At each pin-change interrupt, it will first be analyzed whether the interrupt was caused by a rising or a falling edge. In case of the rising edge, timer 1 will be started with a prescaler of 16 after clearing the current timer value. Then, at the falling edge, the current timer value will be recorded (and timer 1 stopped), the pin-change interrupt will be suspended, and the upper layer will be notified that the incoming PWM measurement data is available.

Function `main()` first initializes the hardware by calling `ioinit()`, and then waits until some incoming PWM value is available. If it is, the output PWM will be adjusted by computing the relative value of the incoming PWM. Finally, the pin-change interrupt is re-enabled, and the CPU is put to sleep.

23.38.2.2 `project.h`

In order for the interrupt service routines to be as fast as possible, some of the CPU registers are set aside completely for use by these routines, so the compiler would not use them for C code. This is arranged for in `project.h`.

The file is divided into one section that will be used by the assembly source code, and another one to be used by C code. The assembly part is distinguished by the preprocessing macro `__ASSEMBLER__` (which will be automatically set by the compiler front-end when preprocessing an assembly-language file), and it contains just macros that give symbolic names to a number of CPU registers. The preprocessor will then replace the symbolic names by their right-hand side definitions before calling the assembler.

In C code, the compiler needs to see variable declarations for these objects. This is done by using declarations that bind a variable permanently to a CPU register (see [How to permanently bind a variable to a register?](#)). Even in case the C code never has a need to access these variables, declaring the register binding that way causes the compiler to not use these registers in C code at all.

The `flags` variable needs to be in the range of r16 through r31 as it is the target of a *load immediate* (or *SER*) instruction that is not applicable to the entire register file.

23.38.2.3 `isrs.S`

This file is a preprocessed assembly source file. The C preprocessor will be run by the compiler front-end first, resolving all `#include`, `#define` etc. directives. The resulting program text will then be passed on to the assembler.

As the C preprocessor strips all C-style comments, preprocessed assembly source files can have both, C-style (`/* ... */`, `// ...`) as well as assembly-style (`;` ...) comments.

At the top, the IO register definition file `avr/io.h` and the project declaration file `project.h` are included. The remainder of the file is conditionally assembled only if the target MCU type is an ATtiny13, so it will be completely ignored for the ATtiny45 option.

Next are the two interrupt service routines for timer 0 compare A match (timer 0 hits *TOP*, as OCR0A is set to 255) and timer 0 overflow (timer 0 hits *BOTTOM*). As discussed above, these are kept as short as possible. They only save `SREG` (as the flags will be modified by the `INC` instruction), increment the `counter_hi` variable which forms the high part of the current time counter (the low part is formed by querying `TCNT0` directly), and clear or set the variable `flags`, respectively, in order to note the current counting direction. The `RETI` instruction terminates these interrupt service routines. Total cycle count is 8 CPU cycles, so together with the 4 CPU cycles needed for interrupt setup, and the 2 cycles for the `RJMP` from the interrupt vector to the handler, these routines will require 14 out of each 256 CPU cycles, or about 5 % of the overall CPU time.

The pin-change interrupt `PCINT0` will be handled in the final part of this file. The basic algorithm is to quickly evaluate the current system time by fetching the current timer value of `TCNT0`, and combining it with the overflow part in `counter_hi`. If the counter is currently counting down rather than up, the value fetched from `TCNT0` must be negated. Finally, if this pin-change interrupt was triggered by a rising edge, the time computed will be recorded as the start time only. Then, at the falling edge, this start time will be subtracted from the current time to compute the actual pulse width seen (left in `pwm_incoming`), and the upper layers are informed of the new value by

setting bit 0 in the `intbits` flags. At the same time, this pin-change interrupt will be disabled so no new measurement can be performed until the upper layer had a chance to process the current value.

23.38.3 The source code

The source code is installed under

```
$prefix/share/doc/avr-libc/examples/asmdemo/,
```

where `$prefix` is a configuration option. For Unix systems, it is usually set to either `/usr` or `/usr/local`.

23.39 A simple project

At this point, you should have the GNU tools configured, built, and installed on your system. In this chapter, we present a simple example of using the GNU tools in an AVR project. After reading this chapter, you should have a better feel as to how the tools are used and how a `Makefile` can be configured.

23.39.1 The Project

This project will use the pulse-width modulator (PWM) to ramp an LED on and off every two seconds. An AT90S2313 processor will be used as the controller. The circuit for this demonstration is shown in the [schematic diagram](#). If you have a development kit, you should be able to use it, rather than build the circuit, for this project.

Note

Meanwhile, the AT90S2313 became obsolete. Either use its successor, the (pin-compatible) ATtiny2313 for the project, or perhaps the ATmega8 or one of its successors (ATmega48/88/168) which have become quite popular since the original demo project had been established. For all these more modern devices, it is no longer necessary to use an external crystal for clocking as they ship with the internal 1 MHz oscillator enabled, so C1, C2, and Q1 can be omitted. Normally, for this experiment, the external circuitry on /RESET (R1, C3) can be omitted as well, leaving only the AVR, the LED, the bypass capacitor C4, and perhaps R2. For the ATmega8/48/88/168, use PB1 (pin 15 at the DIP-28 package) to connect the LED to. Additionally, this demo has been ported to many different other AVRs. The location of the respective OC pin varies between different AVRs, and it is mandated by the AVR hardware.

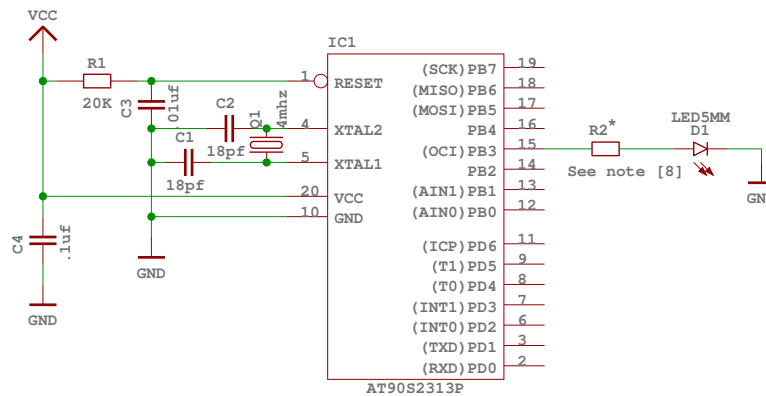


Figure 5: Schematic of circuit for demo project

The source code is given in [demo.c](#). For the sake of this example, create a file called `demo.c` containing this source code. Some of the more important parts of the code are:

Note [1]:

As the AVR microcontroller series has been developed during the past years, new features have been added over time. Even though the basic concepts of the timer/counter1 are still the same as they used to be back in early 2001 when this simple demo was written initially, the names of registers and bits have been changed slightly to reflect the new features. Also, the port and pin mapping of the output compare match 1A (or 1 for older devices) pin which is used to control the LED varies between different AVRs. The file `iocompat.h` tries to abstract between all this differences using some preprocessor `#ifdef` statements, so the actual program itself can operate on a common set of symbolic names. The macros defined by that file are:

- `OCR` the name of the OCR register used to control the PWM (usually either `OCR1` or `OCR1A`)
- `DDROC` the name of the DDR (data direction register) for the OC output
- `OC1` the pin number of the OC1[A] output within its port
- `TIMER1_TOP` the TOP value of the timer used for the PWM (1023 for 10-bit PWMs, 255 for devices that can only handle an 8-bit PWM)
- `TIMER1_PWM_INIT` the initialization bits to be set into control register 1A in order to setup 10-bit (or 8-bit) phase and frequency correct PWM mode
- `TIMER1_CLOCKSOURCE` the clock bits to set in the respective control register to start the PWM timer; usually the timer runs at full CPU clock for 10-bit PWMs, while it runs on a prescaled clock for 8-bit PWMs

Note [2]:

`ISR()` is a macro that marks the function as an interrupt routine. In this case, the function will get called when timer 1 overflows. Setting up interrupts is explained in greater detail in [<avr/interrupt.h>: Interrupts](#).

Note [3]:

The `PWM` is being used in 10-bit mode, so we need a 16-bit variable to remember the current value.

Note [4]:

This section determines the new value of the `PWM`.

Note [5]:

Here's where the newly computed value is loaded into the `PWM` register. Since we are in an interrupt routine, it is safe to use a 16-bit assignment to the register. Outside of an interrupt, the assignment should only be performed with interrupts disabled if there's a chance that an interrupt routine could also access this register (or another register that uses `TEMP`), see the appropriate [FAQ entry](#).

Note [6]:

This routine gets called after a reset. It initializes the `PWM` and enables interrupts.

Note [7]:

The main loop of the program does nothing -- all the work is done by the interrupt routine! The `sleep_mode()` puts the processor on sleep until the next interrupt, to conserve power. Of course, that probably won't be noticeable as we are still driving a LED, it is merely mentioned here to demonstrate the basic principle.

Note [8]:

Early AVR devices saturate their outputs at rather low currents when sourcing current, so the LED can be connected directly, the resulting current through the LED will be about 15 mA. For modern parts (at least for the ATmega 128), however Atmel has drastically increased the IO source capability, so when operating at 5 V `Vcc`, `R2` is needed. Its value should be about 150 Ohms. When operating the circuit at 3 V, it can still be omitted though.

23.39.2 The Source Code

```
/*
 * -----
 * "THE BEER-WARE LICENSE" (Revision 42):
 * <joerg@FreeBSD.ORG> wrote this file. As long as you retain this notice you
 * can do whatever you want with this stuff. If we meet some day, and you think
 * this stuff is worth it, you can buy me a beer in return.      Joerg Wunsch
 * -----
 *
 */
```

```

* Simple AVR demonstration. Controls a LED that can be directly
* connected from OC1/OC1A to GND. The brightness of the LED is
* controlled with the PWM. After each period of the PWM, the PWM
* value is either incremented or decremented, that's all.
*
* $Id$
*/

#include <inttypes.h>
#include <avr/io.h>
#include <avr/interrupt.h>
#include <avr/sleep.h>

#include "iocompat.h"      /* Note [1] */

enum { UP, DOWN };

ISR (TIMER1_OVF_vect)      /* Note [2] */
{
    static uint16_t pwm;    /* Note [3] */
    static uint8_t direction;

    switch (direction)      /* Note [4] */
    {
        case UP:
            if (++pwm == TIMER1_TOP)
                direction = DOWN;
            break;

        case DOWN:
            if (--pwm == 0)
                direction = UP;
            break;
    }

    OCR = pwm;              /* Note [5] */
}

void
iointit (void)             /* Note [6] */
{
    /* Timer 1 is 10-bit PWM (8-bit PWM on some ATtinys). */
    TCCR1A = TIMER1_PWM_INIT;
    /*
     * Start timer 1.
     *
     * NB: TCCR1A and TCCR1B could actually be the same register, so
     * take care to not clobber it.
     */
    TCCR1B |= TIMER1_CLOCKSOURCE;
    /*
     * Run any device-dependent timer 1 setup hook if present.
     */
    #if defined(TIMER1_SETUP_HOOK)
        TIMER1_SETUP_HOOK();
    #endif

    /* Set PWM value to 0. */
    OCR = 0;

    /* Enable OC1 as output. */
    DDROC = _BV (OC1);
}

```



```
/* Enable timer 1 overflow interrupt. */
TIMSK = _BV (TOIE1);
sei ();
}

int
main (void)
{

    ioinit ();

    /* loop forever, the interrupts are doing the rest */

    for (;;)          /* Note [7] */
        sleep_mode();

    return (0);
}
```

23.39.3 Compiling and Linking

This first thing that needs to be done is compile the source. When compiling, the compiler needs to know the processor type so the `-mmcu` option is specified. The `-Os` option will tell the compiler to optimize the code for efficient space usage (at the possible expense of code execution speed). The `-g` is used to embed debug info. The debug info is useful for disassemblies and doesn't end up in the `.hex` files, so I usually specify it. Finally, the `-c` tells the compiler to compile and stop -- don't link. This demo is small enough that we could compile and link in one step. However, real-world projects will have several modules and will typically need to break up the building of the project into several compiles and one link.

```
$ avr-gcc -g -Os -mmcu=atmega8 -c demo.c
```

The compilation will create a `demo.o` file. Next we link it into a binary called `demo.elf`.

```
$ avr-gcc -g -mmcu=atmega8 -o demo.elf demo.o
```

It is important to specify the MCU type when linking. The compiler uses the `-mmcu` option to choose start-up files and run-time libraries that get linked together. If this option isn't specified, the compiler defaults to the 8515 processor environment, which is most certainly what you didn't want.

23.39.4 Examining the Object File

Now we have a binary file. Can we do anything useful with it (besides put it into the processor?) The GNU Binutils suite is made up of many useful tools for manipulating object files that get generated. One tool is `avr-objdump`, which takes information from the object file and displays it in many useful ways. Typing the command by itself will cause it to list out its options.

For instance, to get a feel of the application's size, the `-h` option can be used. The output of this option shows how much space is used in each of the sections (the `.stab` and `.stabstr` sections hold the debugging information and won't make it into the ROM file).

An even more useful option is `-S`. This option disassembles the binary file and intersperses the source code in the output! This method is much better, in my opinion, than using the `-S` with the compiler because this listing includes routines from the libraries and the vector table contents. Also, all the "fix-ups" have been satisfied. In other words, the listing generated by this option reflects the actual code that the processor will run.

```
$ avr-objdump -h -S demo.elf > demo.lst
```

Here's the output as saved in the `demo.lst` file:

```
demo.elf:      file format elf32-avr

Sections:
Idx Name          Size      VMA           LMA           File off  Algn
  0 .text          000000be  00000000  00000000  00000094  2**1
    CONTENTS, ALLOC, LOAD, READONLY, CODE
  1 .data          00000000  00800060  000000be  00000152  2**0
    CONTENTS, ALLOC, LOAD, DATA
  2 .bss           00000003  00800060  00800060  00000152  2**0
    ALLOC
  3 .comment        0000002c  00000000  00000000  00000152  2**0
    CONTENTS, READONLY
  4 .debug_aranges  00000068  00000000  00000000  00000180  2**3
    CONTENTS, READONLY, DEBUGGING
  5 .debug_info     000002d2  00000000  00000000  000001e8  2**0
    CONTENTS, READONLY, DEBUGGING
  6 .debug_abbrev   00000108  00000000  00000000  000004ba  2**0
    CONTENTS, READONLY, DEBUGGING
  7 .debug_line     00000250  00000000  00000000  000005c2  2**0
    CONTENTS, READONLY, DEBUGGING
  8 .debug_frame    00000060  00000000  00000000  00000814  2**2
    CONTENTS, READONLY, DEBUGGING
  9 .debug_str      000000f8  00000000  00000000  00000874  2**0
    CONTENTS, READONLY, DEBUGGING
10 .debug_loc      00000056  00000000  00000000  0000096c  2**0
    CONTENTS, READONLY, DEBUGGING
11 .debug_ranges   00000018  00000000  00000000  000009c2  2**0
    CONTENTS, READONLY, DEBUGGING

Disassembly of section .text:

00000000 <__ctors_end>:
/* __do_clear_bss is only necessary if there is anything in .bss section. */

#ifdef L_clear_bss
.section .init4,"ax",@progbits
DEFUN __do_clear_bss
ldi r18, hi8(__bss_end)
    0: 20 e0      ldi r18, 0x00 ; 0
ldi r26, lo8(__bss_start)
    2: a0 e6      ldi r26, 0x60 ; 96
ldi r27, hi8(__bss_start)
    4: b0 e0      ldi r27, 0x00 ; 0
rjmp .do_clear_bss_start
```

```

        6: 01 c0          rjmp .+2          ; 0xa <.do_clear_bss_start>

00000008 <.do_clear_bss_loop>:
.do_clear_bss_loop:
st X+, __zero_reg__
    8: 1d 92          st X+, r1

0000000a <.do_clear_bss_start>:
.do_clear_bss_start:
cpi r26, lo8(__bss_end)
    a: a3 36          cpi r26, 0x63 ; 99
cpc r27, r18
    c: b2 07          cpc r27, r18
brne .do_clear_bss_loop
    e: e1 f7          brne .-8          ; 0x8 <.do_clear_bss_loop>

00000010 <__vector_8>:
#include "iocompat.h" /* Note [1] */

enum { UP, DOWN };

ISR (TIMER1_OVF_vect) /* Note [2] */
{
    10: 1f 92          push r1
    12: 0f 92          push r0
    14: 0f b6          in r0, 0x3f ; 63
    16: 0f 92          push r0
    18: 11 24          eor r1, r1
    1a: 2f 93          push r18
    1c: 8f 93          push r24
    1e: 9f 93          push r25
    20: 80 91 60 00    lds r24, 0x0060 ; 0x800060 <__DATA_REGION_ORIGIN__>
    24: 90 91 61 00    lds r25, 0x0061 ; 0x800061 <__DATA_REGION_ORIGIN__+0x1>
        static uint16_t pwm; /* Note [3] */
        static uint8_t direction;

        switch (direction) /* Note [4] */
        28: 20 91 62 00    lds r18, 0x0062 ; 0x800062 <direction.1505>
        2c: 22 23          and r18, r18
        2e: a9 f0          breq .+42          ; 0x5a <__SREG__+0x1b>
        30: 21 30          cpi r18, 0x01 ; 1
        32: 49 f4          brne .+18          ; 0x46 <__SREG__+0x7>
            if (++pwm == TIMER1_TOP)
                direction = DOWN;
            break;

        case DOWN:
            if (--pwm == 0)
        34: 01 97          sbiw r24, 0x01 ; 1
        36: 90 93 61 00    sts 0x0061, r25 ; 0x800061 <__DATA_REGION_ORIGIN__+0x1>
        3a: 80 93 60 00    sts 0x0060, r24 ; 0x800060 <__DATA_REGION_ORIGIN__>
        3e: 00 97          sbiw r24, 0x00 ; 0
        40: 11 f4          brne .+4          ; 0x46 <__SREG__+0x7>
                direction = UP;
        42: 10 92 62 00    sts 0x0062, r1 ; 0x800062 <direction.1505>
                break;
            }

        OCR = pwm; /* Note [5] */
        46: 9b bd          out 0x2b, r25 ; 43
        48: 8a bd          out 0x2a, r24 ; 42
}

```

```

4a: 9f 91      pop r25
4c: 8f 91      pop r24
4e: 2f 91      pop r18
50: 0f 90      pop r0
52: 0f be      out 0x3f, r0 ; 63
54: 0f 90      pop r0
56: 1f 90      pop r1
58: 18 95      reti
    static uint8_t direction;

    switch (direction) /* Note [4] */
    {
        case UP:
            if (++pwm == TIMER1_TOP)
2a: 01 96      adiw r24, 0x01 ; 1
5c: 90 93 61 00 sts 0x0061, r25 ; 0x800061 <__DATA_REGION_ORIGIN__+0x1>
60: 80 93 60 00 sts 0x0060, r24 ; 0x800060 <__DATA_REGION_ORIGIN__>
64: 8f 3f      cpi r24, 0xFF ; 255
66: 23 e0      ldi r18, 0x03 ; 3
68: 92 07      cpc r25, r18
6a: 69 f7      brne .-38 ; 0x46 <__SREG__+0x7>
                direction = DOWN;
6c: 21 e0      ldi r18, 0x01 ; 1
6e: 20 93 62 00 sts 0x0062, r18 ; 0x800062 <direction.1505>
72: e9 cf      rjmp .-46 ; 0x46 <__SREG__+0x7>

00000074 <ioint>:

void
ioint (void) /* Note [6] */
{
    /* Timer 1 is 10-bit PWM (8-bit PWM on some ATtinys). */
    TCCR1A = TIMER1_PWM_INIT;
74: 83 e8      ldi r24, 0x83 ; 131
76: 8f bd      out 0x2f, r24 ; 47
    * Start timer 1.
    *
    * NB: TCCR1A and TCCR1B could actually be the same register, so
    * take care to not clobber it.
    */
    TCCR1B |= TIMER1_CLOCKSOURCE;
78: 8e b5      in r24, 0x2e ; 46
7a: 81 60      ori r24, 0x01 ; 1
7c: 8e bd      out 0x2e, r24 ; 46
#ifdef TIMER1_SETUP_HOOK
    TIMER1_SETUP_HOOK();
#endif

    /* Set PWM value to 0. */
    OCR = 0;
7e: 1b bc      out 0x2b, r1 ; 43
80: 1a bc      out 0x2a, r1 ; 42

    /* Enable OC1 as output. */
    DDROC = _BV (OC1);
82: 82 e0      ldi r24, 0x02 ; 2
84: 87 bb      out 0x17, r24 ; 23

    /* Enable timer 1 overflow interrupt. */
    TIMSK = _BV (TOIE1);
86: 84 e0      ldi r24, 0x04 ; 4
88: 89 bf      out 0x39, r24 ; 57

```

```

    sei ();
8a: 78 94      sei
}
8c: 08 95      ret

0000008e <main>:

void
ioinit (void) /* Note [6] */
{
    /* Timer 1 is 10-bit PWM (8-bit PWM on some ATtinys). */
    TCCR1A = TIMER1_PWM_INIT;
8e: 83 e8      ldi r24, 0x83 ; 131
90: 8f bd      out 0x2f, r24 ; 47
    * Start timer 1.
    *
    * NB: TCCR1A and TCCR1B could actually be the same register, so
    * take care to not clobber it.
    */
    TCCR1B |= TIMER1_CLOCKSOURCE;
92: 8e b5      in r24, 0x2e ; 46
94: 81 60      ori r24, 0x01 ; 1
96: 8e bd      out 0x2e, r24 ; 46
#ifdef(TIMER1_SETUP_HOOK)
    TIMER1_SETUP_HOOK();
#endif

    /* Set PWM value to 0. */
    OCR = 0;
98: 1b bc      out 0x2b, r1 ; 43
9a: 1a bc      out 0x2a, r1 ; 42

    /* Enable OC1 as output. */
    DDROC = _BV (OC1);
9c: 82 e0      ldi r24, 0x02 ; 2
9e: 87 bb      out 0x17, r24 ; 23

    /* Enable timer 1 overflow interrupt. */
    TIMSK = _BV (TOIE1);
a0: 84 e0      ldi r24, 0x04 ; 4
a2: 89 bf      out 0x39, r24 ; 57
    sei ();
a4: 78 94      sei
    ioinit ();

    /* loop forever, the interrupts are doing the rest */

    for (;;) /* Note [7] */
        sleep_mode();
a6: 85 b7      in r24, 0x35 ; 53
a8: 80 68      ori r24, 0x80 ; 128
aa: 85 bf      out 0x35, r24 ; 53
ac: 88 95      sleep
ae: 85 b7      in r24, 0x35 ; 53
b0: 8f 77      andi r24, 0x7F ; 127
b2: 85 bf      out 0x35, r24 ; 53
b4: f8 cf      rjmp .-16 ; 0xa6 <main+0x18>

000000b6 <exit>:
b6: f8 94      cli
b8: 00 c0      rjmp .+0 ; 0xba <_exit>

```

```

000000ba <_exit>:
ENDF _exit

/* Code from .fini8 ... .fini1 sections inserted by ld script. */

.section .fini0,"ax",@progbits
cli
    ba: f8 94          cli

000000bc <__stop_program>:
__stop_program:
rjmp __stop_program
    bc: ff cf          rjmp .-2          ; 0xbc <__stop_program>

```

23.39.5 Linker Map Files

`avr-objdump` is very useful, but sometimes it's necessary to see information about the link that can only be generated by the linker. A map file contains this information. A map file is useful for monitoring the sizes of your code and data. It also shows where modules are loaded and which modules were loaded from libraries. It is yet another view of your application. To get a map file, I usually add **-Wl, -Map, demo.map** to my link command. Relink the application using the following command to generate `demo.map` (a portion of which is shown below).

```
$ avr-gcc -g -mmcu=atmega8 -Wl,-Map,demo.map -o demo.elf demo.o
```

Some points of interest in the `demo.map` file are:

```

.rela.plt
*(.rela.plt)

.text          0x0000000000000000          0xbe
*(.vectors)
*(.vectors)
*(.progmem.gcc*)
                0x0000000000000000          . = ALIGN (0x2)
                0x0000000000000000          __trampolines_start = .
*(.trampolines)
.trampolines   0x0000000000000000          0x0 linker stubs
*(.trampolines*)
                0x0000000000000000          __trampolines_end = .
*libprintf_flt.a:*(.progmem.data)
*libc.a:*(.progmem.data)
*(.progmem*)
                0x0000000000000000          . = ALIGN (0x2)
*(.jumptables)
*(.jumptables*)
*(.lowtext)
*(.lowtext*)
                0x0000000000000000          __ctors_start = .

