

amforth 4.8

Technical Documentation

Matthias Trute

amforth 4.8Technical Documentation

by Matthias Trute

Published 2011

Copyright © 2007, 2012 Matthias Trute

Table of Contents

Overview	vi
1. First Steps	1
1.1. User Interface	1
2. Hardware	2
2.1. Controller	2
2.2. Bootloader Support	2
2.3. Fuses	2
3. Source Organisation	3
3.1. Overview	3
3.2. Core system	4
3.2.1. Dictionary files	4
3.2.2. Device Settings	5
3.3. Application Code	5
4. Architecture	6
4.1. Overview	6
4.2. CPU -- Forth VM Mapping	6
4.3. Core System	7
4.3.1. Threading Model	7
4.3.2. Inner Interpreter	7
4.3.3. Interpreter	8
4.3.4. Stacks	10
4.3.5. Interrupts	10
4.3.6. Multitasking	12
4.3.7. Exception Handling	12
4.3.8. User Area	12
4.3.9. Word Lists and Environment Queries	13
4.4. Memory Layout	13
4.4.1. Flash	14
4.4.2. EEPROM	16
4.4.3. RAM	16
5. Implementation	19
5.1. ANS Words	19
5.1.1. Core and Core EXT	19
5.1.2. Block	19
5.1.3. Double Number	19
5.1.4. Exception	19
5.1.5. Facility	20
5.1.6. File Access	20
5.1.7. Floating Point	20
5.1.8. Locals	20
5.1.9. Memory Allocation	20
5.1.10. Programming Tools	20
5.1.11. Word Lists and Search Order	21
5.1.12. Strings	21

5.2. Amforth	22
5.2.1. COLD	22
5.2.2. MCU Access	22
5.2.3. Assembler	23
5.2.4. Memory	23
5.2.5. Input Output	23
5.2.6. Strings	24
6. Library	25
6.1. Hardware Access	25
6.2. Software Modules	25
6.2.1. Multitasking	25
6.2.2. TWI / I2C	26
6.2.3. I2C EEPROM	27
7. Tools	28
7.1. Host	28
7.1.1. Partdescription Converter	28
7.1.2. Documentation	28
7.1.3. Uploader	28
8. Roadmap	30
8.1. More ANS94 Words	30
8.2. More Controller Types	30
8.3. Support	30
8.4. Contributors	30

List of Tables

3-1. Source Filenames	3
4-1. Register Mapping	6
4-2. Extended Forth VM Register Mapping	6
4-3. USER Area	12

Overview

Amforth is a Forth system for the AVR ATmega microcontroller family. It works on the controller itself and does not depend on any additional hard- or software. It places no restrictions on using external hardware.

Amforth implements a large subset of the Forth standard ANS94. Most of the CORE and CORE EXT words and a varying number of words from the other word sets are implemented. It is very easy to extend or shrink the actual word list for a specific application by just editing the dictionary include files.

The dictionary is located in the flash memory. The built-in compiler extends it directly.

amforth provides full access to all interrupts. The interrupt handler routines can be code or forth words.

Amforth is published under the GNU General Public License version 2.

The name amforth has no special meaning.

Amforth is a new implementation. The first code was written in the summer of 2006. It is written "from scratch" using assembly language and forth itself. It does not have a direct relationship to any other forth system.

Chapter 1. First Steps

The first steps require an ATmega microcontroller with an RS232 connection to an PC or a terminal like the VT100 or similar hardware. A customization may change these requirements.

1.1. User Interface

amforth has a simple user interface. Connect your system to a serial terminal (or a PC) and you get the forth prompt >

```
> cold
amforth 4.4 ATmega16 8000 kHz
> words
nr> n>r (i!) !i @i @e !e nip not s>d up! up@ ...
>
```

Chapter 2. Hardware

2.1. Controller

amforth is designed to run on AVR Atmega microcontrollers. It requires ca 7KB flash memory for the basic system and can address 128KB of flash memory.

The ATtiny microcontrollers and a few ATmega types lack the minimum flash capacity. The ATtiny's some machine instructions as well.

2.2. Bootloader Support

Most bootloaders will not work with amforth since they do not provide an application programming interface to rewrite a single flash cell. The default setup will thus replace any bootloader found with some core routines.

It is possible to change the word **!i** to use an API and work with existing bootloaders. **!i** is a deferred word that can be re-targeted to more advanced words that may do address range checkings, write success checks or simply turn on/off LEDs to visualize the flash programming.

2.3. Fuses

Amforth uses the self programming feature of the ATmega microcontrollers to work with the dictionary. It is ok to use the factory default settings plus the changes for the oscillator settings. It is recommended to use a higher CPU frequency to meet the timing requirements of the serial terminal.

Fuses are the main cause for problems with the flash write operations. If the **!i** operation fails, make sure that the code for it is within the boot loader section. It is recommended to make the bootloader section as large as the NRWW section, otherwise the basic machine instruction `spm` may fail silently and the controller becomes unresponsive.

Chapter 3. Source Organisation

3.1. Overview

amforth is written using the standard Atmel AVR 8 bit assembly language. That does not mean that every word is actually written in assembly language however. Most of the words are written in forth itself, but are precompiled into the assembler syntax. This solves the chicken-and-egg problem: how to compile the compiler words.

The source code can be processed with both the AVR Studio and the linux avr assembler avra.

amforth consists of a great number of small source files. Nearly all words are coded in their own source files. These files are organized with include files, named after the pattern `dict*.inc`.

The include hierarchy is as follows: Top level is the application specific file (`template.asm`). It included the file (`core/`)`amforth.asm` only. This file includes the two (`appl/`) files `dict_appl.inc` for the low address words (RWW space) and the file `dict_appl_core.inc` for the high address space words (NRWW). These two include files use other `dict*` files and direct words definition files. Any cross-references are solved by the assembler. The user needs to take care