```

The `.text` segment (where program instructions are stored) starts at location `0x0`.

```

*(.fini2)
*(.fini2)

```

```

*(.fini1)
*(.fini1)
*(.fini0)
.fini0      0x00000000000000ba      0x4 /home/toolsbuild/workspace/avr8-gnu
-toolchain/avr8-gnu-toolchain-linux_x86_64/lib/gcc/avr/7.3.0/avr4/libgcc.a(_exit.
o)
*(.fini0)
      0x00000000000000be      _etext = .

.data      0x0000000000800060      0x0 load address 0x00000000000000be
[!provide]      PROVIDE (__data_start, .)

*(.data)
.data      0x0000000000800060      0x0 demo.o
.data      0x0000000000800060      0x0 /home/toolsbuild/workspace/avr8-gnu
-toolchain/src/avr-libc/avr/lib/avr4/exit.o
.data      0x0000000000800060      0x0 /home/toolsbuild/workspace/avr8-gnu
-toolchain/avr8-gnu-toolchain-linux_x86_64/lib/gcc/avr/7.3.0/avr4/libgcc.a(_exit.
o)
.data      0x0000000000800060      0x0 /home/toolsbuild/workspace/avr8-gnu
-toolchain/avr8-gnu-toolchain-linux_x86_64/lib/gcc/avr/7.3.0/avr4/libgcc.a(_clear
_bss.o)
*(.data*)
*(.gnu.linkonce.d*)
*(.rodata)
*(.rodata*)
*(.gnu.linkonce.r*)
      0x0000000000800060      . = ALIGN (0x2)
      0x0000000000800060      _edata = .
[!provide]      PROVIDE (__data_end, .)

.bss      0x0000000000800060      0x3
      0x0000000000800060      PROVIDE (__bss_start, .)

*(.bss)
.bss      0x0000000000800060      0x3 demo.o
.bss      0x0000000000800063      0x0 /home/toolsbuild/workspace/avr8-gnu
-toolchain/src/avr-libc/avr/lib/avr4/exit.o
.bss      0x0000000000800063      0x0 /home/toolsbuild/workspace/avr8-gnu
-toolchain/avr8-gnu-toolchain-linux_x86_64/lib/gcc/avr/7.3.0/avr4/libgcc.a(_exit.
o)
.bss      0x0000000000800063      0x0 /home/toolsbuild/workspace/avr8-gnu
-toolchain/avr8-gnu-toolchain-linux_x86_64/lib/gcc/avr/7.3.0/avr4/libgcc.a(_clear
_bss.o)
*(.bss*)
*(COMMON)
      0x0000000000800063      PROVIDE (__bss_end, .)
      0x00000000000000be      __data_load_start = LOADADDR (.
data)
      0x00000000000000be      __data_load_end = (__data_load_
start + SIZEOF (.data))

.noinit      0x0000000000800063      0x0
[!provide]      PROVIDE (__noinit_start, .)

*(.noinit*)
[!provide]      PROVIDE (__noinit_end, .)
      0x0000000000800063      _end = .
[!provide]      PROVIDE (__heap_start, .)

.eeprom      0x0000000000810000      0x0
*(.eeprom*)
      0x0000000000810000      __eeprom_end = .

```

The last address in the `.text` segment is location `0x114` (denoted by `_etext`), so the instructions use up 276 bytes of FLASH.

The `.data` segment (where initialized static variables are stored) starts at location `0x60`, which is the first address after the register bank on an ATmega8 processor.

The next available address in the `.data` segment is also location `0x60`, so the application has no initialized data.

The `.bss` segment (where uninitialized data is stored) starts at location `0x60`.

The next available address in the `.bss` segment is location `0x63`, so the application uses 3 bytes of uninitialized data.

The `.eeprom` segment (where EEPROM variables are stored) starts at location `0x0`.

The next available address in the `.eeprom` segment is also location `0x0`, so there aren't any EEPROM variables.

23.39.6 Generating Intel Hex Files

We have a binary of the application, but how do we get it into the processor? Most (if not all) programmers will not accept a GNU executable as an input file, so we need to do a little more processing. The next step is to extract portions of the binary and save the information into `.hex` files. The GNU utility that does this is called `avr-objcopy`.

The ROM contents can be pulled from our project's binary and put into the file `demo.hex` using the following command:

```
$ avr-objcopy -j .text -j .data -O ihex demo.elf demo.hex
```

The resulting `demo.hex` file contains:

```
:1000000020E0A0E6B0E001C01D92A336B207E1F700
:100010001F920F920FB60F9211242F938F939F93DD
:100020008091600090916100209162002223A9F0EC
:10003000213049F40197909361008093600000970C
:1000400011F4109262009BBD8ABD9F918F912F91F8
:100050000F900FBE0F901F9018950196909361001E
:10006000809360008F3F23E0920769F721E020939F
:100070006200E9CF83E88FBD8EB581608EBD1BBC69
:100080001ABC82E087BB84E089BF7894089583E836
:100090008FBD8EB581608EBD1BBC1ABC82E087BB54
:1000A00084E089BF789485B7806885BF889585B7D7
:0E00B0008F7785BFF8CFF89400C0F894FFCF8B
:00000001FF
```

The `-j` option indicates that we want the information from the `.text` and `.data` segment extracted. If we specify the EEPROM segment, we can generate a `.hex` file that can be used to program the EEPROM:

```
$ avr-objcopy -j .eeprom --change-section-lma .eeprom=0 -O ihex demo.elf demo_eeprom.hex
```

There is no `demo_eeprom.hex` file written, as that file would be empty.

Starting with version 2.17 of the GNU binutils, the `avr-objcopy` command that used to generate the empty EEPROM files now aborts because of the empty input section `.eeprom`, so these empty files are not generated. It also signals an error to the Makefile which will be caught there, and makes it print a message about the empty file not being generated.

23.39.7 Letting Make Build the Project

Rather than type these commands over and over, they can all be placed in a make file. To build the demo project using `make`, save the following in a file called `Makefile`.

Note

This `Makefile` can only be used as input for the GNU version of `make`.

```
PRG          = demo
OBJ          = demo.o
#MCU_TARGET  = at90s2313
#MCU_TARGET  = at90s2333
#MCU_TARGET  = at90s4414
#MCU_TARGET  = at90s4433
#MCU_TARGET  = at90s4434
#MCU_TARGET  = at90s8515
#MCU_TARGET  = at90s8535
#MCU_TARGET  = atmega128
#MCU_TARGET  = atmega1280
#MCU_TARGET  = atmega1281
#MCU_TARGET  = atmega1284p
#MCU_TARGET  = atmega16
#MCU_TARGET  = atmega163
#MCU_TARGET  = atmega164p
#MCU_TARGET  = atmega165
#MCU_TARGET  = atmega165p
#MCU_TARGET  = atmega168
#MCU_TARGET  = atmega169
#MCU_TARGET  = atmega169p
#MCU_TARGET  = atmega2560
#MCU_TARGET  = atmega2561
#MCU_TARGET  = atmega32
#MCU_TARGET  = atmega324p
#MCU_TARGET  = atmega325
#MCU_TARGET  = atmega3250
#MCU_TARGET  = atmega329
#MCU_TARGET  = atmega3290
#MCU_TARGET  = atmega32u4
#MCU_TARGET  = atmega48
#MCU_TARGET  = atmega64
#MCU_TARGET  = atmega640
#MCU_TARGET  = atmega644
#MCU_TARGET  = atmega644p
#MCU_TARGET  = atmega645
#MCU_TARGET  = atmega6450
#MCU_TARGET  = atmega649
#MCU_TARGET  = atmega6490
#MCU_TARGET  = atmega8
#MCU_TARGET  = atmega8515
#MCU_TARGET  = atmega8535
#MCU_TARGET  = atmega88
```

```

#MCU_TARGET      = attiny2313
#MCU_TARGET      = attiny24
#MCU_TARGET      = attiny25
#MCU_TARGET      = attiny26
#MCU_TARGET      = attiny261
#MCU_TARGET      = attiny44
#MCU_TARGET      = attiny45
#MCU_TARGET      = attiny461
#MCU_TARGET      = attiny84
#MCU_TARGET      = attiny85
#MCU_TARGET      = attiny861
OPTIMIZE         = -O2

DEFS             =
LIBS             =

# You should not have to change anything below here.

CC              = avr-gcc

# Override is only needed by avr-lib build system.

override CFLAGS      = -g -Wall $(OPTIMIZE) -mmcu=$(MCU_TARGET) $(DEFS)
override LDFLAGS     = -Wl,-Map,$(PRG).map

OBJCOPY          = avr-objcopy
OBJDUMP          = avr-objdump

all: $(PRG).elf lst text eeprom

$(PRG).elf: $(OBJ)
    $(CC) $(CFLAGS) $(LDFLAGS) -o $@ $^ $(LIBS)

# dependency:
demo.o: demo.c iocompat.h

clean:
    rm -rf *.o $(PRG).elf *.eps *.png *.pdf *.bak
    rm -rf *.lst *.map $(EXTRA_CLEAN_FILES)

lst: $(PRG).lst

%.lst: %.elf
    $(OBJDUMP) -h -S $< > $@

# Rules for building the .text rom images

text: hex bin srec

hex: $(PRG).hex
bin: $(PRG).bin
srec: $(PRG).srec

%.hex: %.elf
    $(OBJCOPY) -j .text -j .data -O ihex $< $@

%.srec: %.elf
    $(OBJCOPY) -j .text -j .data -O srec $< $@

%.bin: %.elf
    $(OBJCOPY) -j .text -j .data -O binary $< $@

```

```
# Rules for building the .eeprom rom images

eeprom: ehex ebin esrec

ehex: $(PRG)_eeprom.hex
ebin: $(PRG)_eeprom.bin
esrec: $(PRG)_eeprom.srec

%_eeprom.hex: %.elf
$(OBJCOPY) -j .eeprom --change-section-lma .eeprom=0 -O ihex $< $@ \
|| { echo empty $@ not generated; exit 0; }

%_eeprom.srec: %.elf
$(OBJCOPY) -j .eeprom --change-section-lma .eeprom=0 -O srec $< $@ \
|| { echo empty $@ not generated; exit 0; }

%_eeprom.bin: %.elf
$(OBJCOPY) -j .eeprom --change-section-lma .eeprom=0 -O binary $< $@ \
|| { echo empty $@ not generated; exit 0; }

# Every thing below here is used by avr-libc's build system and can be ignored
# by the casual user.

FIG2DEV                = fig2dev
EXTRA_CLEAN_FILES      = *.hex *.bin *.srec

dox: eps png pdf

eps: $(PRG).eps
png: $(PRG).png
pdf: $(PRG).pdf

%.eps: %.fig
$(FIG2DEV) -L eps $< $@

%.pdf: %.fig
$(FIG2DEV) -L pdf $< $@

%.png: %.fig
$(FIG2DEV) -L png $< $@
```

23.39.8 Reference to the source code

The source code is installed under

`$prefix/share/doc/avr-libc/examples/demo/`,

where `$prefix` is a configuration option. For Unix systems, it is usually set to either `/usr` or `/usr/local`.

23.40 A more sophisticated project

This project extends the basic idea of the [simple project](#) to control a LED with a PWM output, but adds methods to adjust the LED brightness. It employs a lot of the basic concepts of `avr-libc` to achieve that goal.

Understanding this project assumes the simple project has been understood in full, as

well as being acquainted with the basic hardware concepts of an AVR microcontroller.

23.40.1 Hardware setup

The demo is set up in a way so it can be run on the ATmega16 that ships with the STK500 development kit. The only external part needed is a potentiometer attached to the ADC. It is connected to a 10-pin ribbon cable for port A, both ends of the potentiometer to pins 9 (GND) and 10 (VCC), and the wiper to pin 1 (port A0). A bypass capacitor from pin 1 to pin 9 (like 47 nF) is recommendable.




Figure 6: Setup of the STK500

The coloured patch cables are used to provide various interconnections. As there are only four of them in the STK500, there are two options to connect them for this demo. The second option for the yellow-green cable is shown in parenthesis in the table. Alternatively, the "squid" cable from the JTAG ICE kit can be used if available.

Port	Header	Color	Function	Connect to
D0	1	brown	RxD	RxD of the RS-232 header
D1	2	grey	TxD	TXD of the RS-232 header
D2	3	black	button "down"	SW0 (pin 1 switches header)
D3	4	red	button "up"	SW1 (pin 2 switches header)
D4	5	green	button "ADC"	SW2 (pin 3 switches header)
D5	6	blue	LED	LED0 (pin 1 LEDs header)
D6	7	(green)	clock out	LED1 (pin 2 LEDs header)
D7	8	white	1-second flash	LED2 (pin 3 LEDs header)
GND	9		unused	
VCC	10		unused	




Figure 7: Wiring of the STK500

The following picture shows the alternate wiring where LED1 is connected but SW2 is not:




Figure 8: Wiring option #2 of the STK500

As an alternative, this demo can also be run on the popular ATmega8 controller, or its successor ATmega88 as well as the ATmega48 and ATmega168 variants of the latter. These controllers do not have a port named "A", so their ADC inputs are located on port C instead, thus the potentiometer needs to be attached to port C. Likewise, the OC1A output is not on port D pin 5 but on port B pin 1 (PB1). Thus, the above cabling scheme needs to be changed so that PB1 connects to the LED0 pin. (PD6 remains unconnected.) When using the STK500, use one of the jumper cables for this connection. All other port D pins should be connected the same way as described for the ATmega16 above.

When not using an STK500 starter kit, attach the LEDs through some resistor to Vcc (low-active LEDs), and attach pushbuttons from the respective input pins to GND. The internal pull-up resistors are enabled for the pushbutton pins, so no external resistors are needed.

Finally, the demo has been ported to the ATtiny2313 as well. As this AVR does not offer an ADC, everything related to handling the ADC is disabled in the code for that MCU type. Also, port D of this controller type only features 6 pins, so the 1-second flash LED had to be moved from PD6 to PD4. (PD4 is used as the ADC control button on the other MCU types, but that is not needed here.) OC1A is located at PB3 on this device.

The `MCU_TARGET` macro in the Makefile needs to be adjusted appropriately for the alternative controller types.

The flash ROM and RAM consumption of this demo are way below the resources of even an ATmega48, and still well within the capabilities of an ATtiny2313. The major advantage of experimenting with the ATmega16 (in addition that it ships together with an STK500 anyway) is that it can be debugged online via JTAG. Likewise, the ATmega48/88/168 and ATtiny2313 devices can be debugged through debugWire, using the Atmel JTAG ICE mkII or the low-cost AVR Dragon.

Note that in the explanation below, all port/pin names are applicable to the ATmega16

setup.

23.40.2 Functional overview

PD6 will be toggled with each internal clock tick (approx. 10 ms). PD7 will flash once per second.

PD0 and PD1 are configured as UART IO, and can be used to connect the demo kit to a PC (9600 Bd, 8N1 frame format). The demo application talks to the serial port, and it can be controlled from the serial port.

PD2 through PD4 are configured as inputs, and control the application unless control has been taken over by the serial port. Shorting PD2 to GND will decrease the current PWM value, shorting PD3 to GND will increase it.

While PD4 is shorted to GND, one ADC conversion for channel 0 (ADC input is on PA0) will be triggered each internal clock tick, and the resulting value will be used as the PWM value. So the brightness of the LED follows the analog input value on PC0. VAREF on the STK500 should be set to the same value as VCC.

When running in serial control mode, the function of the watchdog timer can be demonstrated by typing an 'r'. This will make the demo application run in a tight loop without retriggering the watchdog so after some seconds, the watchdog will reset the MCU. This situation can be figured out on startup by reading the MCUCSR register.

The current value of the PWM is backed up in an EEPROM cell after about 3 seconds of idle time after the last change. If that EEPROM cell contains a reasonable (i. e. non-erased) value at startup, it is taken as the initial value for the PWM. This virtually preserves the last value across power cycles. By not updating the EEPROM immediately but only after a timeout, EEPROM wear is reduced considerably compared to immediately writing the value at each change.

23.40.3 A code walkthrough

This section explains the ideas behind individual parts of the code. The [source code](#) has been divided into numbered parts, and the following subsections explain each of these parts.

23.40.3.1 Part 1: Macro definitions

A number of preprocessor macros are defined to improve readability and/or portability of the application.

The first macros describe the IO pins our LEDs and pushbuttons are connected to. This provides some kind of mini-HAL (hardware abstraction layer) so should some of the connections be changed, they don't need to be changed inside the code but only on top. Note that the location of the PWM output itself is mandated by the hardware, so it cannot be easily changed. As the ATmega48/88/168 controllers belong to a more recent generation of AVR's, a number of register and bit names have been changed there, so they are mapped back to their ATmega8/16 equivalents to keep the actual program code portable.

The name `F_CPU` is the conventional name to describe the CPU clock frequency of the controller. This demo project just uses the internal calibrated 1 MHz RC oscillator that is enabled by default. Note that when using the `<util/delay.h>` functions, `F_CPU` needs to be defined before including that file.

The remaining macros have their own comments in the source code. The macro `TMR1_SCALE` shows how to use the preprocessor and the compiler's constant expression computation to calculate the value of timer 1's post-scaler in a way so it only depends on `F_CPU` and the desired software clock frequency. While the formula looks a bit complicated, using a macro offers the advantage that the application will automatically scale to new target softclock or master CPU frequencies without having to manually re-calculate hardcoded constants.

23.40.3.2 Part 2: Variable definitions

The `intflags` structure demonstrates a way to allocate bit variables in memory. Each of the interrupt service routines just sets one bit within that structure, and the application's main loop then monitors the bits in order to act appropriately.

Like all variables that are used to communicate values between an interrupt service routine and the main application, it is declared `volatile`.

The variable `ee_pwm` is not a variable in the classical C sense that could be used as an lvalue or within an expression to obtain its value. Instead, the

```
__attribute__((section(".eeprom")))
```

marks it as belonging to the `EEPROM` section. This section is merely used as a placeholder so the compiler can arrange for each individual variable's location in EEPROM. The compiler will also keep track of initial values assigned, and usually the Makefile is arranged to extract these initial values into a separate load file (`largedemo_eeeprom.*` in this case) that can be used to initialize the EEPROM.

The actual EEPROM IO must be performed manually.

Similarly, the variable `mcucsr` is kept in the `.noinit` section in order to prevent it from being cleared upon application startup.

23.40.3.3 Part 3: Interrupt service routines

The ISR to handle timer 1's overflow interrupt arranges for the software clock. While timer 1 runs the PWM, it calls its overflow handler rather frequently, so the `TMR1_SCALE` value is used as a postscaler to reduce the internal software clock frequency further. If the software clock triggers, it sets the `tmr_int` bitfield, and defers all further tasks to the main loop.

The ADC ISR just fetches the value from the ADC conversion, disables the ADC interrupt again, and announces the presence of the new value in the `adc_int` bitfield. The interrupt is kept disabled while not needed, because the ADC will also be triggered by executing the `SLEEP` instruction in idle mode (which is the default sleep mode). Another option would be to turn off the ADC completely here, but that increases the ADC's startup time (not that it would matter much for this application).

23.40.3.4 Part 4: Auxiliary functions

The function `handle_mcucsr()` uses two `__attribute__` declarators to achieve specific goals. First, it will instruct the compiler to place the generated code into the `.init3` section of the output. Thus, it will become part of the application initialization sequence. This is done in order to fetch (and clear) the reason of the last hardware reset from `MCUCSR` as early as possible. There is a short period of time where the next reset could already trigger before the current reason has been evaluated. This also explains why the variable `mcucsr` that mirrors the register's value needs to be placed into the `.noinit` section, because otherwise the default initialization (which happens after `.init3`) would blank the value again.

As the initialization code is not called using `CALL/RET` instructions but rather concatenated together, the compiler needs to be instructed to omit the entire function prologue and epilogue. This is performed by the *naked* attribute. So while syntactically, `handle_mcucsr()` is a function to the compiler, the compiler will just emit the instructions for it without setting up any stack frame, and not even a `RET` instruction at the end.

Function `ioinit()` centralizes all hardware setup. The very last part of that function demonstrates the use of the EEPROM variable `ee_pwm` to obtain an EEPROM address that can in turn be applied as an argument to `eeeprom_read_word()`.

The following functions handle UART character and string output. (UART input is handled by an ISR.) There are two string output functions, `printstr()` and `printstr_p()`. The latter function fetches the string from [program memory](#). Both functions translate a newline character into a carriage return/newline sequence, so a simple `\n` can be used in the source code.

The function `set_pwm()` propagates the new PWM value to the PWM, performing range checking. When the value has been changed, the new percentage will be announced on the serial link. The current value is mirrored in the variable `pwm` so others can use it in calculations. In order to allow for a simple calculation of a percentage value without requiring floating-point mathematics, the maximal value of the PWM is restricted to 1000 rather than 1023, so a simple division by 10 can be used. Due to the nature of the human eye, the difference in LED brightness between 1000 and 1023 is not noticable anyway.

23.40.3.5 Part 5: main()

At the start of `main()`, a variable `mode` is declared to keep the current mode of operation. An enumeration is used to improve the readability. By default, the compiler would allocate a variable of type *int* for an enumeration. The *packed* attribute declarator instructs the compiler to use the smallest possible integer type (which would be an 8-bit type here).

After some initialization actions, the application's main loop follows. In an embedded application, this is normally an infinite loop as there is nothing an application could "exit" into anyway.

At the beginning of the loop, the watchdog timer will be retriggered. If that timer is not triggered for about 2 seconds, it will issue a hardware reset. Care needs to be taken

that no code path blocks longer than this, or it needs to frequently perform watchdog resets of its own. An example of such a code path would be the string IO functions: for an overly large string to print (about 2000 characters at 9600 Bd), they might block for too long.

The loop itself then acts on the interrupt indication bitfields as appropriate, and will eventually put the CPU on sleep at its end to conserve power.

The first interrupt bit that is handled is the (software) timer, at a frequency of approximately 100 Hz. The `CLOCKOUT` pin will be toggled here, so e. g. an oscilloscope can be used on that pin to measure the accuracy of our software clock. Then, the LED flasher for LED2 ("We are alive"-LED) is built. It will flash that LED for about 50 ms, and pause it for another 950 ms. Various actions depending on the operation mode follow. Finally, the 3-second backup timer is implemented that will write the PWM value back to EEPROM once it is not changing anymore.

The ADC interrupt will just adjust the PWM value only.

Finally, the UART Rx interrupt will dispatch on the last character received from the UART.

All the string literals that are used as informational messages within `main()` are placed in [program memory](#) so no SRAM needs to be allocated for them. This is done by using the `PSTR` macro, and passing the string to `printstr_p()`.

23.40.4 The source code

The source code is installed under

```
$prefix/share/doc/avr-libc/examples/largedemo/largedemo.c,
```

where `$prefix` is a configuration option. For Unix systems, it is usually set to either `/usr` or `/usr/local`.

23.41 Using the standard IO facilities

This project illustrates how to use the standard IO facilities (`stdio`) provided by this library. It assumes a basic knowledge of how the `stdio` subsystem is used in standard C applications, and concentrates on the differences in this library's implementation that mainly result from the differences of the microcontroller environment, compared to a hosted environment of a standard computer.

This demo is meant to supplement the [documentation](#), not to replace it.

23.41.1 Hardware setup

The demo is set up in a way so it can be run on the ATmega16 that ships with the STK500 development kit. The UART port needs to be connected to the RS-232 "spare" port by a jumper cable that connects PD0 to RxD and PD1 to TxD. The RS-232 channel is set up as standard input (`stdin`) and standard output (`stdout`), respectively.

In order to have a different device available for a standard error channel (`stderr`), an

industry-standard LCD display with an HD44780-compatible LCD controller has been chosen. This display needs to be connected to port A of the STK500 in the following way:

Port	Header	Function
A0	1	LCD D4
A1	2	LCD D5
A2	3	LCD D6
A3	4	LCD D7
A4	5	LCD R/~W
A5	6	LCD E
A6	7	LCD RS
A7	8	unused
GND	9	GND
VCC	10	Vcc




Figure 9: Wiring of the STK500

The LCD controller is used in 4-bit mode, including polling the "busy" flag so the R/~W line from the LCD controller needs to be connected. Note that the LCD controller has yet another supply pin that is used to adjust the LCD's contrast (V5). Typically, that pin connects to a potentiometer between Vcc and GND. Often, it might work to just connect that pin to GND, while leaving it unconnected usually yields an unreadable display.

Port A has been chosen as 7 pins are needed to connect the LCD, yet all other ports are already partially in use: port B has the pins for in-system programming (ISP), port C has the ports for JTAG (can be used for debugging), and port D is used for the UART connection.

23.41.2 Functional overview

The project consists of the following files:

- `stdiodemo.c` This is the main example file.
- `defines.h` Contains some global defines, like the LCD wiring
- `hd44780.c` Implementation of an HD44780 LCD display driver
- `hd44780.h` Interface declarations for the HD44780 driver
- `lcd.c` Implementation of LCD character IO on top of the HD44780 driver
- `lcd.h` Interface declarations for the LCD driver
- `uart.c` Implementation of a character IO driver for the internal UART
- `uart.h` Interface declarations for the UART driver

23.41.3 A code walkthrough

23.41.3.1 `stdiodemo.c`

As usual, include files go first. While conventionally, system header files (those in angular brackets `<...>`) go before application-specific header files (in double quotes), `defines.h` comes as the first header file here. The main reason is that this file defines the value of `F_CPU` which needs to be known before including `<utils/delay.h>`.

The function `ioinit()` summarizes all hardware initialization tasks. As this function is declared to be module-internal only (`static`), the compiler will notice its simplicity, and with a reasonable optimization level in effect, it will inline that function. That needs to be kept in mind when debugging, because the inlining might cause the debugger to "jump around wildly" at a first glance when single-stepping.

The definitions of `uart_str` and `lcd_str` set up two stdio streams. The initialization is done using the `FDEV_SETUP_STREAM()` initializer template macro, so a static object can be constructed that can be used for IO purposes. This initializer macro takes three arguments, two function macros to connect the corresponding output and input functions, respectively, the third one describes the intent of the stream (read, write, or both). Those functions that are not required by the specified intent (like the input function for `lcd_str` which is specified to only perform output operations) can be given as `NULL`.

The stream `uart_str` corresponds to input and output operations performed over the RS-232 connection to a terminal (e.g. from/to a PC running a terminal program), while the `lcd_str` stream provides a method to display character data on the LCD text display.

The function `delay_1s()` suspends program execution for approximately one second. This is done using the `_delay_ms()` function from `<util/delay.h>` which in turn needs the `F_CPU` macro in order to adjust the cycle counts. As the `_delay_ms()` function has a limited range of allowable argument values (depending

on `F_CPU`), a value of 10 ms has been chosen as the base delay which would be safe for CPU frequencies of up to about 26 MHz. This function is then called 100 times to accomodate for the actual one-second delay.

In a practical application, long delays like this one were better be handled by a hardware timer, so the main CPU would be free for other tasks while waiting, or could be put on sleep.

At the beginning of `main()`, after initializing the peripheral devices, the default stdio streams `stdin`, `stdout`, and `stderr` are set up by using the existing static `FILE` stream objects. While this is not mandatory, the availability of `stdin` and `stdout` allows to use the shorthand functions (e.g. `printf()` instead of `fprintf()`), and `stderr` can mnemonically be referred to when sending out diagnostic messages.

Just for demonstration purposes, `stdin` and `stdout` are connected to a stream that will perform UART IO, while `stderr` is arranged to output its data to the LCD text display.

Finally, a main loop follows that accepts simple "commands" entered via the RS-232 connection, and performs a few simple actions based on the commands.

First, a prompt is sent out using `printf_P()` (which takes a [program space string](#)). The string is read into an internal buffer as one line of input, using `fgets()`. While it would be also possible to use `gets()` (which implicitly reads from `stdin`), `gets()` has no control that the user's input does not overflow the input buffer provided so it should never be used at all.

If `fgets()` fails to read anything, the main loop is left. Of course, normally the main loop of a microcontroller application is supposed to never finish, but again, for demonstrational purposes, this explains the error handling of stdio. `fgets()` will return `NULL` in case of an input error or end-of-file condition on input. Both these conditions are in the domain of the function that is used to establish the stream, `uart_putchar()` in this case. In short, this function returns `EOF` in case of a serial line "break" condition (extended start condition) has been recognized on the serial line. Common PC terminal programs allow to assert this condition as some kind of out-of-band signalling on an RS-232 connection.

When leaving the main loop, a goodbye message is sent to standard error output (i.e. to the LCD), followed by three dots in one-second spacing, followed by a sequence that will clear the LCD. Finally, `main()` will be terminated, and the library will add an infinite loop, so only a CPU reset will be able to restart the application.

There are three "commands" recognized, each determined by the first letter of the line entered (converted to lower case):

- The 'q' (quit) command has the same effect of leaving the main loop.
- The 'l' (LCD) command takes its second argument, and sends it to the LCD.
- The 'u' (UART) command takes its second argument, and sends it back to the UART connection.

Command recognition is done using `sscanf()` where the first format in the format string just skips over the command itself (as the assignment suppression modifier `*` is given).

23.41.3.2 defines.h

This file just contains a few peripheral definitions.

The `F_CPU` macro defines the CPU clock frequency, to be used in delay loops, as well as in the UART baud rate calculation.

The macro `UART_BAUD` defines the RS-232 baud rate. Depending on the actual CPU frequency, only a limited range of baud rates can be supported.

The remaining macros customize the IO port and pins used for the HD44780 LCD driver. Each definition consists of a letter naming the port this pin is attached to, and a respective bit number. For accessing the data lines, only the first data line gets its own macro (line D4 on the HD44780, lines D0 through D3 are not used in 4-bit mode), all other data lines are expected to be in ascending order next to D4.

23.41.3.3 hd44780.h

This file describes the public interface of the low-level LCD driver that interfaces to the HD44780 LCD controller. Public functions are available to initialize the controller into 4-bit mode, to wait for the controller's busy bit to be clear, and to read or write one byte from or to the controller.

As there are two different forms of controller IO, one to send a command or receive the controller status (RS signal clear), and one to send or receive data to/from the controller's SRAM (RS asserted), macros are provided that build on the mentioned function primitives.

Finally, macros are provided for all the controller commands to allow them to be used symbolically. The HD44780 datasheet explains these basic functions of the controller in more detail.

23.41.3.4 hd44780.c

This is the implementation of the low-level HD44780 LCD controller driver.

On top, a few preprocessor glueing tricks are used to establish symbolic access to the hardware port pins the LCD controller is attached to, based on the application's definitions made in [defines.h](#).

The `hd44780_pulse_e()` function asserts a short pulse to the controller's E (enable) pin. Since reading back the data asserted by the LCD controller needs to be performed while E is active, this function reads and returns the input data if the parameter `readback` is true. When called with a compile-time constant parameter that is false, the compiler will completely eliminate the unused readback operation, as well as the return value as part of its optimizations.

As the controller is used in 4-bit interface mode, all byte IO to/from the controller needs to be handled as two nibble IOs. The functions `hd44780_outnibble()` and `hd44780_innibble()` implement this. They do not belong to the public interface, so they are declared static.

Building upon these, the public functions `hd44780_outbyte()` and `hd44780_in-`

`inbyte()` transfer one byte to/from the controller.

The function `hd44780_wait_ready()` waits for the controller to become ready, by continuously polling the controller's status (which is read by performing a byte read with the RS signal cleared), and examining the BUSY flag within the status byte. This function needs to be called before performing any controller IO.

Finally, `hd44780_init()` initializes the LCD controller into 4-bit mode, based on the initialization sequence mandated by the datasheet. As the BUSY flag cannot be examined yet at this point, this is the only part of this code where timed delays are used. While the controller can perform a power-on reset when certain constraints on the power supply rise time are met, always calling the software initialization routine at startup ensures the controller will be in a known state. This function also puts the interface into 4-bit mode (which would not be done automatically after a power-on reset).

23.41.3.5 `lcd.h`

This function declares the public interface of the higher-level (character IO) LCD driver.

23.41.3.6 `lcd.c`

The implementation of the higher-level LCD driver. This driver builds on top of the HD44780 low-level LCD controller driver, and offers a character IO interface suitable for direct use by the standard IO facilities. Where the low-level HD44780 driver deals with setting up controller SRAM addresses, writing data to the controller's SRAM, and controlling display functions like clearing the display, or moving the cursor, this high-level driver allows to just write a character to the LCD, in the assumption this will somehow show up on the display.

Control characters can be handled at this level, and used to perform specific actions on the LCD. Currently, there is only one control character that is being dealt with: a newline character (`\n`) is taken as an indication to clear the display and set the cursor into its initial position upon reception of the next character, so a "new line" of text can be displayed. Therefore, a received newline character is remembered until more characters have been sent by the application, and will only then cause the display to be cleared before continuing. This provides a convenient abstraction where full lines of text can be sent to the driver, and will remain visible at the LCD until the next line is to be displayed.

Further control characters could be implemented, e. g. using a set of escape sequences. That way, it would be possible to implement self-scrolling display lines etc.

The public function `lcd_init()` first calls the initialization entry point of the lower-level HD44780 driver, and then sets up the LCD in a way we'd like to (display cleared, non-blinking cursor enabled, SRAM addresses are increasing so characters will be written left to right).

The public function `lcd_putchar()` takes arguments that make it suitable for being passed as a `put()` function pointer to the stdio stream initialization functions and

macros (`fdevopen()`, `FDEV_SETUP_STREAM()` etc.). Thus, it takes two arguments, the character to display itself, and a reference to the underlying stream object, and it is expected to return 0 upon success.

This function remembers the last unprocessed newline character seen in the function-local static variable `nl_seen`. If a newline character is encountered, it will simply set this variable to a true value, and return to the caller. As soon as the first non-newline character is to be displayed with `nl_seen` still true, the LCD controller is told to clear the display, put the cursor home, and restart at SRAM address 0. All other characters are sent to the display.

The single static function-internal variable `nl_seen` works for this purpose. If multiple LCDs should be controlled using the same set of driver functions, that would not work anymore, as a way is needed to distinguish between the various displays. This is where the second parameter can be used, the reference to the stream itself: instead of keeping the state inside a private variable of the function, it can be kept inside a private object that is attached to the stream itself. A reference to that private object can be attached to the stream (e.g. inside the function `lcd_init()` that then also needs to be passed a reference to the stream) using `fdev_set_udata()`, and can be accessed inside `lcd_putchar()` using `fdev_get_udata()`.

23.41.3.7 `uart.h`

Public interface definition for the RS-232 UART driver, much like in `lcd.h` except there is now also a character input function available.

As the RS-232 input is line-buffered in this example, the macro `RX_BUFSIZE` determines the size of that buffer.

23.41.3.8 `uart.c`

This implements an stdio-compatible RS-232 driver using an AVR's standard UART (or USART in asynchronous operation mode). Both, character output as well as character input operations are implemented. Character output takes care of converting the internal newline `\n` into its external representation carriage return/line feed (`\r\n`).

Character input is organized as a line-buffered operation that allows to minimally edit the current line until it is "sent" to the application when either a carriage return (`\r`) or newline (`\n`) character is received from the terminal. The line editing functions implemented are:

- `\b` (back space) or `\177` (delete) deletes the previous character
- `^u` (control-U, ASCII NAK) deletes the entire input buffer
- `^w` (control-W, ASCII ETB) deletes the previous input word, delimited by white space
- `^r` (control-R, ASCII DC2) sends a `\r`, then reprints the buffer (refresh)
- `\t` (tabulator) will be replaced by a single space

The function `uart_init()` takes care of all hardware initialization that is required to put the UART into a mode with 8 data bits, no parity, one stop bit (commonly referred to as 8N1) at the baud rate configured in [defines.h](#). At low CPU clock frequencies, the `U2X` bit in the UART is set, reducing the oversampling from 16x to 8x, which allows for a 9600 Bd rate to be achieved with tolerable error using the default 1 MHz RC oscillator.

The public function `uart_putchar()` again has suitable arguments for direct use by the stdio stream interface. It performs the `\n` into `\r\n` translation by recursively calling itself when it sees a `\n` character. Just for demonstration purposes, the `\a` (audible bell, ASCII BEL) character is implemented by sending a string to `stderr`, so it will be displayed on the LCD.

The public function `uart_getchar()` implements the line editor. If there are characters available in the line buffer (variable `rxp` is not `NULL`), the next character will be returned from the buffer without any UART interaction.

If there are no characters inside the line buffer, the input loop will be entered. Characters will be read from the UART, and processed accordingly. If the UART signalled a framing error (FE bit set), typically caused by the terminal sending a *line break* condition (start condition held much longer than one character period), the function will return an end-of-file condition using `_FDEV_EOF`. If there was a data overrun condition on input (DOR bit set), an error condition will be returned as `_FDEV_ERR`.

Line editing characters are handled inside the loop, potentially modifying the buffer status. If characters are attempted to be entered beyond the size of the line buffer, their reception is refused, and a `\a` character is sent to the terminal. If a `\r` or `\n` character is seen, the variable `rxp` (receive pointer) is set to the beginning of the buffer, the loop is left, and the first character of the buffer will be returned to the application. (If no other characters have been entered, this will just be the newline character, and the buffer is marked as being exhausted immediately again.)

23.41.4 The source code

The source code is installed under

```
$prefix/share/doc/avr-libc/examples/stdiodemo/,
```

where `$prefix` is a configuration option. For Unix systems, it is usually set to either `/usr` or `/usr/local`.

23.42 Example using the two-wire interface (TWI)

Some newer devices of the ATmega series contain builtin support for interfacing the microcontroller to a two-wire bus, called TWI. This is essentially the same called I2C by Philips, but that term is avoided in Atmel's documentation due to patenting issues.

For further documentation, see:

http://www.nxp.com/documents/user_manual/UM10204.pdf

23.42.1 Introduction into TWI

The two-wire interface consists of two signal lines named *SDA* (serial data) and *SCL* (serial clock) (plus a ground line, of course). All devices participating in the bus are connected together, using open-drain driver circuitry, so the wires must be terminated using appropriate pullup resistors. The pullups must be small enough to recharge the line capacity in short enough time compared to the desired maximal clock frequency, yet large enough so all drivers will not be overloaded. There are formulas in the datasheet that help selecting the pullups.

Devices can either act as a master to the bus (i. e., they initiate a transfer), or as a slave (they only act when being called by a master). The bus is multi-master capable, and a particular device implementation can act as either master or slave at different times. Devices are addressed using a 7-bit address (coordinated by Philips) transferred as the first byte after the so-called start condition. The LSB of that byte is R/~W, i. e. it determines whether the request to the slave is to read or write data during the next cycles. (There is also an option to have devices using 10-bit addresses but that is not covered by this example.)

23.42.2 The TWI example project

The ATmega TWI hardware supports both, master and slave operation. This example will only demonstrate how to use an AVR microcontroller as TWI master. The implementation is kept simple in order to concentrate on the steps that are required to talk to a TWI slave, so all processing is done in polled-mode, waiting for the TWI interface to indicate that the next processing step is due (by setting the TWINT interrupt bit). If it is desired to have the entire TWI communication happen in "background", all this can be implemented in an interrupt-controlled way, where only the start condition needs to be triggered from outside the interrupt routine.

There is a variety of slave devices available that can be connected to a TWI bus. For the purpose of this example, an EEPROM device out of the industry-standard **24Cxx** series has been chosen (where *xx* can be one of **01**, **02**, **04**, **08**, or **16**) which are available from various vendors. The choice was almost arbitrary, mainly triggered by the fact that an EEPROM device is being talked to in both directions, reading and writing the slave device, so the example will demonstrate the details of both.

Usually, there is probably not much need to add more EEPROM to an ATmega system that way: the smallest possible AVR device that offers hardware TWI support is the ATmega8 which comes with 512 bytes of EEPROM, which is equivalent to an 24C04 device. The ATmega128 already comes with twice as much EEPROM as the 24C16 would offer. One exception might be to use an externally connected EEPROM device that is removable; e. g. SDRAM PC memory comes with an integrated TWI EEPROM that carries the RAM configuration information.

23.42.3 The Source Code

The source code is installed under

```
$prefix/share/doc/avr-libc/examples/twitest/twitest.c,
```

where `$prefix` is a configuration option. For Unix systems, it is usually set to either `/usr` or `/usr/local`.

Note [1]

The header file `<util/twi.h>` contains some macro definitions for symbolic constants used in the TWI status register. These definitions match the names used in the Atmel datasheet except that all names have been prefixed with `TW_`.

Note [2]

The clock is used in timer calculations done by the compiler, for the UART baud rate and the TWI clock rate.

Note [3]

The address assigned for the 24Cxx EEPROM consists of 1010 in the upper four bits. The following three bits are normally available as slave sub-addresses, allowing to operate more than one device of the same type on a single bus, where the actual sub-address used for each device is configured by hardware strapping. However, since the next data packet following the device selection only allows for 8 bits that are used as an EEPROM address, devices that require more than 8 address bits (24C04 and above) "steal" subaddress bits and use them for the EEPROM cell address bits 9 to 11 as required. This example simply assumes all subaddress bits are 0 for the smaller devices, so the E0, E1, and E2 inputs of the 24Cxx must be grounded.

Note [3a]

EEPROMs of type 24C32 and above cannot be addressed anymore even with the sub-address bit trick. Thus, they require the upper address bits being sent separately on the bus. When activating the `WORD_ADDRESS_16BIT` define, the algorithm implements that auxiliary address byte transmission.

Note [4]

For slow clocks, enable the 2 x U[S]ART clock multiplier, to improve the baud rate error. This will allow a 9600 Bd communication using the standard 1 MHz calibrated RC oscillator. See also the Baud rate tables in the datasheets.

Note [5]

The datasheet explains why a minimum TWBR value of 10 should be maintained when running in master mode. Thus, for system clocks below 3.6 MHz, we cannot run the bus at the intended clock rate of 100 kHz but have to slow down accordingly.

Note [6]

This function is used by the standard output facilities that are utilized in this example for debugging and demonstration purposes.

Note [7]

In order to shorten the data to be sent over the TWI bus, the 24Cxx EEPROMs support multiple data bytes transferred within a single request, maintaining an internal address counter that is updated after each data byte transferred successfully. When reading data, one request can read the entire device memory if desired (the counter would wrap around and start back from 0 when reaching the end of the device).

Note [8]

When reading the EEPROM, a first device selection must be made with write intent (R/~W bit set to 0 indicating a write operation) in order to transfer the EEPROM address to start reading from. This is called *master transmitter mode*. Each completion of a particular step in TWI communication is indicated by an asserted TWINT bit in TWCR. (An interrupt would be generated if allowed.) After performing any actions that are needed for the next communication step, the interrupt condition must be manually cleared by *setting* the TWINT bit. Unlike with many other interrupt sources, this would even be required when using a true interrupt routine, since as soon as TWINT is re-asserted, the next bus transaction will start.

Note [9]

Since the TWI bus is multi-master capable, there is potential for a bus contention when one master starts to access the bus. Normally, the TWI bus interface unit will detect this situation, and will not initiate a start condition while the bus is busy. However, in case two masters were starting at exactly the same time, the way bus arbitration works, there is always a chance that one master could lose arbitration of the bus during any transmit operation. A master that has lost arbitration is required by the protocol to immediately cease talking on the bus; in particular it must not initiate a stop condition in order to not corrupt the ongoing transfer from the active master. In this example, upon detecting a lost arbitration condition, the entire transfer is going to be restarted. This will cause a new start condition to be initiated, which will normally be delayed until the currently active master has released the bus.

Note [10]

Next, the device slave is going to be reselected (using a so-called repeated start condition which is meant to guarantee that the bus arbitration will remain at the current master) using the same slave address (SLA), but this time with read intent (R/~W bit set to 1) in order to request the device slave to start transferring data from the slave to the master in the next packet.

Note [11]

If the EEPROM device is still busy writing one or more cells after a previous write request, it will simply leave its bus interface drivers at high impedance, and does not respond to a selection in any way at all. The master selecting the device will see the high level at SDA after transferring the SLA+R/W packet as a NACK to its selection request. Thus, the select process is simply started over (effectively causing a *repeated start condition*), until the device will eventually respond. This polling procedure is recommended in the 24Cxx datasheet in order to minimize the busy wait time when writing. Note that in case a device is broken and never responds to a selection (e. g. since it is no longer present at all), this will cause an infinite loop. Thus the maximal number of iterations made until the device is declared to be not responding at all, and an error is returned, will be limited to MAX_ITER.

Note [12]

This is called *master receiver mode*: the bus master still supplies the SCL clock, but the device slave drives the SDA line with the appropriate data. After 8 data bits, the master responds with an ACK bit (SDA driven low) in order to request another data transfer from the slave, or it can leave the SDA line high (NACK), indicating to the slave that it is going to stop the transfer now. Assertion of ACK is handled by setting the TWEA bit in TWCR when starting the current transfer.

Note [13]

The control word sent out in order to initiate the transfer of the next data packet is initially set up to assert the TWEA bit. During the last loop iteration, TWEA is de-asserted so the client will get informed that no further transfer is desired.

Note [14]

Except in the case of lost arbitration, all bus transactions must properly be terminated by the master initiating a stop condition.

Note [15]

Writing to the EEPROM device is simpler than reading, since only a master transmitter mode transfer is needed. Note that the first packet after the SLA+W selection is always considered to be the EEPROM address for the next operation. (This packet is exactly the same as the one above sent before starting to read the device.) In case a master transmitter mode transfer is going to send more than one data packet, all following packets will be considered data bytes to write at the indicated address. The internal address pointer will be incremented after each write operation.

Note [16]

24Cxx devices can become write-protected by strapping their \sim WC pin to logic high. (Leaving it unconnected is explicitly allowed, and constitutes logic low level, i. e. no write protection.) In case of a write protected device, all data transfer attempts will be NACKed by the device. Note that some devices might not implement this.

24 Data Structure Documentation

24.1 div_t Struct Reference

Data Fields

- int [quot](#)
- int [rem](#)

24.1.1 Detailed Description

Result type for function [div\(\)](#).

24.1.2 Field Documentation

24.1.2.1 int div_t::quot

The Quotient.

24.1.2.2 int div_t::rem

The Remainder.

The documentation for this struct was generated from the following file:

- [stdlib.h](#)

24.2 ldiv_t Struct Reference

Data Fields

- long [quot](#)
- long [rem](#)

24.2.1 Detailed Description

Result type for function [ldiv\(\)](#).

24.2.2 Field Documentation

24.2.2.1 long ldiv_t::quot

The Quotient.

24.2.2.2 long ldiv_t::rem

The Remainder.

The documentation for this struct was generated from the following file:

- [stdlib.h](#)

24.3 tm Struct Reference

Data Fields

- [int8_t](#) [tm_sec](#)
- [int8_t](#) [tm_min](#)
- [int8_t](#) [tm_hour](#)
- [int8_t](#) [tm_mday](#)
- [int8_t](#) [tm_wday](#)
- [int8_t](#) [tm_mon](#)
- [int16_t](#) [tm_year](#)
- [int16_t](#) [tm_yday](#)
- [int16_t](#) [tm_isdst](#)

24.3.1 Detailed Description

The tm structure contains a representation of time 'broken down' into components of the Gregorian calendar.

The value of [tm_isdst](#) is zero if Daylight Saving Time is not in effect, and is negative if the information is not available.

When Daylight Saving Time is in effect, the value represents the number of seconds the clock is advanced.

See the [set_dst\(\)](#) function for more information about Daylight Saving.

24.3.2 Field Documentation

24.3.2.1 int8_t tm::tm_hour

hours since midnight - [0 to 23]

24.3.2.2 int16_t tm::tm_isdst

Daylight Saving Time flag

24.3.2.3 int8_t tm::tm_mday

day of the month - [1 to 31]

24.3.2.4 int8_t tm::tm_min

minutes after the hour - [0 to 59]

24.3.2.5 int8_t tm::tm_mon

months since January - [0 to 11]

24.3.2.6 int8_t tm::tm_sec

seconds after the minute - [0 to 59]

24.3.2.7 int8_t tm::tm_wday

days since Sunday - [0 to 6]

24.3.2.8 int16_t tm::tm_yday

days since January 1 - [0 to 365]

24.3.2.9 int16_t tm::tm_year

years since 1900

The documentation for this struct was generated from the following file:

- [time.h](#)

24.4 week_date Struct Reference

Data Fields

- int [year](#)
- int [week](#)
- int [day](#)

24.4.1 Detailed Description

Structure which represents a date as a year, week number of that year, and day of week. See http://en.wikipedia.org/wiki/ISO_week_date for more information.

24.4.2 Field Documentation

24.4.2.1 int week_date::day

day within week

24.4.2.2 int week_date::week

week number (#1 is where first Thursday is in)

24.4.2.3 int week_date::year

year number (Gregorian calendar)

The documentation for this struct was generated from the following file:

- [time.h](#)

25 File Documentation

25.1 assert.h File Reference

Defines

- #define [assert](#)(expression)

25.1.1 Detailed Description

25.2 atoi.S File Reference

25.2.1 Detailed Description

25.3 atol.S File Reference

25.3.1 Detailed Description

25.4 atomic.h File Reference

Defines

- #define [ATOMIC_BLOCK](#)(type)
- #define [NONATOMIC_BLOCK](#)(type)
- #define [ATOMIC_RESTORESTATE](#)
- #define [ATOMIC_FORCEON](#)
- #define [NONATOMIC_RESTORESTATE](#)
- #define [NONATOMIC_FORCEOFF](#)

25.4.1 Detailed Description

25.5 boot.h File Reference

Defines

- #define [BOOTLOADER_SECTION](#) __attribute__((section(".bootloader")))
- #define [boot_spm_interrupt_enable](#)() (__SPM_REG |= (uint8_t)_BV(SPMIE))
- #define [boot_spm_interrupt_disable](#)() (__SPM_REG &= (uint8_t)~_BV(SPMIE))
- #define [boot_is_spm_interrupt](#)() (__SPM_REG & (uint8_t)_BV(SPMIE))
- #define [boot_rww_busy](#)() (__SPM_REG & (uint8_t)_BV(__COMMON_ASB))
- #define [boot_spm_busy](#)() (__SPM_REG & (uint8_t)_BV(__SPM_ENABLE))
- #define [boot_spm_busy_wait](#)() do{ }while(boot_spm_busy())
- #define [GET_LOW_FUSE_BITS](#) (0x0000)
- #define [GET_LOCK_BITS](#) (0x0001)
- #define [GET_EXTENDED_FUSE_BITS](#) (0x0002)
- #define [GET_HIGH_FUSE_BITS](#) (0x0003)
- #define [boot_lock_fuse_bits_get](#)(address)
- #define [boot_signature_byte_get](#)(addr)
- #define [boot_page_fill](#)(address, data) __boot_page_fill_normal(address, data)
- #define [boot_page_erase](#)(address) __boot_page_erase_normal(address)
- #define [boot_page_write](#)(address) __boot_page_write_normal(address)
- #define [boot_rww_enable](#)() __boot_rww_enable()
- #define [boot_lock_bits_set](#)(lock_bits) __boot_lock_bits_set(lock_bits)
- #define [boot_page_fill_safe](#)(address, data)

- #define [boot_page_erase_safe](#)(address)
- #define [boot_page_write_safe](#)(address)
- #define [boot_rww_enable_safe](#)()
- #define [boot_lock_bits_set_safe](#)(lock_bits)

25.5.1 Detailed Description

25.6 cpufunc.h File Reference

Defines

- #define [_NOP](#)()
- #define [_MemoryBarrier](#)()

Functions

- void [ccp_write_io](#) (uint8_t *__ioaddr, uint8_t __value)

25.6.1 Detailed Description

25.7 crc16.h File Reference

Functions

- static __inline__ [uint16_t_crc16_update](#) (uint16_t __crc, uint8_t __data)
- static __inline__ [uint16_t_crc_xmodem_update](#) (uint16_t __crc, uint8_t __data)
- static __inline__ [uint16_t_crc_ccitt_update](#) (uint16_t __crc, uint8_t __data)
- static __inline__ [uint8_t_crc_ibutton_update](#) (uint8_t __crc, uint8_t __data)
- static __inline__ [uint8_t_crc8_ccitt_update](#) (uint8_t __crc, uint8_t __data)

25.7.1 Detailed Description

25.8 ctype.h File Reference

Functions

Character classification routines

These functions perform character classification. They return true or false status depending whether the character passed to the function falls into the function's classification (i.e. [isdigit](#)() returns true if its argument is any value '0' through '9', inclusive). If the input is not an unsigned char value, all of this function return false.

- int [isalnum](#) (int __c)
- int [isalpha](#) (int __c)
- int [isascii](#) (int __c)

- int [isblank](#) (int __c)
- int [iscntrl](#) (int __c)
- int [isdigit](#) (int __c)
- int [isgraph](#) (int __c)
- int [islower](#) (int __c)
- int [isprint](#) (int __c)
- int [ispunct](#) (int __c)
- int [isspace](#) (int __c)
- int [isupper](#) (int __c)
- int [isxdigit](#) (int __c)

Character conversion routines

This realization permits all possible values of integer argument. The [toascii\(\)](#) function clears all highest bits. The [tolower\(\)](#) and [toupper\(\)](#) functions return an input argument as is, if it is not an unsigned char value.

- int [toascii](#) (int __c)
- int [tolower](#) (int __c)
- int [toupper](#) (int __c)

25.8.1 Detailed Description

25.9 delay.h File Reference

Defines

- #define [F_CPU](#) 1000000UL

Functions

- void [_delay_ms](#) (double __ms)
- void [_delay_us](#) (double __us)

25.9.1 Detailed Description

25.10 delay_basic.h File Reference

Functions

- void [_delay_loop_1](#) (uint8_t __count)
- void [_delay_loop_2](#) (uint16_t __count)

25.10.1 Detailed Description

25.11 errno.h File Reference

Defines

- #define [EDOM](#) 33

- #define [ERANGE](#) 34

Variables

- int [errno](#)

25.11.1 Detailed Description

25.12 fdevopen.c File Reference

Functions

- [FILE *](#) [fdevopen](#) (int(*put)(char, [FILE *](#)), int(*get)([FILE *](#)))

25.12.1 Detailed Description

25.13 fuse.h File Reference

25.13.1 Detailed Description

25.14 interrupt.h File Reference

Defines

Global manipulation of the interrupt flag

The global interrupt flag is maintained in the I bit of the status register (SREG).

Handling interrupts frequently requires attention regarding atomic access to objects that could be altered by code running within an interrupt context, see [<util/atomic.h>](#).

Frequently, interrupts are being disabled for periods of time in order to perform certain operations without being disturbed; see [Problems with reordering code](#) for things to be taken into account with respect to compiler optimizations.

- #define [sei](#)()
- #define [cli](#)()

Macros for writing interrupt handler functions

- #define [ISR](#)(vector, attributes)
- #define [SIGNAL](#)(vector)
- #define [EMPTY_INTERRUPT](#)(vector)
- #define [ISR_ALIAS](#)(vector, target_vector)
- #define [reti](#)()
- #define [BADISR_vect](#)

ISR attributes

- #define [ISR_BLOCK](#)
- #define [ISR_NOBLOCK](#)
- #define [ISR_NAKED](#)
- #define [ISR_ALIASOF](#)(target_vector)

25.14.1 Detailed Description

@{

25.15 inttypes.h File Reference

Defines

macros for printf and scanf format specifiers

For C++, these are only included if `__STDC_LIMIT_MACROS` is defined before including `<inttypes.h>`.

- #define [PRId8](#) "d"
- #define [PRIdLEAST8](#) "d"
- #define [PRIdFAST8](#) "d"
- #define [PRIi8](#) "i"
- #define [PRIiLEAST8](#) "i"
- #define [PRIiFAST8](#) "i"
- #define [PRId16](#) "d"
- #define [PRIdLEAST16](#) "d"
- #define [PRIdFAST16](#) "d"
- #define [PRIi16](#) "i"
- #define [PRIiLEAST16](#) "i"
- #define [PRIiFAST16](#) "i"
- #define [PRId32](#) "ld"
- #define [PRIdLEAST32](#) "ld"
- #define [PRIdFAST32](#) "ld"
- #define [PRIi32](#) "li"
- #define [PRIiLEAST32](#) "li"
- #define [PRIiFAST32](#) "li"
- #define [PRIdPTR](#) [PRId16](#)
- #define [PRIiPTR](#) [PRIi16](#)
- #define [PRIo8](#) "o"
- #define [PRIoLEAST8](#) "o"
- #define [PRIoFAST8](#) "o"
- #define [PRIu8](#) "u"
- #define [PRIuLEAST8](#) "u"
- #define [PRIuFAST8](#) "u"
- #define [PRIx8](#) "x"
- #define [PRIxLEAST8](#) "x"
- #define [PRIxFAST8](#) "x"
- #define [PRIX8](#) "X"
- #define [PRIXLEAST8](#) "X"
- #define [PRIXFAST8](#) "X"
- #define [PRIo16](#) "o"
- #define [PRIoLEAST16](#) "o"
- #define [PRIoFAST16](#) "o"

- #define [PRIu16](#) "u"
- #define [PRIuLEAST16](#) "u"
- #define [PRIuFAST16](#) "u"
- #define [PRIx16](#) "x"
- #define [PRIxLEAST16](#) "x"
- #define [PRIxFAST16](#) "x"
- #define [PRIX16](#) "X"
- #define [PRIXLEAST16](#) "X"
- #define [PRIXFAST16](#) "X"
- #define [PRIo32](#) "lo"
- #define [PRIoLEAST32](#) "lo"
- #define [PRIoFAST32](#) "lo"
- #define [PRIu32](#) "lu"
- #define [PRIuLEAST32](#) "lu"
- #define [PRIuFAST32](#) "lu"
- #define [PRIx32](#) "lx"
- #define [PRIxLEAST32](#) "lx"
- #define [PRIxFAST32](#) "lx"
- #define [PRIX32](#) "IX"
- #define [PRIXLEAST32](#) "IX"
- #define [PRIXFAST32](#) "IX"
- #define [PRIoPTR](#) PRIo16
- #define [PRIuPTR](#) PRIu16
- #define [PRIxPTR](#) PRIx16
- #define [PRIXPTR](#) PRIX16
- #define [SCNd8](#) "hhd"
- #define [SCNdLEAST8](#) "hhd"
- #define [SCNdFAST8](#) "hhd"
- #define [SCNi8](#) "hhi"
- #define [SCNiLEAST8](#) "hhi"
- #define [SCNiFAST8](#) "hhi"
- #define [SCNd16](#) "d"
- #define [SCNdLEAST16](#) "d"
- #define [SCNdFAST16](#) "d"
- #define [SCNi16](#) "i"
- #define [SCNiLEAST16](#) "i"
- #define [SCNiFAST16](#) "i"
- #define [SCNd32](#) "ld"
- #define [SCNdLEAST32](#) "ld"
- #define [SCNdFAST32](#) "ld"
- #define [SCNi32](#) "li"
- #define [SCNiLEAST32](#) "li"
- #define [SCNiFAST32](#) "li"
- #define [SCNdPTR](#) SCNd16
- #define [SCNiPTR](#) SCNi16
- #define [SCNo8](#) "hho"
- #define [SCNoLEAST8](#) "hho"
- #define [SCNoFAST8](#) "hho"
- #define [SCNu8](#) "hhu"
- #define [SCNuLEAST8](#) "hhu"
- #define [SCNuFAST8](#) "hhu"
- #define [SCNx8](#) "hhx"
- #define [SCNxLEAST8](#) "hhx"
- #define [SCNxFAST8](#) "hhx"
- #define [SCNo16](#) "o"

- #define [SCNoLEAST16](#) "o"
- #define [SCNoFAST16](#) "o"
- #define [SCNu16](#) "u"
- #define [SCNuLEAST16](#) "u"
- #define [SCNuFAST16](#) "u"
- #define [SCNx16](#) "x"
- #define [SCNxLEAST16](#) "x"
- #define [SCNxFAST16](#) "x"
- #define [SCNo32](#) "lo"
- #define [SCNoLEAST32](#) "lo"
- #define [SCNoFAST32](#) "lo"
- #define [SCNu32](#) "lu"
- #define [SCNuLEAST32](#) "lu"
- #define [SCNuFAST32](#) "lu"
- #define [SCNx32](#) "lx"
- #define [SCNxLEAST32](#) "lx"
- #define [SCNxFAST32](#) "lx"
- #define [SCNoPTR](#) SCNo16
- #define [SCNuPTR](#) SCNu16
- #define [SCNxPTR](#) SCNx16

Typedefs

Far pointers for memory access >64K

- typedef [int32_t](#) [int_farptr_t](#)
- typedef [uint32_t](#) [uint_farptr_t](#)

25.15.1 Detailed Description

25.16 io.h File Reference

25.16.1 Detailed Description

25.17 lock.h File Reference

25.17.1 Detailed Description

25.18 math.h File Reference

Defines

- #define [M_E](#) 2.7182818284590452354
- #define [M_LOG2E](#) 1.4426950408889634074
- #define [M_LOG10E](#) 0.43429448190325182765
- #define [M_LN2](#) 0.69314718055994530942
- #define [M_LN10](#) 2.30258509299404568402
- #define [M_PI](#) 3.14159265358979323846
- #define [M_PI_2](#) 1.57079632679489661923

- #define [M_PI_4](#) 0.78539816339744830962
- #define [M_1_PI](#) 0.31830988618379067154
- #define [M_2_PI](#) 0.63661977236758134308
- #define [M_2_SQRTPI](#) 1.12837916709551257390
- #define [M_SQRT2](#) 1.41421356237309504880
- #define [M_SQRT1_2](#) 0.70710678118654752440
- #define [NAN](#) __builtin_nan("")
- #define [INFINITY](#) __builtin_inf()
- #define [cosf](#) cos
- #define [sinf](#) sin
- #define [tanf](#) tan
- #define [fabsf](#) fabs
- #define [fmodf](#) fmod
- #define [cbrtf](#) cbrt
- #define [hypotf](#) hypot
- #define [squaref](#) square
- #define [floorf](#) floor
- #define [ceilf](#) ceil
- #define [frexpf](#) frexp
- #define [ldexpf](#) ldexp
- #define [expf](#) exp
- #define [coshf](#) cosh
- #define [sinhf](#) sinh
- #define [tanhf](#) tanh
- #define [acosf](#) acos
- #define [asinf](#) asin
- #define [atanf](#) atan
- #define [atan2f](#) atan2
- #define [logf](#) log
- #define [log10f](#) log10
- #define [powf](#) pow
- #define [isnanf](#) isnan
- #define [isinf](#) isinf
- #define [isfinitef](#) isfinite
- #define [copysignf](#) copysign
- #define [signbitf](#) signbit
- #define [fdimf](#) fdim
- #define [fmaf](#) fma
- #define [fmaxf](#) fmax
- #define [fminf](#) fmin
- #define [truncf](#) trunc
- #define [roundf](#) round
- #define [lroundf](#) lround
- #define [lrintf](#) lrint

Functions

- double [cos](#) (double __x)
- double [sin](#) (double __x)
- double [tan](#) (double __x)
- double [fabs](#) (double __x)
- double [fmod](#) (double __x, double __y)
- double [modf](#) (double __x, double *__iptr)
- float [modff](#) (float __x, float *__iptr)
- double [sqrt](#) (double __x)
- float [sqrtf](#) (float)
- double [cbrt](#) (double __x)
- double [hypot](#) (double __x, double __y)
- double [square](#) (double __x)
- double [floor](#) (double __x)
- double [ceil](#) (double __x)
- double [frexp](#) (double __x, int *__pexp)
- double [ldexp](#) (double __x, int __exp)
- double [exp](#) (double __x)
- double [cosh](#) (double __x)
- double [sinh](#) (double __x)
- double [tanh](#) (double __x)
- double [acos](#) (double __x)
- double [asin](#) (double __x)
- double [atan](#) (double __x)
- double [atan2](#) (double __y, double __x)
- double [log](#) (double __x)
- double [log10](#) (double __x)
- double [pow](#) (double __x, double __y)
- int [isnan](#) (double __x)
- int [isinf](#) (double __x)
- static int [isfinite](#) (double __x)
- static double [copysign](#) (double __x, double __y)
- int [signbit](#) (double __x)
- double [fdim](#) (double __x, double __y)
- double [fma](#) (double __x, double __y, double __z)
- double [fmax](#) (double __x, double __y)
- double [fmin](#) (double __x, double __y)
- double [trunc](#) (double __x)
- double [round](#) (double __x)
- long [lround](#) (double __x)
- long [lrint](#) (double __x)

25.18.1 Detailed Description

25.19 parity.h File Reference

Defines

- #define `parity_even_bit`(val)

25.19.1 Detailed Description

25.20 pgmspace.h File Reference

Defines

- #define `PROGMEM` `__ATTR_PROGMEM__`
- #define `PGM_P` `const char *`
- #define `PGM_VOID_P` `const void *`
- #define `PSTR`(s) `((const PROGMEM char *) (s))`
- #define `pgm_read_byte_near`(address_short) `__LPM__((uint16_t)(address_short))`
- #define `pgm_read_word_near`(address_short) `__LPM_word__((uint16_t)(address_short))`
- #define `pgm_read_dword_near`(address_short) `__LPM_dword__((uint16_t)(address_short))`
- #define `pgm_read_float_near`(address_short) `__LPM_float__((uint16_t)(address_short))`
- #define `pgm_read_ptr_near`(address_short) `(void*)__LPM_word__((uint16_t)(address_short))`
- #define `pgm_read_byte_far`(address_long) `__ELPM__((uint32_t)(address_long))`
- #define `pgm_read_word_far`(address_long) `__ELPM_word__((uint32_t)(address_long))`
- #define `pgm_read_dword_far`(address_long) `__ELPM_dword__((uint32_t)(address_long))`
- #define `pgm_read_float_far`(address_long) `__ELPM_float__((uint32_t)(address_long))`
- #define `pgm_read_ptr_far`(address_long) `(void*)__ELPM_word__((uint32_t)(address_long))`
- #define `pgm_read_byte`(address_short) `pgm_read_byte_near(address_short)`
- #define `pgm_read_word`(address_short) `pgm_read_word_near(address_short)`
- #define `pgm_read_dword`(address_short) `pgm_read_dword_near(address_short)`
- #define `pgm_read_float`(address_short) `pgm_read_float_near(address_short)`
- #define `pgm_read_ptr`(address_short) `pgm_read_ptr_near(address_short)`
- #define `pgm_get_far_address`(var)

Typedefs

- typedef void PROGMEM [prog_void](#)
- typedef char PROGMEM [prog_char](#)
- typedef unsigned char PROGMEM [prog_uchar](#)
- typedef int8_t PROGMEM [prog_int8_t](#)
- typedef uint8_t PROGMEM [prog_uint8_t](#)
- typedef int16_t PROGMEM [prog_int16_t](#)
- typedef uint16_t PROGMEM [prog_uint16_t](#)
- typedef int32_t PROGMEM [prog_int32_t](#)
- typedef uint32_t PROGMEM [prog_uint32_t](#)
- typedef int64_t PROGMEM [prog_int64_t](#)
- typedef uint64_t PROGMEM [prog_uint64_t](#)

Functions

- const void * [memchr_P](#) (const void *, int __val, size_t __len)
- int [memcmp_P](#) (const void *, const void *, size_t) __ATTR_PURE__
- void * [memcpy_P](#) (void *, const void *, int __val, size_t)
- void * [memcpy_P](#) (void *, const void *, size_t)
- void * [memcpy_P](#) (const void *, size_t, const void *, size_t) __ATTR_PURE__
- const void * [memrchr_P](#) (const void *, int __val, size_t __len)
- char * [strcat_P](#) (char *, const char *)
- const char * [strchr_P](#) (const char *, int __val)
- const char * [strchrnul_P](#) (const char *, int __val)
- int [strcmp_P](#) (const char *, const char *) __ATTR_PURE__
- char * [strcpy_P](#) (char *, const char *)
- int [strcasestr_P](#) (const char *, const char *) __ATTR_PURE__
- char * [strcasestr_P](#) (const char *, const char *) __ATTR_PURE__
- size_t [strcspn_P](#) (const char *__s, const char *__reject) __ATTR_PURE__
- size_t [strlcat_P](#) (char *, const char *, size_t)
- size_t [strncpy_P](#) (char *, const char *, size_t)
- size_t [strlen_P](#) (const char *, size_t)
- int [strncmp_P](#) (const char *, const char *, size_t) __ATTR_PURE__
- int [strncasecmp_P](#) (const char *, const char *, size_t) __ATTR_PURE__
- char * [strncat_P](#) (char *, const char *, size_t)
- char * [strncpy_P](#) (char *, const char *, size_t)
- char * [strpbrk_P](#) (const char *__s, const char *__accept) __ATTR_PURE__
- const char * [strrchr_P](#) (const char *, int __val)
- char * [strsep_P](#) (char **__sp, const char *__delim)
- size_t [strspn_P](#) (const char *__s, const char *__accept) __ATTR_PURE__
- char * [strstr_P](#) (const char *, const char *) __ATTR_PURE__
- char * [strtok_P](#) (char *__s, const char *__delim)
- char * [strtok_r_P](#) (char *__s, const char *__delim, char **__last)
- size_t [strlen_PF](#) (uint_farptr_t src)

- `size_t strlen_PF (uint_farptr_t src, size_t len)`
- `void * memcpy_PF (void *dest, uint_farptr_t src, size_t len)`
- `char * strcpy_PF (char *dest, uint_farptr_t src)`
- `char * strncpy_PF (char *dest, uint_farptr_t src, size_t len)`
- `char * strcat_PF (char *dest, uint_farptr_t src)`
- `size_t strlcat_PF (char *dst, uint_farptr_t src, size_t siz)`
- `char * strncat_PF (char *dest, uint_farptr_t src, size_t len)`
- `int strcmp_PF (const char *s1, uint_farptr_t s2) __ATTR_PURE__`
- `int strncmp_PF (const char *s1, uint_farptr_t s2, size_t n) __ATTR_PURE__`
- `int strcasecmp_PF (const char *s1, uint_farptr_t s2) __ATTR_PURE__`
- `int strncasecmp_PF (const char *s1, uint_farptr_t s2, size_t n) __ATTR_PURE__`
- `char * strstr_PF (const char *s1, uint_farptr_t s2)`
- `size_t strlcpy_PF (char *dst, uint_farptr_t src, size_t siz)`
- `int memcmp_PF (const void *, uint_farptr_t, size_t) __ATTR_PURE__`
- `static size_t strlen_P (const char *s)`

25.20.1 Detailed Description

25.21 power.h File Reference

Defines

- `#define clock_prescale_get() (clock_div_t)(CLKPR & (uint8_t)((1<<CLKPS0)|(1<<CLKPS1)|(1<<CLKPS2)|(1<<CLKPS3)))`

Functions

- `static __inline void __attribute__((__always_inline__)) __power_all_enable()`
- `void clock_prescale_set (clock_div_t __x)`

25.21.1 Detailed Description

25.21.2 Define Documentation

25.21.2.1 `#define clock_prescale_get() (clock_div_t)(CLKPR & (uint8_t)((1<<CLKPS0)|(1<<CLKPS1)|(1<<CLKPS2)|(1<<CLKPS3)))`

Gets and returns the clock prescaler register setting. The return type is `clock_div_t`.

Note

For device with XTAL Divide Control Register (XDIV), return can actually range from 1 to 129. Care should be taken has the return value could differ from the typedef enum `clock_div_t`. This should only happen if `clock_prescale_set` was previously called with a value other than those defined by `clock_div_t`.

25.22 setbaud.h File Reference

Defines

- #define [BAUD_TOL](#) 2
- #define [UBRR_VALUE](#)
- #define [UBRRL_VALUE](#)
- #define [UBRRH_VALUE](#)
- #define [USE_2X](#) 0

25.22.1 Detailed Description

25.23 setjmp.h File Reference

Functions

- int [setjmp](#) (jmp_buf __jmpb)
- void [longjmp](#) (jmp_buf __jmpb, int __ret) [__ATTR_NORETURN__](#)

25.23.1 Detailed Description

25.24 signature.h File Reference

25.24.1 Detailed Description

25.25 sleep.h File Reference

Functions

- void [sleep_enable](#) (void)
- void [sleep_disable](#) (void)
- void [sleep_cpu](#) (void)
- void [sleep_mode](#) (void)
- void [sleep_bod_disable](#) (void)

25.25.1 Detailed Description

25.26 stdint.h File Reference

Defines

Limits of specified-width integer types

C++ implementations should define these macros only when `__STDC_LIMIT_MACROS` is defined before `<stdint.h>` is included

- #define [INT8_MAX](#) 0x7f

- #define `INT8_MIN` $(-INT8_MAX - 1)$
- #define `UINT8_MAX` $(INT8_MAX * 2 + 1)$
- #define `INT16_MAX` `0x7fff`
- #define `INT16_MIN` $(-INT16_MAX - 1)$
- #define `UINT16_MAX` $(_CONCAT(INT16_MAX, U) * 2U + 1U)$
- #define `INT32_MAX` `0x7fffffffL`
- #define `INT32_MIN` $(-INT32_MAX - 1L)$
- #define `UINT32_MAX` $(_CONCAT(INT32_MAX, U) * 2UL + 1UL)$
- #define `INT64_MAX` `0x7fffffffffffffffLL`
- #define `INT64_MIN` $(-INT64_MAX - 1LL)$
- #define `UINT64_MAX` $(_CONCAT(INT64_MAX, U) * 2ULL + 1ULL)$

Limits of minimum-width integer types

- #define `INT_LEAST8_MAX` `INT8_MAX`
- #define `INT_LEAST8_MIN` `INT8_MIN`
- #define `UINT_LEAST8_MAX` `UINT8_MAX`
- #define `INT_LEAST16_MAX` `INT16_MAX`
- #define `INT_LEAST16_MIN` `INT16_MIN`
- #define `UINT_LEAST16_MAX` `UINT16_MAX`
- #define `INT_LEAST32_MAX` `INT32_MAX`
- #define `INT_LEAST32_MIN` `INT32_MIN`
- #define `UINT_LEAST32_MAX` `UINT32_MAX`
- #define `INT_LEAST64_MAX` `INT64_MAX`
- #define `INT_LEAST64_MIN` `INT64_MIN`
- #define `UINT_LEAST64_MAX` `UINT64_MAX`

Limits of fastest minimum-width integer types

- #define `INT_FAST8_MAX` `INT8_MAX`
- #define `INT_FAST8_MIN` `INT8_MIN`
- #define `UINT_FAST8_MAX` `UINT8_MAX`
- #define `INT_FAST16_MAX` `INT16_MAX`
- #define `INT_FAST16_MIN` `INT16_MIN`
- #define `UINT_FAST16_MAX` `UINT16_MAX`
- #define `INT_FAST32_MAX` `INT32_MAX`
- #define `INT_FAST32_MIN` `INT32_MIN`
- #define `UINT_FAST32_MAX` `UINT32_MAX`
- #define `INT_FAST64_MAX` `INT64_MAX`
- #define `INT_FAST64_MIN` `INT64_MIN`
- #define `UINT_FAST64_MAX` `UINT64_MAX`

Limits of integer types capable of holding object pointers

- #define `INTPTR_MAX` `INT16_MAX`
- #define `INTPTR_MIN` `INT16_MIN`
- #define `UINTPTR_MAX` `UINT16_MAX`

Limits of greatest-width integer types

- #define `INTMAX_MAX` `INT64_MAX`
- #define `INTMAX_MIN` `INT64_MIN`
- #define `UINTMAX_MAX` `UINT64_MAX`

Limits of other integer types

C++ implementations should define these macros only when `__STDC_LIMIT_MACROS` is defined before `<stdint.h>` is included

- `#define PTRDIFF_MAX INT16_MAX`
- `#define PTRDIFF_MIN INT16_MIN`
- `#define SIG_ATOMIC_MAX INT8_MAX`
- `#define SIG_ATOMIC_MIN INT8_MIN`
- `#define SIZE_MAX UINT16_MAX`
- `#define WCHAR_MAX __WCHAR_MAX__`
- `#define WCHAR_MIN __WCHAR_MIN__`
- `#define WINT_MAX __WINT_MAX__`
- `#define WINT_MIN __WINT_MIN__`

Macros for integer constants

C++ implementations should define these macros only when `__STDC_CONSTANT_MACROS` is defined before `<stdint.h>` is included.

These definitions are valid for integer constants without suffix and for macros defined as integer constant without suffix

- `#define INT8_C(value) ((int8_t) value)`
- `#define UINT8_C(value) ((uint8_t) __CONCAT(value, U))`
- `#define INT16_C(value) value`
- `#define UINT16_C(value) __CONCAT(value, U)`
- `#define INT32_C(value) __CONCAT(value, L)`
- `#define UINT32_C(value) __CONCAT(value, UL)`
- `#define INT64_C(value) __CONCAT(value, LL)`
- `#define UINT64_C(value) __CONCAT(value, ULL)`
- `#define INTMAX_C(value) __CONCAT(value, LL)`
- `#define UINTMAX_C(value) __CONCAT(value, ULL)`

Typedefs**Exact-width integer types**

Integer types having exactly the specified width

- `typedef signed char int8_t`
- `typedef unsigned char uint8_t`
- `typedef signed int int16_t`
- `typedef unsigned int uint16_t`
- `typedef signed long int int32_t`
- `typedef unsigned long int uint32_t`
- `typedef signed long long int int64_t`
- `typedef unsigned long long int uint64_t`

Integer types capable of holding object pointers

These allow you to declare variables of the same size as a pointer.

- `typedef int16_t intptr_t`
- `typedef uint16_t uintptr_t`

Minimum-width integer types*Integer types having at least the specified width*

- typedef [int8_t](#) [int_least8_t](#)
- typedef [uint8_t](#) [uint_least8_t](#)
- typedef [int16_t](#) [int_least16_t](#)
- typedef [uint16_t](#) [uint_least16_t](#)
- typedef [int32_t](#) [int_least32_t](#)
- typedef [uint32_t](#) [uint_least32_t](#)
- typedef [int64_t](#) [int_least64_t](#)
- typedef [uint64_t](#) [uint_least64_t](#)

Fastest minimum-width integer types*Integer types being usually fastest having at least the specified width*

- typedef [int8_t](#) [int_fast8_t](#)
- typedef [uint8_t](#) [uint_fast8_t](#)
- typedef [int16_t](#) [int_fast16_t](#)
- typedef [uint16_t](#) [uint_fast16_t](#)
- typedef [int32_t](#) [int_fast32_t](#)
- typedef [uint32_t](#) [uint_fast32_t](#)
- typedef [int64_t](#) [int_fast64_t](#)
- typedef [uint64_t](#) [uint_fast64_t](#)

Greatest-width integer types*Types designating integer data capable of representing any value of any integer type in the corresponding signed or unsigned category*

- typedef [int64_t](#) [intmax_t](#)
- typedef [uint64_t](#) [uintmax_t](#)

25.26.1 Detailed Description**25.27 stdio.h File Reference****Defines**

- #define [stdin](#) ([__iob](#)[0])
- #define [stdout](#) ([__iob](#)[1])
- #define [stderr](#) ([__iob](#)[2])
- #define [EOF](#) (-1)
- #define [fdev_set_udata](#)(stream, u) do { (stream)->udata = u; } while(0)
- #define [fdev_get_udata](#)(stream) ((stream)->udata)
- #define [fdev_setup_stream](#)(stream, put, get, rflag)
- #define [_FDEV_SETUP_READ](#) [__SRD](#)
- #define [_FDEV_SETUP_WRITE](#) [__SWR](#)
- #define [_FDEV_SETUP_RW](#) ([__SRD](#)|[__SWR](#))
- #define [_FDEV_ERR](#) (-1)
- #define [_FDEV_EOF](#) (-2)

- #define [FDEV_SETUP_STREAM](#)(put, get, rwflag)
- #define [fdev_close](#)()
- #define [putc](#)(__c, __stream) [fputc](#)(__c, __stream)
- #define [putchar](#)(__c) [fputc](#)(__c, stdout)
- #define [getc](#)(__stream) [fgetc](#)(__stream)
- #define [getchar](#)() [fgetc](#)(stdin)

Typedefs

- typedef struct __file [FILE](#)

Functions

- int [fclose](#) ([FILE](#) *__stream)
- int [vfprintf](#) ([FILE](#) *__stream, const char *__fmt, va_list __ap)
- int [vfprintf_P](#) ([FILE](#) *__stream, const char *__fmt, va_list __ap)
- int [fputc](#) (int __c, [FILE](#) *__stream)
- int [printf](#) (const char *__fmt,...)
- int [printf_P](#) (const char *__fmt,...)
- int [vprintf](#) (const char *__fmt, va_list __ap)
- int [sprintf](#) (char *__s, const char *__fmt,...)
- int [sprintf_P](#) (char *__s, const char *__fmt,...)
- int [snprintf](#) (char *__s, size_t __n, const char *__fmt,...)
- int [snprintf_P](#) (char *__s, size_t __n, const char *__fmt,...)
- int [vsprintf](#) (char *__s, const char *__fmt, va_list ap)
- int [vsprintf_P](#) (char *__s, const char *__fmt, va_list ap)
- int [vsnprintf](#) (char *__s, size_t __n, const char *__fmt, va_list ap)
- int [vsnprintf_P](#) (char *__s, size_t __n, const char *__fmt, va_list ap)
- int [fprintf](#) ([FILE](#) *__stream, const char *__fmt,...)
- int [fprintf_P](#) ([FILE](#) *__stream, const char *__fmt,...)
- int [fputs](#) (const char *__str, [FILE](#) *__stream)
- int [fputs_P](#) (const char *__str, [FILE](#) *__stream)
- int [puts](#) (const char *__str)
- int [puts_P](#) (const char *__str)
- size_t [fwrite](#) (const void *__ptr, size_t __size, size_t __nmemb, [FILE](#) *__stream)
- int [fgetc](#) ([FILE](#) *__stream)
- int [ungetc](#) (int __c, [FILE](#) *__stream)
- char * [fgets](#) (char *__str, int __size, [FILE](#) *__stream)
- char * [gets](#) (char *__str)
- size_t [fread](#) (void *__ptr, size_t __size, size_t __nmemb, [FILE](#) *__stream)
- void [clearerr](#) ([FILE](#) *__stream)
- int [feof](#) ([FILE](#) *__stream)
- int [ferror](#) ([FILE](#) *__stream)
- int [vfscanf](#) ([FILE](#) *__stream, const char *__fmt, va_list __ap)
- int [vfscanf_P](#) ([FILE](#) *__stream, const char *__fmt, va_list __ap)

- int [fscanf](#) ([FILE](#) *__stream, const char *__fmt,...)
- int [fscanf_P](#) ([FILE](#) *__stream, const char *__fmt,...)
- int [scanf](#) (const char *__fmt,...)
- int [scanf_P](#) (const char *__fmt,...)
- int [vscanf](#) (const char *__fmt, va_list __ap)
- int [sscanf](#) (const char *__buf, const char *__fmt,...)
- int [sscanf_P](#) (const char *__buf, const char *__fmt,...)
- int [fflush](#) ([FILE](#) *stream)

25.27.1 Detailed Description

25.28 `stdlib.h` File Reference

Data Structures

- struct [div_t](#)
- struct [ldiv_t](#)

Defines

- #define [RAND_MAX](#) 0x7FFF

Typedefs

- typedef int(* [__compar_fn_t](#))(const void *, const void *)

Functions

- void [abort](#) (void) [__ATTR_NORETURN__](#)
- int [abs](#) (int __i)
- long [labs](#) (long __i)
- void * [bsearch](#) (const void *__key, const void *__base, size_t __nmemb, size_t __size, int(*__compar)(const void *, const void *))
- [div_t](#) [div](#) (int __num, int __denom) [__asm__](#)("__divmodhi4")
- [ldiv_t](#) [ldiv](#) (long __num, long __denom) [__asm__](#)("__divmodsi4")
- void [qsort](#) (void *__base, size_t __nmemb, size_t __size, [__compar_fn_t](#) __compar)
- long [strtol](#) (const char *__nptr, char **__endptr, int __base)
- unsigned long [strtoul](#) (const char *__nptr, char **__endptr, int __base)
- long [atol](#) (const char *__s) [__ATTR_PURE__](#)
- int [atoi](#) (const char *__s) [__ATTR_PURE__](#)
- void [exit](#) (int __status) [__ATTR_NORETURN__](#)
- void * [malloc](#) (size_t __size) [__ATTR_MALLOC__](#)
- void [free](#) (void *__ptr)
- void * [calloc](#) (size_t __nele, size_t __size) [__ATTR_MALLOC__](#)

- void * [realloc](#) (void *__ptr, size_t __size) __ATTR_MALLOC__
- double [strtod](#) (const char *__nptr, char **__endptr)
- double [atof](#) (const char *__nptr)
- int [rand](#) (void)
- void [srand](#) (unsigned int __seed)
- int [rand_r](#) (unsigned long *__ctx)

Variables

- size_t [__malloc_margin](#)
- char * [__malloc_heap_start](#)
- char * [__malloc_heap_end](#)

Non-standard (i.e. non-ISO C) functions.

- #define [RANDOM_MAX](#) 0x7FFFFFFF
- char * [itoa](#) (int val, char *s, int radix)
- char * [ltoa](#) (long val, char *s, int radix)
- char * [utoa](#) (unsigned int val, char *s, int radix)
- char * [ultoa](#) (unsigned long val, char *s, int radix)
- long [random](#) (void)
- void [srandom](#) (unsigned long __seed)
- long [random_r](#) (unsigned long *__ctx)

Conversion functions for double arguments.

Note that these functions are not located in the default library, `libc.a`, but in the mathematical library, `libm.a`. So when linking the application, the `-lm` option needs to be specified.

- #define [DOSTR_ALWAYS_SIGN](#) 0x01
- #define [DOSTR_PLUS_SIGN](#) 0x02
- #define [DOSTR_UPPERCASE](#) 0x04
- #define [EXIT_SUCCESS](#) 0
- #define [EXIT_FAILURE](#) 1
- char * [dostre](#) (double __val, char *__s, unsigned char __prec, unsigned char __flags)
- char * [dostrf](#) (double __val, signed char __width, unsigned char __prec, char *__s)

25.28.1 Detailed Description

25.29 string.h File Reference

Defines

- #define [_FFS\(x\)](#)

Functions

- int [ffs](#) (int __val)
- int [ffsl](#) (long __val)
- __extension__ int [ffsll](#) (long long __val)
- void * [memccpy](#) (void *, const void *, int, size_t)
- void * [memchr](#) (const void *, int, size_t) __ATTR_PURE__
- int [memcmp](#) (const void *, const void *, size_t) __ATTR_PURE__
- void * [memcpy](#) (void *, const void *, size_t)
- void * [memmem](#) (const void *, size_t, const void *, size_t) __ATTR_PURE__
- void * [memmove](#) (void *, const void *, size_t)
- void * [memrchr](#) (const void *, int, size_t) __ATTR_PURE__
- void * [memset](#) (void *, int, size_t)
- char * [strcat](#) (char *, const char *)
- char * [strchr](#) (const char *, int) __ATTR_PURE__
- char * [strchrnul](#) (const char *, int) __ATTR_PURE__
- int [strcmp](#) (const char *, const char *) __ATTR_PURE__
- char * [strcpy](#) (char *, const char *)
- int [strcasecmp](#) (const char *, const char *) __ATTR_PURE__
- char * [strcasestr](#) (const char *, const char *) __ATTR_PURE__
- size_t [strcspn](#) (const char *__s, const char *__reject) __ATTR_PURE__
- char * [strdup](#) (const char *s1)
- size_t [strlcat](#) (char *, const char *, size_t)
- size_t [strlcpy](#) (char *, const char *, size_t)
- size_t [strlen](#) (const char *) __ATTR_PURE__
- char * [strlwr](#) (char *)
- char * [strncat](#) (char *, const char *, size_t)
- int [strncmp](#) (const char *, const char *, size_t) __ATTR_PURE__
- char * [strncpy](#) (char *, const char *, size_t)
- int [strncasecmp](#) (const char *, const char *, size_t) __ATTR_PURE__
- size_t [strnlen](#) (const char *, size_t) __ATTR_PURE__
- char * [strpbrk](#) (const char *__s, const char *__accept) __ATTR_PURE__
- char * [strrchr](#) (const char *, int) __ATTR_PURE__
- char * [strrev](#) (char *)
- char * [strsep](#) (char **, const char *)
- size_t [strspn](#) (const char *__s, const char *__accept) __ATTR_PURE__
- char * [strstr](#) (const char *, const char *) __ATTR_PURE__
- char * [strtok](#) (char *, const char *)
- char * [strtok_r](#) (char *, const char *, char **)
- char * [strupr](#) (char *)

25.29.1 Detailed Description

25.30 time.h File Reference

Data Structures

- struct [tm](#)
- struct [week_date](#)

Defines

- `#define ONE_HOUR 3600`
- `#define ONE_DEGREE 3600`
- `#define ONE_DAY 86400`
- `#define UNIX_OFFSET 946684800`
- `#define NTP_OFFSET 3155673600`

Typedefs

- `typedef uint32_t time_t`

Enumerations

- `enum _WEEK_DAYS_ {
 SUNDAY, MONDAY, TUESDAY, WEDNESDAY,
 THURSDAY, FRIDAY, SATURDAY }`
- `enum _MONTHS_ {
 JANUARY, FEBRUARY, MARCH, APRIL,
 MAY, JUNE, JULY, AUGUST,
 SEPTEMBER, OCTOBER, NOVEMBER, DECEMBER }`

Functions

- `time_t time (time_t *timer)`
- `int32_t difftime (time_t time1, time_t time0)`
- `time_t mktime (struct tm *timeptr)`
- `time_t mk_gmtime (const struct tm *timeptr)`
- `struct tm * gmtime (const time_t *timer)`
- `void gmtime_r (const time_t *timer, struct tm *timeptr)`
- `struct tm * localtime (const time_t *timer)`
- `void localtime_r (const time_t *timer, struct tm *timeptr)`
- `char * asctime (const struct tm *timeptr)`
- `void asctime_r (const struct tm *timeptr, char *buf)`
- `char * ctime (const time_t *timer)`
- `void ctime_r (const time_t *timer, char *buf)`
- `char * isotime (const struct tm *timeptr)`
- `void isotime_r (const struct tm *, char *)`
- `size_t strftime (char *s, size_t maxsize, const char *format, const struct tm *timeptr)`
- `void set_dst (int(*) (const time_t *, int32_t *))`
- `void set_zone (int32_t)`
- `void set_system_time (time_t timestamp)`
- `void system_tick (void)`
- `uint8_t is_leap_year (int16_t year)`

- `uint8_t month_length` (`int16_t` year, `uint8_t` month)
- `uint8_t week_of_year` (`const struct tm *timeptr`, `uint8_t` start)
- `uint8_t week_of_month` (`const struct tm *timeptr`, `uint8_t` start)
- `struct week_date * iso_week_date` (`int` year, `int` yday)
- `void iso_week_date_r` (`int` year, `int` yday, `struct week_date *`)
- `uint32_t fatfs_time` (`const struct tm *timeptr`)
- `void set_position` (`int32_t` latitude, `int32_t` longitude)
- `int16_t equation_of_time` (`const time_t *timer`)
- `int32_t daylight_seconds` (`const time_t *timer`)
- `time_t solar_noon` (`const time_t *timer`)
- `time_t sun_rise` (`const time_t *timer`)
- `time_t sun_set` (`const time_t *timer`)
- `double solar_declination` (`const time_t *timer`)
- `int8_t moon_phase` (`const time_t *timer`)
- `unsigned long gm_sidereal` (`const time_t *timer`)
- `unsigned long lm_sidereal` (`const time_t *timer`)

25.30.1 Detailed Description

25.31 twi.h File Reference

Defines

TWSR values

Mnemonics:

TW_MT_xxx - master transmitter

TW_MR_xxx - master receiver

TW_ST_xxx - slave transmitter

TW_SR_xxx - slave receiver

- `#define TW_START` 0x08
- `#define TW_REP_START` 0x10
- `#define TW_MT_SLA_ACK` 0x18
- `#define TW_MT_SLA_NACK` 0x20
- `#define TW_MT_DATA_ACK` 0x28
- `#define TW_MT_DATA_NACK` 0x30
- `#define TW_MT_ARB_LOST` 0x38
- `#define TW_MR_ARB_LOST` 0x38
- `#define TW_MR_SLA_ACK` 0x40
- `#define TW_MR_SLA_NACK` 0x48
- `#define TW_MR_DATA_ACK` 0x50
- `#define TW_MR_DATA_NACK` 0x58
- `#define TW_ST_SLA_ACK` 0xA8
- `#define TW_ST_ARB_LOST_SLA_ACK` 0xB0
- `#define TW_ST_DATA_ACK` 0xB8
- `#define TW_ST_DATA_NACK` 0xC0
- `#define TW_ST_LAST_DATA` 0xC8
- `#define TW_SR_SLA_ACK` 0x60

- #define [TW_SR_ARB_LOST_SLA_ACK](#) 0x68
- #define [TW_SR_GCALL_ACK](#) 0x70
- #define [TW_SR_ARB_LOST_GCALL_ACK](#) 0x78
- #define [TW_SR_DATA_ACK](#) 0x80
- #define [TW_SR_DATA_NACK](#) 0x88
- #define [TW_SR_GCALL_DATA_ACK](#) 0x90
- #define [TW_SR_GCALL_DATA_NACK](#) 0x98
- #define [TW_SR_STOP](#) 0xA0
- #define [TW_NO_INFO](#) 0xF8
- #define [TW_BUS_ERROR](#) 0x00
- #define [TW_STATUS_MASK](#)
- #define [TW_STATUS](#) (TWSR & TW_STATUS_MASK)

R/~W bit in SLA+R/W address field.

- #define [TW_READ](#) 1
- #define [TW_WRITE](#) 0

25.31.1 Detailed Description

25.32 wdt.h File Reference

Defines

- #define [wdt_reset\(\)](#) __asm__ __volatile__ ("wdr")
- #define [WDTO_15MS](#) 0
- #define [WDTO_30MS](#) 1
- #define [WDTO_60MS](#) 2
- #define [WDTO_120MS](#) 3
- #define [WDTO_250MS](#) 4
- #define [WDTO_500MS](#) 5
- #define [WDTO_1S](#) 6
- #define [WDTO_2S](#) 7
- #define [WDTO_4S](#) 8
- #define [WDTO_8S](#) 9

Functions

- static __inline__ __attribute__((__always_inline__)) void [wdt_enable](#)(const [uint8_t](#) value)

25.32.1 Detailed Description

Index

- <alloca.h>: Allocate space in the stack, [128](#)
- <assert.h>: Diagnostics, [129](#)
- <avr/boot.h>: Bootloader Support Utilities, [225](#)
- <avr/cpufunc.h>: Special AVR CPU functions, [231](#)
- <avr/eeprom.h>: EEPROM handling, [232](#)
- <avr/fuse.h>: Fuse Support, [237](#)
- <avr/interrupt.h>: Interrupts, [240](#)
- <avr/io.h>: AVR device-specific IO definitions, [260](#)
- <avr/lock.h>: Lockbit Support, [261](#)
- <avr/pgmspace.h>: Program Space Utilities, [264](#)
- <avr/power.h>: Power Reduction Management, [288](#)
- <avr/sfr_defs.h>: Special function registers, [293](#)
- <avr/signature.h>: Signature Support, [295](#)
- <avr/sleep.h>: Power Management and Sleep Modes, [296](#)
- <avr/version.h>: avr-libc version macros, [298](#)
- <avr/wdt.h>: Watchdog timer handling, [299](#)
- <compat/deprecated.h>: Deprecated items, [321](#)
- <compat/ina90.h>: Compatibility with IAR EWB 3.x, [324](#)
- <ctype.h>: Character Operations, [130](#)
- <errno.h>: System Errors, [132](#)
- <inttypes.h>: Integer Type conversions, [133](#)
- <math.h>: Mathematics, [147](#)
- <setjmp.h>: Non-local goto, [160](#)
- <stdint.h>: Standard Integer Types, [162](#)
- <stdio.h>: Standard IO facilities, [174](#)
- <stdlib.h>: General utilities, [192](#)
- <string.h>: Strings, [203](#)
- <time.h>: Time, [216](#)
- <util/atomic.h>: Atomically and Non-Atomically Executed Code Blocks, [303](#)
- <util/crc16.h>: CRC Computations, [306](#)
- <util/delay.h>: Convenience functions for busy-wait delay loops, [310](#)
- <util/delay_basic.h>: Basic busy-wait delay loops, [313](#)
- <util/parity.h>: Parity bit generation, [314](#)
- <util/setbaud.h>: Helper macros for baud rate calculations, [314](#)
- <util/twi.h>: TWI bit mask definitions, [317](#)
- \$PATH, [87](#)
- \$PREFIX, [87](#)
- prefix, [87](#)
- _BV
 - avr_sfr, [294](#)
- _EGET
 - avr_eeprom, [234](#)
- _EEPUT
 - avr_eeprom, [234](#)
- _FDEV_EOF
 - avr_stdio, [178](#)
- _FDEV_ERR
 - avr_stdio, [178](#)
- _FDEV_SETUP_READ
 - avr_stdio, [178](#)
- _FDEV_SETUP_RW
 - avr_stdio, [178](#)
- _FDEV_SETUP_WRITE
 - avr_stdio, [178](#)
- _FFS
 - avr_string, [204](#)
- _MONTHS_
 - avr_time, [219](#)
- _MemoryBarrier
 - avr_cpufunc, [232](#)
- _NOP
 - avr_cpufunc, [232](#)
- _PROTECTED_WRITE
 - avr_io, [261](#)
- _PROTECTED_WRITE_SPM
 - avr_io, [261](#)
- _WEEK_DAYS_
 - avr_time, [219](#)
- __AVR_LIBC_DATE_
 - avr_version, [299](#)
- __AVR_LIBC_DATE_STRING__

- avr_version, [299](#)
- __AVR_LIBC_MAJOR__
 - avr_version, [299](#)
- __AVR_LIBC_MINOR__
 - avr_version, [299](#)
- __AVR_LIBC_REVISION__
 - avr_version, [299](#)
- __AVR_LIBC_VERSION_STRING__
 - avr_version, [299](#)
- __AVR_LIBC_VERSION__
 - avr_version, [299](#)
- __EGET
 - avr_eeprom, [234](#)
- __EPUT
 - avr_eeprom, [234](#)
- __attribute__
 - avr_watchdog, [303](#)
- __compar_fn_t
 - avr_stdlib, [194](#)
- __malloc_heap_end
 - avr_stdlib, [202](#)
- __malloc_heap_start
 - avr_stdlib, [202](#)
- __malloc_margin
 - avr_stdlib, [203](#)
- _crc16_update
 - util_crc, [307](#)
- _crc8_ccitt_update
 - util_crc, [308](#)
- _crc_ccitt_update
 - util_crc, [308](#)
- _crc_ibutton_update
 - util_crc, [309](#)
- _crc_xmodem_update
 - util_crc, [310](#)
- _delay_loop_1
 - util_delay_basic, [313](#)
- _delay_loop_2
 - util_delay_basic, [313](#)
- _delay_ms
 - util_delay, [311](#)
- _delay_us
 - util_delay, [312](#)
- A more sophisticated project, [343](#)
- A simple project, [329](#)
- abort
 - avr_stdlib, [194](#)
- abs
 - avr_stdlib, [194](#)
- acos
 - avr_math, [154](#)
- acosf
 - avr_math, [149](#)
- Additional notes from <avr/sfr_defs.h>, [292](#)
- alloca
 - alloca, [128](#)
- asctime
 - avr_time, [220](#)
- asctime_r
 - avr_time, [220](#)
- asin
 - avr_math, [154](#)
- asinf
 - avr_math, [149](#)
- assert
 - avr_assert, [129](#)
- assert.h, [365](#)
- atan
 - avr_math, [154](#)
- atan2
 - avr_math, [154](#)
- atan2f
 - avr_math, [149](#)
- atanf
 - avr_math, [150](#)
- atof
 - avr_stdlib, [194](#)
- atoi
 - avr_stdlib, [195](#)
- atoi.S, [366](#)
- atol
 - avr_stdlib, [195](#)
- atol.S, [366](#)
- atomic.h, [366](#)
- ATOMIC_BLOCK
 - util_atomic, [305](#)
- ATOMIC_FORCEON
 - util_atomic, [305](#)
- ATOMIC_RESTORESTATE
 - util_atomic, [305](#)
- avr_assert
 - assert, [129](#)
- avr_boot
 - boot_is_spm_interrupt, [227](#)
 - boot_lock_bits_set, [227](#)
 - boot_lock_bits_set_safe, [227](#)
 - boot_lock_fuse_bits_get, [227](#)
 - boot_page_erase, [228](#)

- boot_page_erase_safe, 228
- boot_page_fill, 228
- boot_page_fill_safe, 229
- boot_page_write, 229
- boot_page_write_safe, 229
- boot_rww_busy, 229
- boot_rww_enable, 230
- boot_rww_enable_safe, 230
- boot_signature_byte_get, 230
- boot_spm_busy, 230
- boot_spm_busy_wait, 230
- boot_spm_interrupt_disable, 231
- boot_spm_interrupt_enable, 231
- BOOTLOADER_SECTION, 231
- GET_EXTENDED_FUSE_BITS, 231
- GET_HIGH_FUSE_BITS, 231
- GET_LOCK_BITS, 231
- GET_LOW_FUSE_BITS, 231
- avr_cpufunc
 - _MemoryBarrier, 232
 - _NOP, 232
 - ccp_write_io, 232
- avr_eeprom
 - _EEGET, 234
 - _EEPUT, 234
 - __EEGET, 234
 - __EEPUT, 234
 - EEMEM, 234
 - eeprom_busy_wait, 234
 - eeprom_is_ready, 235
 - eeprom_read_block, 235
 - eeprom_read_byte, 235
 - eeprom_read_dword, 235
 - eeprom_read_float, 235
 - eeprom_read_word, 235
 - eeprom_update_block, 235
 - eeprom_update_byte, 235
 - eeprom_update_dword, 236
 - eeprom_update_float, 236
 - eeprom_update_word, 236
 - eeprom_write_block, 236
 - eeprom_write_byte, 236
 - eeprom_write_dword, 236
 - eeprom_write_float, 236
 - eeprom_write_word, 236
- avr_errno
 - EDOM, 133
 - ERANGE, 133
 - errno, 133
- avr_interrupts
 - BADISR_vect, 257
 - cli, 257
 - EMPTY_INTERRUPT, 257
 - ISR, 257
 - ISR_ALIAS, 258
 - ISR_ALIASOF, 258
 - ISR_BLOCK, 258
 - ISR_NAKED, 259
 - ISR_NOBLOCK, 259
 - reti, 259
 - sei, 259
 - SIGNAL, 259
- avr_inttypes
 - int_farptr_t, 147
 - PRId16, 137
 - PRId32, 137
 - PRId8, 137
 - PRIdFAST16, 137
 - PRIdFAST32, 137
 - PRIdFAST8, 137
 - PRIdLEAST16, 137
 - PRIdLEAST32, 137
 - PRIdLEAST8, 137
 - PRIdPTR, 137
 - PRi16, 137
 - PRi32, 138
 - PRi8, 138
 - PRiFAST16, 138
 - PRiFAST32, 138
 - PRiFAST8, 138
 - PRiLEAST16, 138
 - PRiLEAST32, 138
 - PRiLEAST8, 138
 - PRiPTR, 138
 - PRIo16, 138
 - PRIo32, 138
 - PRIo8, 139
 - PRIoFAST16, 139
 - PRIoFAST32, 139
 - PRIoFAST8, 139
 - PRIoLEAST16, 139
 - PRIoLEAST32, 139
 - PRIoLEAST8, 139
 - PRIoPTR, 139
 - PRiU16, 139
 - PRiU32, 139
 - PRiU8, 139
 - PRiUFAST16, 140
 - PRiUFAST32, 140

- PRIuFAST8, [140](#)
 - PRIuLEAST16, [140](#)
 - PRIuLEAST32, [140](#)
 - PRIuLEAST8, [140](#)
 - PRIuPTR, [140](#)
 - PRIX16, [140](#)
 - PRIX16, [140](#)
 - PRIX32, [140](#)
 - PRIX32, [140](#)
 - PRIX8, [141](#)
 - PRIX8, [141](#)
 - PRIXFAST16, [141](#)
 - PRIXFAST16, [141](#)
 - PRIXFAST32, [141](#)
 - PRIXFAST32, [141](#)
 - PRIXFAST8, [141](#)
 - PRIXFAST8, [141](#)
 - PRIXLEAST16, [141](#)
 - PRIXLEAST16, [141](#)
 - PRIXLEAST32, [141](#)
 - PRIXLEAST32, [142](#)
 - PRIXLEAST8, [142](#)
 - PRIXLEAST8, [142](#)
 - PRIXPTR, [142](#)
 - PRIXPTR, [142](#)
 - SCNd16, [142](#)
 - SCNd32, [142](#)
 - SCNd8, [142](#)
 - SCNdFAST16, [142](#)
 - SCNdFAST32, [142](#)
 - SCNdFAST8, [142](#)
 - SCNdLEAST16, [143](#)
 - SCNdLEAST32, [143](#)
 - SCNdLEAST8, [143](#)
 - SCNdPTR, [143](#)
 - SCNi16, [143](#)
 - SCNi32, [143](#)
 - SCNi8, [143](#)
 - SCNiFAST16, [143](#)
 - SCNiFAST32, [143](#)
 - SCNiFAST8, [143](#)
 - SCNiLEAST16, [143](#)
 - SCNiLEAST32, [144](#)
 - SCNiLEAST8, [144](#)
 - SCNiPTR, [144](#)
 - SCNo16, [144](#)
 - SCNo32, [144](#)
 - SCNo8, [144](#)
 - SCNoFAST16, [144](#)
 - SCNoFAST32, [144](#)
 - SCNoFAST8, [144](#)
 - SCNoLEAST16, [144](#)
 - SCNoLEAST32, [144](#)
 - SCNoLEAST8, [145](#)
 - SCNoPTR, [145](#)
 - SCNu16, [145](#)
 - SCNu32, [145](#)
 - SCNu8, [145](#)
 - SCNuFAST16, [145](#)
 - SCNuFAST32, [145](#)
 - SCNuFAST8, [145](#)
 - SCNuLEAST16, [145](#)
 - SCNuLEAST32, [145](#)
 - SCNuLEAST8, [145](#)
 - SCNuPTR, [146](#)
 - SCNx16, [146](#)
 - SCNx32, [146](#)
 - SCNx8, [146](#)
 - SCNxFAST16, [146](#)
 - SCNxFAST32, [146](#)
 - SCNxFAST8, [146](#)
 - SCNxLEAST16, [146](#)
 - SCNxLEAST32, [146](#)
 - SCNxLEAST8, [146](#)
 - SCNxPTR, [146](#)
 - uint_farp_ptr_t, [147](#)
- avr_io
- _PROTECTED_WRITE, [261](#)
 - _PROTECTED_WRITE_SPM, [261](#)
- avr_math
- acos, [154](#)
 - acosf, [149](#)
 - asin, [154](#)
 - asinf, [149](#)
 - atan, [154](#)
 - atan2, [154](#)
 - atan2f, [149](#)
 - atanf, [150](#)
 - cbrt, [154](#)
 - cbrtf, [150](#)
 - ceil, [155](#)
 - ceilf, [150](#)
 - copysign, [155](#)
 - copysignf, [150](#)
 - cos, [155](#)
 - cosf, [150](#)
 - cosh, [155](#)
 - coshf, [150](#)
 - exp, [155](#)
 - expf, [150](#)

- [fabs](#), 155
- [fabsf](#), 150
- [fdim](#), 155
- [fdimf](#), 150
- [floor](#), 155
- [floorf](#), 150
- [fma](#), 155
- [fmaf](#), 150
- [fmax](#), 156
- [fmaxf](#), 151
- [fmin](#), 156
- [fminf](#), 151
- [fmod](#), 156
- [fmodf](#), 151
- [frexp](#), 156
- [frexpf](#), 151
- [hypot](#), 156
- [hypotf](#), 151
- [INFINITY](#), 151
- [isfinite](#), 156
- [isfinitef](#), 151
- [isinf](#), 157
- [isinff](#), 151
- [isnan](#), 157
- [isnanf](#), 151
- [ldexp](#), 157
- [ldexpf](#), 151
- [log](#), 157
- [log10](#), 157
- [log10f](#), 151
- [logf](#), 152
- [lrint](#), 157
- [lrintf](#), 152
- [lround](#), 157
- [lroundf](#), 152
- [M_1_PI](#), 152
- [M_2_PI](#), 152
- [M_2_SQRTPI](#), 152
- [M_E](#), 152
- [M_LN10](#), 152
- [M_LN2](#), 152
- [M_LOG10E](#), 152
- [M_LOG2E](#), 152
- [M_PI](#), 153
- [M_PI_2](#), 153
- [M_PI_4](#), 153
- [M_SQRT1_2](#), 153
- [M_SQRT2](#), 153
- [modf](#), 158
- [modff](#), 158
- [NAN](#), 153
- [pow](#), 158
- [powf](#), 153
- [round](#), 158
- [roundf](#), 153
- [signbit](#), 158
- [signbitf](#), 153
- [sin](#), 159
- [sinf](#), 153
- [sinh](#), 159
- [sinhf](#), 153
- [sqrt](#), 159
- [sqrtf](#), 159
- [square](#), 159
- [squaref](#), 154
- [tan](#), 159
- [tanf](#), 154
- [tanh](#), 159
- [tanhf](#), 154
- [trunc](#), 159
- [truncf](#), 154
- [avr_pgmspace](#)
 - [memccpy_P](#), 274
 - [memchr_P](#), 274
 - [memcmp_P](#), 275
 - [memcmp_PF](#), 275
 - [memcpy_P](#), 275
 - [memcpy_PF](#), 275
 - [memmem_P](#), 276
 - [memrchr_P](#), 276
 - [pgm_get_far_address](#), 266
 - [PGM_P](#), 267
 - [pgm_read_byte](#), 267
 - [pgm_read_byte_far](#), 267
 - [pgm_read_byte_near](#), 268
 - [pgm_read_dword](#), 268
 - [pgm_read_dword_far](#), 268
 - [pgm_read_dword_near](#), 268
 - [pgm_read_float](#), 268
 - [pgm_read_float_far](#), 268
 - [pgm_read_float_near](#), 269
 - [pgm_read_ptr](#), 269
 - [pgm_read_ptr_far](#), 269
 - [pgm_read_ptr_near](#), 269
 - [pgm_read_word](#), 269
 - [pgm_read_word_far](#), 270
 - [pgm_read_word_near](#), 270
 - [PGM_VOID_P](#), 270
 - [prog_char](#), 270
 - [prog_int16_t](#), 271

- prog_int32_t, [271](#)
- prog_int64_t, [271](#)
- prog_int8_t, [272](#)
- prog_uchar, [272](#)
- prog_uint16_t, [272](#)
- prog_uint32_t, [273](#)
- prog_uint64_t, [273](#)
- prog_uint8_t, [274](#)
- prog_void, [274](#)
- PROGMEM, [270](#)
- PSTR, [270](#)
- strcasecmp_P, [276](#)
- strcasecmp_PF, [277](#)
- strcasestr_P, [277](#)
- strcat_P, [277](#)
- strcat_PF, [277](#)
- strchr_P, [278](#)
- strchrnul_P, [278](#)
- strcmp_P, [278](#)
- strcmp_PF, [279](#)
- strcpy_P, [279](#)
- strcpy_PF, [279](#)
- strcspn_P, [280](#)
- strlcat_P, [280](#)
- strlcat_PF, [280](#)
- strncpy_P, [281](#)
- strncpy_PF, [281](#)
- strlen_P, [281](#)
- strlen_PF, [282](#)
- strncasecmp_P, [282](#)
- strncasecmp_PF, [282](#)
- strncat_P, [283](#)
- strncat_PF, [283](#)
- strncmp_P, [283](#)
- strncmp_PF, [284](#)
- strncpy_P, [284](#)
- strncpy_PF, [284](#)
- strlen_P, [285](#)
- strlen_PF, [285](#)
- strpbrk_P, [286](#)
- strchr_P, [286](#)
- strsep_P, [286](#)
- strspn_P, [286](#)
- strstr_P, [287](#)
- strstr_PF, [287](#)
- strtok_P, [287](#)
- strtok_rP, [288](#)
- avr_power
 - clock_prescale_set, [291](#)
- avr_sfr
 - _BV, [294](#)
 - bit_is_clear, [294](#)
 - bit_is_set, [294](#)
 - loop_until_bit_is_clear, [295](#)
 - loop_until_bit_is_set, [295](#)
- avr_sleep
 - sleep_bod_disable, [297](#)
 - sleep_cpu, [297](#)
 - sleep_disable, [298](#)
 - sleep_enable, [298](#)
 - sleep_mode, [298](#)
- avr_stdint
 - INT16_C, [165](#)
 - INT16_MAX, [165](#)
 - INT16_MIN, [165](#)
 - int16_t, [171](#)
 - INT32_C, [165](#)
 - INT32_MAX, [165](#)
 - INT32_MIN, [166](#)
 - int32_t, [171](#)
 - INT64_C, [166](#)
 - INT64_MAX, [166](#)
 - INT64_MIN, [166](#)
 - int64_t, [171](#)
 - INT8_C, [166](#)
 - INT8_MAX, [166](#)
 - INT8_MIN, [166](#)
 - int8_t, [171](#)
 - INT_FAST16_MAX, [166](#)
 - INT_FAST16_MIN, [166](#)
 - int_fast16_t, [171](#)
 - INT_FAST32_MAX, [166](#)
 - INT_FAST32_MIN, [166](#)
 - int_fast32_t, [171](#)
 - INT_FAST64_MAX, [167](#)
 - INT_FAST64_MIN, [167](#)
 - int_fast64_t, [171](#)
 - INT_FAST8_MAX, [167](#)
 - INT_FAST8_MIN, [167](#)
 - int_fast8_t, [171](#)
 - INT_LEAST16_MAX, [167](#)
 - INT_LEAST16_MIN, [167](#)
 - int_least16_t, [171](#)
 - INT_LEAST32_MAX, [167](#)
 - INT_LEAST32_MIN, [167](#)
 - int_least32_t, [171](#)
 - INT_LEAST64_MAX, [167](#)
 - INT_LEAST64_MIN, [167](#)
 - int_least64_t, [172](#)
 - INT_LEAST8_MAX, [167](#)

- INT_LEAST8_MIN, 168
- int_least8_t, 172
- INTMAX_C, 168
- INTMAX_MAX, 168
- INTMAX_MIN, 168
- intmax_t, 172
- INTPTR_MAX, 168
- INTPTR_MIN, 168
- intptr_t, 172
- PTRDIFF_MAX, 168
- PTRDIFF_MIN, 168
- SIG_ATOMIC_MAX, 168
- SIG_ATOMIC_MIN, 168
- SIZE_MAX, 168
- UINT16_C, 169
- UINT16_MAX, 169
- uint16_t, 172
- UINT32_C, 169
- UINT32_MAX, 169
- uint32_t, 172
- UINT64_C, 169
- UINT64_MAX, 169
- uint64_t, 172
- UINT8_C, 169
- UINT8_MAX, 169
- uint8_t, 172
- UINT_FAST16_MAX, 169
- uint_fast16_t, 172
- UINT_FAST32_MAX, 169
- uint_fast32_t, 173
- UINT_FAST64_MAX, 170
- uint_fast64_t, 173
- UINT_FAST8_MAX, 170
- uint_fast8_t, 173
- UINT_LEAST16_MAX, 170
- uint_least16_t, 173
- UINT_LEAST32_MAX, 170
- uint_least32_t, 173
- UINT_LEAST64_MAX, 170
- uint_least64_t, 173
- UINT_LEAST8_MAX, 170
- uint_least8_t, 173
- UINTMAX_C, 170
- UINTMAX_MAX, 170
- uintmax_t, 173
- UINTPTR_MAX, 170
- uintptr_t, 173
- avr_stdio
 - _FDEV_EOF, 178
 - _FDEV_ERR, 178
 - _FDEV_SETUP_READ, 178
 - _FDEV_SETUP_RW, 178
 - _FDEV_SETUP_WRITE, 178
 - clearerr, 181
 - EOF, 179
 - fclose, 181
 - fdev_close, 179
 - fdev_get_udata, 179
 - fdev_set_udata, 179
 - FDEV_SETUP_STREAM, 179
 - fdev_setup_stream, 179
 - fdevopen, 181
 - feof, 182
 - ferror, 182
 - fflush, 182
 - fgetc, 182
 - fgets, 182
 - FILE, 181
 - fprintf, 183
 - fprintf_P, 183
 - fputc, 183
 - fputs, 183
 - fputs_P, 183
 - fread, 183
 - fscanf, 183
 - fscanf_P, 184
 - fwrite, 184
 - getc, 180
 - getchar, 180
 - gets, 184
 - printf, 184
 - printf_P, 184
 - putc, 180
 - putchar, 180
 - puts, 184
 - puts_P, 184
 - scanf, 184
 - scanf_P, 185
 - snprintf, 185
 - snprintf_P, 185
 - sprintf, 185
 - sprintf_P, 185
 - sscanf, 185
 - sscanf_P, 185
 - stderr, 180
 - stdin, 180
 - stdout, 181
 - ungetc, 185
 - vfprintf, 186
 - vfprintf_P, 188

- vfscanf, [188](#)
- vfscanf_P, [191](#)
- vprintf, [191](#)
- vscanf, [191](#)
- vsnprintf, [191](#)
- vsnprintf_P, [191](#)
- vsprintf, [191](#)
- vsprintf_P, [191](#)
- avr_stdlib
 - __compar_fn_t, [194](#)
 - __malloc_heap_end, [202](#)
 - __malloc_heap_start, [202](#)
 - __malloc_margin, [203](#)
 - abort, [194](#)
 - abs, [194](#)
 - atof, [194](#)
 - atoi, [195](#)
 - atol, [195](#)
 - bsearch, [195](#)
 - calloc, [196](#)
 - div, [196](#)
 - DTOSTR_ALWAYS_SIGN, [193](#)
 - DTOSTR_PLUS_SIGN, [193](#)
 - DTOSTR_UPPERCASE, [193](#)
 - dtostre, [196](#)
 - dtostrf, [196](#)
 - exit, [196](#)
 - EXIT_FAILURE, [194](#)
 - EXIT_SUCCESS, [194](#)
 - free, [197](#)
 - itoa, [197](#)
 - labs, [197](#)
 - ldiv, [197](#)
 - ltoa, [198](#)
 - malloc, [198](#)
 - qsort, [198](#)
 - rand, [199](#)
 - RAND_MAX, [194](#)
 - rand_r, [199](#)
 - random, [199](#)
 - RANDOM_MAX, [194](#)
 - random_r, [199](#)
 - realloc, [199](#)
 - srand, [200](#)
 - srandom, [200](#)
 - strtod, [200](#)
 - strtol, [200](#)
 - strtoul, [201](#)
 - ultoa, [201](#)
 - utoa, [202](#)
- avr_string
 - _FFS, [204](#)
 - ffs, [204](#)
 - ffsl, [205](#)
 - ffsll, [205](#)
 - memccpy, [205](#)
 - memchr, [205](#)
 - memcmp, [205](#)
 - memcpy, [206](#)
 - memmem, [206](#)
 - memmove, [206](#)
 - memrchr, [207](#)
 - memset, [207](#)
 - strcasecmp, [207](#)
 - strcasestr, [207](#)
 - strcat, [208](#)
 - strchr, [208](#)
 - strchrnul, [208](#)
 - strcmp, [209](#)
 - strcpy, [209](#)
 - strcspn, [209](#)
 - strdup, [209](#)
 - strlcat, [210](#)
 - strncpy, [210](#)
 - strlen, [211](#)
 - strlwr, [211](#)
 - strncasecmp, [211](#)
 - strncat, [212](#)
 - strncmp, [212](#)
 - strncpy, [212](#)
 - strnlen, [212](#)
 - strpbrk, [213](#)
 - strrchr, [213](#)
 - strrev, [213](#)
 - strsep, [214](#)
 - strspn, [214](#)
 - strstr, [214](#)
 - strtok, [214](#)
 - strtok_r, [215](#)
 - strupr, [215](#)
- avr_time
 - _MONTHS_, [219](#)
 - _WEEK_DAYS_, [219](#)
 - asctime, [220](#)
 - asctime_r, [220](#)
 - ctime, [220](#)
 - ctime_r, [220](#)
 - daylight_seconds, [220](#)
 - difftime, [220](#)
 - equation_of_time, [220](#)

- fatfs_time, [220](#)
- gm_sidereal, [220](#)
- gmtime, [221](#)
- gmtime_r, [221](#)
- is_leap_year, [221](#)
- iso_week_date, [221](#)
- iso_week_date_r, [221](#)
- isotime, [221](#)
- isotime_r, [221](#)
- lm_sidereal, [221](#)
- localtime, [221](#)
- localtime_r, [222](#)
- mk_gmtime, [222](#)
- mktime, [222](#)
- month_length, [222](#)
- moon_phase, [222](#)
- NTP_OFFSET, [218](#)
- ONE_DAY, [219](#)
- ONE_DEGREE, [219](#)
- ONE_HOUR, [219](#)
- set_dst, [222](#)
- set_position, [223](#)
- set_system_time, [223](#)
- set_zone, [223](#)
- solar_declination, [223](#)
- solar_noon, [224](#)
- strftime, [224](#)
- sun_rise, [224](#)
- sun_set, [224](#)
- system_tick, [224](#)
- time, [224](#)
- time_t, [219](#)
- UNIX_OFFSET, [219](#)
- week_of_month, [224](#)
- week_of_year, [225](#)
- avr_version
 - __AVR_LIBC_DATE__, [299](#)
 - __AVR_LIBC_DATE_STRING__, [299](#)
 - __AVR_LIBC_MAJOR__, [299](#)
 - __AVR_LIBC_MINOR__, [299](#)
 - __AVR_LIBC_REVISION__, [299](#)
 - __AVR_LIBC_VERSION_-
STRING__, [299](#)
 - __AVR_LIBC_VERSION__, [299](#)
- avr_watchdog
 - __attribute__, [303](#)
 - wdt_reset, [301](#)
 - WDTO_120MS, [301](#)
 - WDTO_15MS, [301](#)
 - WDTO_1S, [301](#)
 - WDTO_250MS, [301](#)
 - WDTO_2S, [301](#)
 - WDTO_30MS, [301](#)
 - WDTO_4S, [302](#)
 - WDTO_500MS, [302](#)
 - WDTO_60MS, [302](#)
 - WDTO_8S, [302](#)
- avrdude, usage, [117](#)
- avrprog, usage, [117](#)
- BADISR_vect
 - avr_interrupts, [257](#)
- BAUD_TOL
 - util_setbaud, [316](#)
- bit_is_clear
 - avr_sfr, [294](#)
- bit_is_set
 - avr_sfr, [294](#)
- boot.h, [366](#)
- boot_is_spm_interrupt
 - avr_boot, [227](#)
- boot_lock_bits_set
 - avr_boot, [227](#)
- boot_lock_bits_set_safe
 - avr_boot, [227](#)
- boot_lock_fuse_bits_get
 - avr_boot, [227](#)
- boot_page_erase
 - avr_boot, [228](#)
- boot_page_erase_safe
 - avr_boot, [228](#)
- boot_page_fill
 - avr_boot, [228](#)
- boot_page_fill_safe
 - avr_boot, [229](#)
- boot_page_write
 - avr_boot, [229](#)
- boot_page_write_safe
 - avr_boot, [229](#)
- boot_rww_busy
 - avr_boot, [229](#)
- boot_rww_enable
 - avr_boot, [230](#)
- boot_rww_enable_safe
 - avr_boot, [230](#)
- boot_signature_byte_get
 - avr_boot, [230](#)
- boot_spm_busy
 - avr_boot, [230](#)

- boot_spm_busy_wait
 - avr_boot, [230](#)
- boot_spm_interrupt_disable
 - avr_boot, [231](#)
- boot_spm_interrupt_enable
 - avr_boot, [231](#)
- BOOTLOADER_SECTION
 - avr_boot, [231](#)
- bsearch
 - avr_stdlib, [195](#)
- calloc
 - avr_stdlib, [196](#)
- cbi
 - deprecated_items, [322](#)
- cbrt
 - avr_math, [154](#)
- cbrtf
 - avr_math, [150](#)
- ccp_write_io
 - avr_cpufunc, [232](#)
- ceil
 - avr_math, [155](#)
- ceilf
 - avr_math, [150](#)
- clearerr
 - avr_stdio, [181](#)
- cli
 - avr_interrupts, [257](#)
- clock_prescale_get
 - power.h, [377](#)
- clock_prescale_set
 - avr_power, [291](#)
- Combining C and assembly source files, [326](#)
- copysign
 - avr_math, [155](#)
- copysignf
 - avr_math, [150](#)
- cos
 - avr_math, [155](#)
- cosf
 - avr_math, [150](#)
- cosh
 - avr_math, [155](#)
- coshf
 - avr_math, [150](#)
- cpufunc.h, [367](#)
- crc16.h, [367](#)
- ctime
 - avr_time, [220](#)
- ctime_r
 - avr_time, [220](#)
- ctype
 - isalnum, [131](#)
 - isalpha, [131](#)
 - isascii, [131](#)
 - isblank, [131](#)
 - isctrl, [131](#)
 - isdigit, [131](#)
 - isgraph, [131](#)
 - islower, [131](#)
 - isprint, [131](#)
 - ispunct, [131](#)
 - isspace, [132](#)
 - isupper, [132](#)
 - isxdigit, [132](#)
 - toascii, [132](#)
 - tolower, [132](#)
 - toupper, [132](#)
- ctype.h, [367](#)
- day
 - week_date, [365](#)
- daylight_seconds
 - avr_time, [220](#)
- delay.h, [368](#)
- delay_basic.h, [368](#)
- Demo projects, [325](#)
- deprecated_items
 - cbi, [322](#)
 - enable_external_int, [322](#)
 - inb, [322](#)
 - inp, [323](#)
 - INTERRUPT, [323](#)
 - outb, [323](#)
 - outp, [323](#)
 - sbi, [324](#)
 - timer_enable_int, [324](#)
- difftime
 - avr_time, [220](#)
- disassembling, [333](#)
- div
 - avr_stdlib, [196](#)
- div_t, [362](#)
- quot, [362](#)
- rem, [362](#)
- DTOSTR_ALWAYS_SIGN
 - avr_stdlib, [193](#)
- DTOSTR_PLUS_SIGN

- avr_stdlib, [193](#)
- DTOSTR_UPPERCASE
 - avr_stdlib, [193](#)
- dtostre
 - avr_stdlib, [196](#)
- dtostrf
 - avr_stdlib, [196](#)
- EDOM
 - avr_errno, [133](#)
- EEMEM
 - avr_eeprom, [234](#)
- eeprom_busy_wait
 - avr_eeprom, [234](#)
- eeprom_is_ready
 - avr_eeprom, [235](#)
- eeprom_read_block
 - avr_eeprom, [235](#)
- eeprom_read_byte
 - avr_eeprom, [235](#)
- eeprom_read_dword
 - avr_eeprom, [235](#)
- eeprom_read_float
 - avr_eeprom, [235](#)
- eeprom_read_word
 - avr_eeprom, [235](#)
- eeprom_update_block
 - avr_eeprom, [235](#)
- eeprom_update_byte
 - avr_eeprom, [235](#)
- eeprom_update_dword
 - avr_eeprom, [236](#)
- eeprom_update_float
 - avr_eeprom, [236](#)
- eeprom_update_word
 - avr_eeprom, [236](#)
- eeprom_write_block
 - avr_eeprom, [236](#)
- eeprom_write_byte
 - avr_eeprom, [236](#)
- eeprom_write_dword
 - avr_eeprom, [236](#)
- eeprom_write_float
 - avr_eeprom, [236](#)
- eeprom_write_word
 - avr_eeprom, [236](#)
- EMPTY_INTERRUPT
 - avr_interrups, [257](#)
- enable_external_int
 - deprecated_items, [322](#)
- EOF
 - avr_stdio, [179](#)
- equation_of_time
 - avr_time, [220](#)
- ERANGE
 - avr_errno, [133](#)
- errno
 - avr_errno, [133](#)
- errno.h, [368](#)
- Example using the two-wire interface (TWI), [357](#)
- exit
 - avr_stdlib, [196](#)
- EXIT_FAILURE
 - avr_stdlib, [194](#)
- EXIT_SUCCESS
 - avr_stdlib, [194](#)
- exp
 - avr_math, [155](#)
- expf
 - avr_math, [150](#)
- F_CPU
 - util_delay, [311](#)
- fabs
 - avr_math, [155](#)
- fabsf
 - avr_math, [150](#)
- FAQ, [59](#)
- fatfs_time
 - avr_time, [220](#)
- fclose
 - avr_stdio, [181](#)
- fdev_close
 - avr_stdio, [179](#)
- fdev_get_udata
 - avr_stdio, [179](#)
- fdev_set_udata
 - avr_stdio, [179](#)
- FDEV_SETUP_STREAM
 - avr_stdio, [179](#)
- fdev_setup_stream
 - avr_stdio, [179](#)
- fdevopen
 - avr_stdio, [181](#)
- fdevopen.c, [369](#)
- fdim
 - avr_math, [155](#)
- fdimf
 - avr_math, [150](#)

- feof
 - avr_stdio, [182](#)
- ferror
 - avr_stdio, [182](#)
- fflush
 - avr_stdio, [182](#)
- ffs
 - avr_string, [204](#)
- ffsl
 - avr_string, [205](#)
- ffsll
 - avr_string, [205](#)
- fgetc
 - avr_stdio, [182](#)
- fgets
 - avr_stdio, [182](#)
- FILE
 - avr_stdio, [181](#)
- floor
 - avr_math, [155](#)
- floorf
 - avr_math, [150](#)
- fma
 - avr_math, [155](#)
- fmaf
 - avr_math, [150](#)
- fmax
 - avr_math, [156](#)
- fmaxf
 - avr_math, [151](#)
- fmin
 - avr_math, [156](#)
- fminf
 - avr_math, [151](#)
- fmod
 - avr_math, [156](#)
- fmodf
 - avr_math, [151](#)
- fprintf
 - avr_stdio, [183](#)
- fprintf_P
 - avr_stdio, [183](#)
- fputc
 - avr_stdio, [183](#)
- fputs
 - avr_stdio, [183](#)
- fputs_P
 - avr_stdio, [183](#)
- fread
 - avr_stdio, [183](#)
- free
 - avr_stdlib, [197](#)
- frexp
 - avr_math, [156](#)
- frexpf
 - avr_math, [151](#)
- fscanf
 - avr_stdio, [183](#)
- fscanf_P
 - avr_stdio, [184](#)
- fuse.h, [369](#)
- fwrite
 - avr_stdio, [184](#)
- GET_EXTENDED_FUSE_BITS
 - avr_boot, [231](#)
- GET_HIGH_FUSE_BITS
 - avr_boot, [231](#)
- GET_LOCK_BITS
 - avr_boot, [231](#)
- GET_LOW_FUSE_BITS
 - avr_boot, [231](#)
- getc
 - avr_stdio, [180](#)
- getchar
 - avr_stdio, [180](#)
- gets
 - avr_stdio, [184](#)
- gm_siderreal
 - avr_time, [220](#)
- gmtime
 - avr_time, [221](#)
- gmtime_r
 - avr_time, [221](#)
- hypot
 - avr_math, [156](#)
- hypotf
 - avr_math, [151](#)
- inb
 - deprecated_items, [322](#)
- INFINITY
 - avr_math, [151](#)
- inp
 - deprecated_items, [323](#)
- installation, [87](#)
- installation, avarice, [92](#)
- installation, avr-libc, [90](#)
- installation, avrdude, [91](#)

- installation, avrprog, [91](#)
- installation, binutils, [89](#)
- installation, gcc, [90](#)
- Installation, gdb, [91](#)
- installation, simulavr, [92](#)
- INT16_C
 - avr_stdint, [165](#)
- INT16_MAX
 - avr_stdint, [165](#)
- INT16_MIN
 - avr_stdint, [165](#)
- int16_t
 - avr_stdint, [171](#)
- INT32_C
 - avr_stdint, [165](#)
- INT32_MAX
 - avr_stdint, [165](#)
- INT32_MIN
 - avr_stdint, [166](#)
- int32_t
 - avr_stdint, [171](#)
- INT64_C
 - avr_stdint, [166](#)
- INT64_MAX
 - avr_stdint, [166](#)
- INT64_MIN
 - avr_stdint, [166](#)
- int64_t
 - avr_stdint, [171](#)
- INT8_C
 - avr_stdint, [166](#)
- INT8_MAX
 - avr_stdint, [166](#)
- INT8_MIN
 - avr_stdint, [166](#)
- int8_t
 - avr_stdint, [171](#)
- int_farptr_t
 - avr_inttypes, [147](#)
- INT_FAST16_MAX
 - avr_stdint, [166](#)
- INT_FAST16_MIN
 - avr_stdint, [166](#)
- int_fast16_t
 - avr_stdint, [171](#)
- INT_FAST32_MAX
 - avr_stdint, [166](#)
- INT_FAST32_MIN
 - avr_stdint, [166](#)
- int_fast32_t
 - avr_stdint, [171](#)
- INT_FAST64_MAX
 - avr_stdint, [167](#)
- INT_FAST64_MIN
 - avr_stdint, [167](#)
- int_fast64_t
 - avr_stdint, [171](#)
- INT_FAST8_MAX
 - avr_stdint, [167](#)
- INT_FAST8_MIN
 - avr_stdint, [167](#)
- int_fast8_t
 - avr_stdint, [171](#)
- INT_LEAST16_MAX
 - avr_stdint, [167](#)
- INT_LEAST16_MIN
 - avr_stdint, [167](#)
- int_least16_t
 - avr_stdint, [171](#)
- INT_LEAST32_MAX
 - avr_stdint, [167](#)
- INT_LEAST32_MIN
 - avr_stdint, [167](#)
- int_least32_t
 - avr_stdint, [171](#)
- INT_LEAST64_MAX
 - avr_stdint, [167](#)
- INT_LEAST64_MIN
 - avr_stdint, [167](#)
- int_least64_t
 - avr_stdint, [172](#)
- INT_LEAST8_MAX
 - avr_stdint, [167](#)
- INT_LEAST8_MIN
 - avr_stdint, [168](#)
- int_least8_t
 - avr_stdint, [172](#)
- INTERRUPT
 - deprecated_items, [323](#)
- interrupt.h, [369](#)
- INTMAX_C
 - avr_stdint, [168](#)
- INTMAX_MAX
 - avr_stdint, [168](#)
- INTMAX_MIN
 - avr_stdint, [168](#)
- intmax_t
 - avr_stdint, [172](#)
- INTPTR_MAX
 - avr_stdint, [168](#)

- INTPTR_MIN
 - avr_stdint, 168
- intptr_t
 - avr_stdint, 172
- inttypes.h, 370
- io.h, 372
- is_leap_year
 - avr_time, 221
- isalnum
 - ctype, 131
- isalpha
 - ctype, 131
- isascii
 - ctype, 131
- isblank
 - ctype, 131
- isctrl
 - ctype, 131
- isdigit
 - ctype, 131
- isfinite
 - avr_math, 156
- isfinitef
 - avr_math, 151
- isgraph
 - ctype, 131
- isinf
 - avr_math, 157
- isinff
 - avr_math, 151
- islower
 - ctype, 131
- isnan
 - avr_math, 157
- isnanf
 - avr_math, 151
- iso_week_date
 - avr_time, 221
- iso_week_date_r
 - avr_time, 221
- isotime
 - avr_time, 221
- isotime_r
 - avr_time, 221
- isprint
 - ctype, 131
- ispunct
 - ctype, 131
- ISR
 - avr_interrupts, 257
- ISR_ALIAS
 - avr_interrupts, 258
- ISR_ALIASOF
 - avr_interrupts, 258
- ISR_BLOCK
 - avr_interrupts, 258
- ISR_NAKED
 - avr_interrupts, 259
- ISR_NOBLOCK
 - avr_interrupts, 259
- isspace
 - ctype, 132
- isupper
 - ctype, 132
- isxdigit
 - ctype, 132
- itoa
 - avr_stdlib, 197
- labs
 - avr_stdlib, 197
- ldexp
 - avr_math, 157
- ldexpf
 - avr_math, 151
- ldiv
 - avr_stdlib, 197
- ldiv_t, 363
 - quot, 363
 - rem, 363
- lm_siderear
 - avr_time, 221
- localtime
 - avr_time, 221
- localtime_r
 - avr_time, 222
- lock.h, 372
- log
 - avr_math, 157
- log10
 - avr_math, 157
- log10f
 - avr_math, 151
- logf
 - avr_math, 152
- longjmp
 - setjmp, 161
- loop_until_bit_is_clear
 - avr_sfr, 295
- loop_until_bit_is_set

- avr_sfr, 295
- lrint
 - avr_math, 157
- lrintf
 - avr_math, 152
- lround
 - avr_math, 157
- lroundf
 - avr_math, 152
- ltoa
 - avr_stdlib, 198
- M_1_PI
 - avr_math, 152
- M_2_PI
 - avr_math, 152
- M_2_SQRTPI
 - avr_math, 152
- M_E
 - avr_math, 152
- M_LN10
 - avr_math, 152
- M_LN2
 - avr_math, 152
- M_LOG10E
 - avr_math, 152
- M_LOG2E
 - avr_math, 152
- M_PI
 - avr_math, 153
- M_PI_2
 - avr_math, 153
- M_PI_4
 - avr_math, 153
- M_SQRT1_2
 - avr_math, 153
- M_SQRT2
 - avr_math, 153
- malloc
 - avr_stdlib, 198
- math.h, 372
- memcpy
 - avr_string, 205
- memcpy_P
 - avr_pgmspace, 274
- memchr
 - avr_string, 205
- memchr_P
 - avr_pgmspace, 274
- memcmp
 - avr_string, 205
- memcmp_P
 - avr_pgmspace, 275
- memcmp_PF
 - avr_pgmspace, 275
- memcpyy
 - avr_string, 206
- memcpyy_P
 - avr_pgmspace, 275
- memcpyy_PF
 - avr_pgmspace, 275
- memmem
 - avr_string, 206
- memmem_P
 - avr_pgmspace, 276
- memmove
 - avr_string, 206
- memrchr
 - avr_string, 207
- memrchr_P
 - avr_pgmspace, 276
- memset
 - avr_string, 207
- mk_gmtime
 - avr_time, 222
- mktime
 - avr_time, 222
- modf
 - avr_math, 158
- modff
 - avr_math, 158
- month_length
 - avr_time, 222
- moon_phase
 - avr_time, 222
- NAN
 - avr_math, 153
- NONATOMIC_BLOCK
 - util_atomic, 305
- NONATOMIC_FORCEOFF
 - util_atomic, 306
- NONATOMIC_RESTORESTATE
 - util_atomic, 306
- NTP_OFFSET
 - avr_time, 218
- ONE_DAY
 - avr_time, 219
- ONE_DEGREE

- avr_time, 219
- ONE_HOUR
 - avr_time, 219
- outb
 - deprecated_items, 323
- outp
 - deprecated_items, 323
- parity.h, 375
- parity_even_bit
 - util_parity, 314
- pgm_get_far_address
 - avr_pgmspace, 266
- PGM_P
 - avr_pgmspace, 267
- pgm_read_byte
 - avr_pgmspace, 267
- pgm_read_byte_far
 - avr_pgmspace, 267
- pgm_read_byte_near
 - avr_pgmspace, 268
- pgm_read_dword
 - avr_pgmspace, 268
- pgm_read_dword_far
 - avr_pgmspace, 268
- pgm_read_dword_near
 - avr_pgmspace, 268
- pgm_read_float
 - avr_pgmspace, 268
- pgm_read_float_far
 - avr_pgmspace, 268
- pgm_read_float_near
 - avr_pgmspace, 269
- pgm_read_ptr
 - avr_pgmspace, 269
- pgm_read_ptr_far
 - avr_pgmspace, 269
- pgm_read_ptr_near
 - avr_pgmspace, 269
- pgm_read_word
 - avr_pgmspace, 269
- pgm_read_word_far
 - avr_pgmspace, 270
- pgm_read_word_near
 - avr_pgmspace, 270
- PGM_VOID_P
 - avr_pgmspace, 270
- pgmspace.h, 375
- pow
 - avr_math, 158
- power.h, 377
 - clock_prescale_get, 377
- powf
 - avr_math, 153
- PRId16
 - avr_inttypes, 137
- PRId32
 - avr_inttypes, 137
- PRId8
 - avr_inttypes, 137
- PRIdFAST16
 - avr_inttypes, 137
- PRIdFAST32
 - avr_inttypes, 137
- PRIdFAST8
 - avr_inttypes, 137
- PRIdLEAST16
 - avr_inttypes, 137
- PRIdLEAST32
 - avr_inttypes, 137
- PRIdLEAST8
 - avr_inttypes, 137
- PRIdPTR
 - avr_inttypes, 137
- PRi16
 - avr_inttypes, 137
- PRi32
 - avr_inttypes, 138
- PRi8
 - avr_inttypes, 138
- PRiFAST16
 - avr_inttypes, 138
- PRiFAST32
 - avr_inttypes, 138
- PRiFAST8
 - avr_inttypes, 138
- PRiLEAST16
 - avr_inttypes, 138
- PRiLEAST32
 - avr_inttypes, 138
- PRiLEAST8
 - avr_inttypes, 138
- PRiPTR
 - avr_inttypes, 138
- printf
 - avr_stdio, 184
- printf_P
 - avr_stdio, 184
- PRIo16
 - avr_inttypes, 138

- PRIo32
 - avr_inttypes, [138](#)
- PRIo8
 - avr_inttypes, [139](#)
- PRIoFAST16
 - avr_inttypes, [139](#)
- PRIoFAST32
 - avr_inttypes, [139](#)
- PRIoFAST8
 - avr_inttypes, [139](#)
- PRIoLEAST16
 - avr_inttypes, [139](#)
- PRIoLEAST32
 - avr_inttypes, [139](#)
- PRIoLEAST8
 - avr_inttypes, [139](#)
- PRIoPTR
 - avr_inttypes, [139](#)
- PRIo16
 - avr_inttypes, [139](#)
- PRIo32
 - avr_inttypes, [139](#)
- PRIo8
 - avr_inttypes, [139](#)
- PRIoFAST16
 - avr_inttypes, [140](#)
- PRIoFAST32
 - avr_inttypes, [140](#)
- PRIoFAST8
 - avr_inttypes, [140](#)
- PRIoLEAST16
 - avr_inttypes, [140](#)
- PRIoLEAST32
 - avr_inttypes, [140](#)
- PRIoLEAST8
 - avr_inttypes, [140](#)
- PRIoPTR
 - avr_inttypes, [140](#)
- PRIX16
 - avr_inttypes, [140](#)
- PRIX16
 - avr_inttypes, [140](#)
- PRIX32
 - avr_inttypes, [140](#)
- PRIX32
 - avr_inttypes, [140](#)
- PRIX8
 - avr_inttypes, [141](#)
- PRIX8
 - avr_inttypes, [141](#)
- PRIXFAST16
 - avr_inttypes, [141](#)
- PRIXFAST16
 - avr_inttypes, [141](#)
- PRIXFAST32
 - avr_inttypes, [141](#)
- PRIXFAST32
 - avr_inttypes, [141](#)
- PRIXFAST8
 - avr_inttypes, [141](#)
- PRIXFAST8
 - avr_inttypes, [141](#)
- PRIXLEAST16
 - avr_inttypes, [141](#)
- PRIXLEAST16
 - avr_inttypes, [141](#)
- PRIXLEAST32
 - avr_inttypes, [141](#)
- PRIXLEAST32
 - avr_inttypes, [142](#)
- PRIXLEAST8
 - avr_inttypes, [142](#)
- PRIXLEAST8
 - avr_inttypes, [142](#)
- PRIXPTR
 - avr_inttypes, [142](#)
- PRIXPTR
 - avr_inttypes, [142](#)
- prog_char
 - avr_pgmspace, [270](#)
- prog_int16_t
 - avr_pgmspace, [271](#)
- prog_int32_t
 - avr_pgmspace, [271](#)
- prog_int64_t
 - avr_pgmspace, [271](#)
- prog_int8_t
 - avr_pgmspace, [272](#)
- prog_uchar
 - avr_pgmspace, [272](#)
- prog_uint16_t
 - avr_pgmspace, [272](#)
- prog_uint32_t
 - avr_pgmspace, [273](#)
- prog_uint64_t
 - avr_pgmspace, [273](#)
- prog_uint8_t
 - avr_pgmspace, [274](#)
- prog_void
 - avr_pgmspace, [274](#)

PROGMEM
 avr_pgmspace, 270
PSTR
 avr_pgmspace, 270
PTRDIFF_MAX
 avr_stdint, 168
PTRDIFF_MIN
 avr_stdint, 168
putc
 avr_stdio, 180
putchar
 avr_stdio, 180
puts
 avr_stdio, 184
puts_P
 avr_stdio, 184

qsort
 avr_stdlib, 198
quot
 div_t, 362
 ldiv_t, 363

rand
 avr_stdlib, 199
RAND_MAX
 avr_stdlib, 194
rand_r
 avr_stdlib, 199
random
 avr_stdlib, 199
RANDOM_MAX
 avr_stdlib, 194
random_r
 avr_stdlib, 199
realloc
 avr_stdlib, 199
rem
 div_t, 362
 ldiv_t, 363
reti
 avr_interrupts, 259
round
 avr_math, 158
roundf
 avr_math, 153

sbi
 deprecated_items, 324
scanf
 avr_stdio, 184
scanf_P
 avr_stdio, 185
SCNd16
 avr_inttypes, 142
SCNd32
 avr_inttypes, 142
SCNd8
 avr_inttypes, 142
SCNdFAST16
 avr_inttypes, 142
SCNdFAST32
 avr_inttypes, 142
SCNdFAST8
 avr_inttypes, 142
SCNdLEAST16
 avr_inttypes, 143
SCNdLEAST32
 avr_inttypes, 143
SCNdLEAST8
 avr_inttypes, 143
SCNdPTR
 avr_inttypes, 143
SCNi16
 avr_inttypes, 143
SCNi32
 avr_inttypes, 143
SCNi8
 avr_inttypes, 143
SCNiFAST16
 avr_inttypes, 143
SCNiFAST32
 avr_inttypes, 143
SCNiFAST8
 avr_inttypes, 143
SCNiLEAST16
 avr_inttypes, 143
SCNiLEAST32
 avr_inttypes, 144
SCNiLEAST8
 avr_inttypes, 144
SCNiPTR
 avr_inttypes, 144
SCNo16
 avr_inttypes, 144
SCNo32
 avr_inttypes, 144
SCNo8
 avr_inttypes, 144
SCNoFAST16

- [avr_inttypes, 144](#)
- [SCNoFAST32](#)
 - [avr_inttypes, 144](#)
- [SCNoFAST8](#)
 - [avr_inttypes, 144](#)
- [SCNoLEAST16](#)
 - [avr_inttypes, 144](#)
- [SCNoLEAST32](#)
 - [avr_inttypes, 144](#)
- [SCNoLEAST8](#)
 - [avr_inttypes, 145](#)
- [SCNoPTR](#)
 - [avr_inttypes, 145](#)
- [SCNu16](#)
 - [avr_inttypes, 145](#)
- [SCNu32](#)
 - [avr_inttypes, 145](#)
- [SCNu8](#)
 - [avr_inttypes, 145](#)
- [SCNuFAST16](#)
 - [avr_inttypes, 145](#)
- [SCNuFAST32](#)
 - [avr_inttypes, 145](#)
- [SCNuFAST8](#)
 - [avr_inttypes, 145](#)
- [SCNuLEAST16](#)
 - [avr_inttypes, 145](#)
- [SCNuLEAST32](#)
 - [avr_inttypes, 145](#)
- [SCNuLEAST8](#)
 - [avr_inttypes, 145](#)
- [SCNuPTR](#)
 - [avr_inttypes, 146](#)
- [SCNx16](#)
 - [avr_inttypes, 146](#)
- [SCNx32](#)
 - [avr_inttypes, 146](#)
- [SCNx8](#)
 - [avr_inttypes, 146](#)
- [SCNxFAST16](#)
 - [avr_inttypes, 146](#)
- [SCNxFAST32](#)
 - [avr_inttypes, 146](#)
- [SCNxFAST8](#)
 - [avr_inttypes, 146](#)
- [SCNxLEAST16](#)
 - [avr_inttypes, 146](#)
- [SCNxLEAST32](#)
 - [avr_inttypes, 146](#)
- [SCNxLEAST8](#)
 - [avr_inttypes, 146](#)
- [SCNxPTR](#)
 - [avr_inttypes, 146](#)
- [sei](#)
 - [avr_interrupts, 259](#)
- [set_dst](#)
 - [avr_time, 222](#)
- [set_position](#)
 - [avr_time, 223](#)
- [set_system_time](#)
 - [avr_time, 223](#)
- [set_zone](#)
 - [avr_time, 223](#)
- [setbaud.h, 378](#)
- [setjmp](#)
 - [longjmp, 161](#)
 - [setjmp, 161](#)
- [setjmp.h, 378](#)
- [SIG_ATOMIC_MAX](#)
 - [avr_stdint, 168](#)
- [SIG_ATOMIC_MIN](#)
 - [avr_stdint, 168](#)
- [SIGNAL](#)
 - [avr_interrupts, 259](#)
- [signature.h, 378](#)
- [signbit](#)
 - [avr_math, 158](#)
- [signbitf](#)
 - [avr_math, 153](#)
- [sin](#)
 - [avr_math, 159](#)
- [sinf](#)
 - [avr_math, 153](#)
- [sinh](#)
 - [avr_math, 159](#)
- [sinhf](#)
 - [avr_math, 153](#)
- [SIZE_MAX](#)
 - [avr_stdint, 168](#)
- [sleep.h, 378](#)
- [sleep_bod_disable](#)
 - [avr_sleep, 297](#)
- [sleep_cpu](#)
 - [avr_sleep, 297](#)
- [sleep_disable](#)
 - [avr_sleep, 298](#)
- [sleep_enable](#)
 - [avr_sleep, 298](#)
- [sleep_mode](#)
 - [avr_sleep, 298](#)

snprintf
 avr_stdio, 185
snprintf_P
 avr_stdio, 185
solar_declination
 avr_time, 223
solar_noon
 avr_time, 224
sprintf
 avr_stdio, 185
sprintf_P
 avr_stdio, 185
sqrt
 avr_math, 159
sqrtf
 avr_math, 159
square
 avr_math, 159
squaref
 avr_math, 154
srand
 avr_stdlib, 200
srandom
 avr_stdlib, 200
sscanf
 avr_stdio, 185
sscanf_P
 avr_stdio, 185
stderr
 avr_stdio, 180
stdin
 avr_stdio, 180
stdint.h, 378
stdio.h, 381
stdlib.h, 383
stdout
 avr_stdio, 181
strcasecmp
 avr_string, 207
strcasecmp_P
 avr_pgmspace, 276
strcasecmp_PF
 avr_pgmspace, 277
strcasestr
 avr_string, 207
strcasestr_P
 avr_pgmspace, 277
strcat
 avr_string, 208
strcat_P
 avr_pgmspace, 277
strcat_PF
 avr_pgmspace, 277
strchr
 avr_string, 208
strchr_P
 avr_pgmspace, 278
strchrnul
 avr_string, 208
strchrnul_P
 avr_pgmspace, 278
strcmp
 avr_string, 209
strcmp_P
 avr_pgmspace, 278
strcmp_PF
 avr_pgmspace, 279
strcpy
 avr_string, 209
strcpy_P
 avr_pgmspace, 279
strcpy_PF
 avr_pgmspace, 279
strcspn
 avr_string, 209
strcspn_P
 avr_pgmspace, 280
strdup
 avr_string, 209
strftime
 avr_time, 224
string.h, 384
strlcat
 avr_string, 210
strlcat_P
 avr_pgmspace, 280
strlcat_PF
 avr_pgmspace, 280
strlcpy
 avr_string, 210
strlcpy_P
 avr_pgmspace, 281
strlcpy_PF
 avr_pgmspace, 281
strlen
 avr_string, 211
strlen_P
 avr_pgmspace, 281
strlen_PF
 avr_pgmspace, 282

- strlwr
 - avr_string, [211](#)
- strncasecmp
 - avr_string, [211](#)
- strncasecmp_P
 - avr_pgmspace, [282](#)
- strncasecmp_PF
 - avr_pgmspace, [282](#)
- strncat
 - avr_string, [212](#)
- strncat_P
 - avr_pgmspace, [283](#)
- strncat_PF
 - avr_pgmspace, [283](#)
- strncmp
 - avr_string, [212](#)
- strncmp_P
 - avr_pgmspace, [283](#)
- strncmp_PF
 - avr_pgmspace, [284](#)
- strncpy
 - avr_string, [212](#)
- strncpy_P
 - avr_pgmspace, [284](#)
- strncpy_PF
 - avr_pgmspace, [284](#)
- strlen
 - avr_string, [212](#)
- strlen_P
 - avr_pgmspace, [285](#)
- strlen_PF
 - avr_pgmspace, [285](#)
- strpbrk
 - avr_string, [213](#)
- strpbrk_P
 - avr_pgmspace, [286](#)
- strchr
 - avr_string, [213](#)
- strchr_P
 - avr_pgmspace, [286](#)
- strrev
 - avr_string, [213](#)
- strsep
 - avr_string, [214](#)
- strsep_P
 - avr_pgmspace, [286](#)
- strspn
 - avr_string, [214](#)
- strspn_P
 - avr_pgmspace, [286](#)
- strstr
 - avr_string, [214](#)
- strstr_P
 - avr_pgmspace, [287](#)
- strstr_PF
 - avr_pgmspace, [287](#)
- strtod
 - avr_stdlib, [200](#)
- strtok
 - avr_string, [214](#)
- strtok_P
 - avr_pgmspace, [287](#)
- strtok_r
 - avr_string, [215](#)
- strtok_rP
 - avr_pgmspace, [288](#)
- strtol
 - avr_stdlib, [200](#)
- strtoul
 - avr_stdlib, [201](#)
- strupr
 - avr_string, [215](#)
- sun_rise
 - avr_time, [224](#)
- sun_set
 - avr_time, [224](#)
- supported devices, [2](#)
- system_tick
 - avr_time, [224](#)
- tan
 - avr_math, [159](#)
- tanf
 - avr_math, [154](#)
- tanh
 - avr_math, [159](#)
- tanhf
 - avr_math, [154](#)
- time
 - avr_time, [224](#)
- time.h, [385](#)
- time_t
 - avr_time, [219](#)
- timer_enable_int
 - deprecated_items, [324](#)
- tm, [363](#)
 - tm_hour, [364](#)
 - tm_isdst, [364](#)
 - tm_mday, [364](#)
 - tm_min, [364](#)

- tm_mon, [364](#)
- tm_sec, [364](#)
- tm_wday, [364](#)
- tm_yday, [364](#)
- tm_year, [364](#)
- tm_hour
 - tm, [364](#)
- tm_isdst
 - tm, [364](#)
- tm_mday
 - tm, [364](#)
- tm_min
 - tm, [364](#)
- tm_mon
 - tm, [364](#)
- tm_sec
 - tm, [364](#)
- tm_wday
 - tm, [364](#)
- tm_yday
 - tm, [364](#)
- tm_year
 - tm, [364](#)
- toascii
 - ctype, [132](#)
- tolower
 - ctype, [132](#)
- tools, optional, [88](#)
- tools, required, [88](#)
- toupper
 - ctype, [132](#)
- trunc
 - avr_math, [159](#)
- truncf
 - avr_math, [154](#)
- TW_BUS_ERROR
 - util_twi, [318](#)
- TW_MR_ARB_LOST
 - util_twi, [318](#)
- TW_MR_DATA_ACK
 - util_twi, [318](#)
- TW_MR_DATA_NACK
 - util_twi, [318](#)
- TW_MR_SLA_ACK
 - util_twi, [318](#)
- TW_MR_SLA_NACK
 - util_twi, [318](#)
- TW_MT_ARB_LOST
 - util_twi, [318](#)
- TW_MT_DATA_ACK
 - util_twi, [318](#)
- TW_MT_DATA_NACK
 - util_twi, [318](#)
- TW_MT_SLA_ACK
 - util_twi, [318](#)
- TW_MT_SLA_NACK
 - util_twi, [319](#)
- TW_NO_INFO
 - util_twi, [319](#)
- TW_READ
 - util_twi, [319](#)
- TW_REP_START
 - util_twi, [319](#)
- TW_SR_ARB_LOST_GCALL_ACK
 - util_twi, [319](#)
- TW_SR_ARB_LOST_SLA_ACK
 - util_twi, [319](#)
- TW_SR_DATA_ACK
 - util_twi, [319](#)
- TW_SR_DATA_NACK
 - util_twi, [319](#)
- TW_SR_GCALL_ACK
 - util_twi, [319](#)
- TW_SR_GCALL_DATA_ACK
 - util_twi, [319](#)
- TW_SR_GCALL_DATA_NACK
 - util_twi, [319](#)
- TW_SR_SLA_ACK
 - util_twi, [320](#)
- TW_SR_STOP
 - util_twi, [320](#)
- TW_ST_ARB_LOST_SLA_ACK
 - util_twi, [320](#)
- TW_ST_DATA_ACK
 - util_twi, [320](#)
- TW_ST_DATA_NACK
 - util_twi, [320](#)
- TW_ST_LAST_DATA
 - util_twi, [320](#)
- TW_ST_SLA_ACK
 - util_twi, [320](#)
- TW_START
 - util_twi, [320](#)
- TW_STATUS
 - util_twi, [320](#)
- TW_STATUS_MASK
 - util_twi, [320](#)
- TW_WRITE
 - util_twi, [321](#)
- twi.h, [387](#)

- UBRR_VALUE
 - util_setbaud, [316](#)
- UBRRH_VALUE
 - util_setbaud, [316](#)
- UBRRL_VALUE
 - util_setbaud, [316](#)
- UINT16_C
 - avr_stdint, [169](#)
- UINT16_MAX
 - avr_stdint, [169](#)
- uint16_t
 - avr_stdint, [172](#)
- UINT32_C
 - avr_stdint, [169](#)
- UINT32_MAX
 - avr_stdint, [169](#)
- uint32_t
 - avr_stdint, [172](#)
- UINT64_C
 - avr_stdint, [169](#)
- UINT64_MAX
 - avr_stdint, [169](#)
- uint64_t
 - avr_stdint, [172](#)
- UINT8_C
 - avr_stdint, [169](#)
- UINT8_MAX
 - avr_stdint, [169](#)
- uint8_t
 - avr_stdint, [172](#)
- uint_farptr_t
 - avr_inttypes, [147](#)
- UINT_FAST16_MAX
 - avr_stdint, [169](#)
- uint_fast16_t
 - avr_stdint, [172](#)
- UINT_FAST32_MAX
 - avr_stdint, [169](#)
- uint_fast32_t
 - avr_stdint, [173](#)
- UINT_FAST64_MAX
 - avr_stdint, [170](#)
- uint_fast64_t
 - avr_stdint, [173](#)
- UINT_FAST8_MAX
 - avr_stdint, [170](#)
- uint_fast8_t
 - avr_stdint, [173](#)
- UINT_LEAST16_MAX
 - avr_stdint, [170](#)
- uint_least16_t
 - avr_stdint, [173](#)
- UINT_LEAST32_MAX
 - avr_stdint, [170](#)
- uint_least32_t
 - avr_stdint, [173](#)
- UINT_LEAST64_MAX
 - avr_stdint, [170](#)
- uint_least64_t
 - avr_stdint, [173](#)
- UINT_LEAST8_MAX
 - avr_stdint, [170](#)
- uint_least8_t
 - avr_stdint, [173](#)
- UINTMAX_C
 - avr_stdint, [170](#)
- UINTMAX_MAX
 - avr_stdint, [170](#)
- uintmax_t
 - avr_stdint, [173](#)
- UINTPTR_MAX
 - avr_stdint, [170](#)
- uintptr_t
 - avr_stdint, [173](#)
- ultoa
 - avr_stdlib, [201](#)
- ungetc
 - avr_stdio, [185](#)
- UNIX_OFFSET
 - avr_time, [219](#)
- USE_2X
 - util_setbaud, [316](#)
- Using the standard IO facilities, [350](#)
- util_atomic
 - ATOMIC_BLOCK, [305](#)
 - ATOMIC_FORCEON, [305](#)
 - ATOMIC_RESTORESTATE, [305](#)
 - NONATOMIC_BLOCK, [305](#)
 - NONATOMIC_FORCEOFF, [306](#)
 - NONATOMIC_RESTORESTATE, [306](#)
- util_crc
 - _crc16_update, [307](#)
 - _crc8_ccitt_update, [308](#)
 - _crc_ccitt_update, [308](#)
 - _crc_ibutton_update, [309](#)
 - _crc_xmodem_update, [310](#)
- util_delay
 - _delay_ms, [311](#)
 - _delay_us, [312](#)

- F_CPU, 311
- util_delay_basic
 - _delay_loop_1, 313
 - _delay_loop_2, 313
- util_parity
 - parity_even_bit, 314
- util_setbaud
 - BAUD_TOL, 316
 - UBRR_VALUE, 316
 - UBRRH_VALUE, 316
 - UBRRL_VALUE, 316
 - USE_2X, 316
- util_twi
 - TW_BUS_ERROR, 318
 - TW_MR_ARB_LOST, 318
 - TW_MR_DATA_ACK, 318
 - TW_MR_DATA_NACK, 318
 - TW_MR_SLA_ACK, 318
 - TW_MR_SLA_NACK, 318
 - TW_MT_ARB_LOST, 318
 - TW_MT_DATA_ACK, 318
 - TW_MT_DATA_NACK, 318
 - TW_MT_SLA_ACK, 318
 - TW_MT_SLA_NACK, 319
 - TW_NO_INFO, 319
 - TW_READ, 319
 - TW_REP_START, 319
 - TW_SR_ARB_LOST_GCALL_-
ACK, 319
 - TW_SR_ARB_LOST_SLA_ACK,
319
 - TW_SR_DATA_ACK, 319
 - TW_SR_DATA_NACK, 319
 - TW_SR_GCALL_ACK, 319
 - TW_SR_GCALL_DATA_ACK, 319
 - TW_SR_GCALL_DATA_NACK,
319
 - TW_SR_SLA_ACK, 320
 - TW_SR_STOP, 320
 - TW_ST_ARB_LOST_SLA_ACK,
320
 - TW_ST_DATA_ACK, 320
 - TW_ST_DATA_NACK, 320
 - TW_ST_LAST_DATA, 320
 - TW_ST_SLA_ACK, 320
 - TW_START, 320
 - TW_STATUS, 320
 - TW_STATUS_MASK, 320
 - TW_WRITE, 321
- utoa
 - avr_stdlib, 202
- vfprintf
 - avr_stdio, 186
- vfprintf_P
 - avr_stdio, 188
- vfscanf
 - avr_stdio, 188
- vfscanf_P
 - avr_stdio, 191
- vprintf
 - avr_stdio, 191
- vscanf
 - avr_stdio, 191
- vsnprintf
 - avr_stdio, 191
- vsnprintf_P
 - avr_stdio, 191
- vsprintf
 - avr_stdio, 191
- vsprintf_P
 - avr_stdio, 191
- wdt.h, 388
- wdt_reset
 - avr_watchdog, 301
- WDTO_120MS
 - avr_watchdog, 301
- WDTO_15MS
 - avr_watchdog, 301
- WDTO_1S
 - avr_watchdog, 301
- WDTO_250MS
 - avr_watchdog, 301
- WDTO_2S
 - avr_watchdog, 301
- WDTO_30MS
 - avr_watchdog, 301
- WDTO_4S
 - avr_watchdog, 302
- WDTO_500MS
 - avr_watchdog, 302
- WDTO_60MS
 - avr_watchdog, 302
- WDTO_8S
 - avr_watchdog, 302
- week
 - week_date, 365
- week_date, 365
 - day, 365

- week, [365](#)
 - year, [365](#)
- week_of_month
 - avr_time, [224](#)
- week_of_year
 - avr_time, [225](#)
- year
 - week_date, [365](#)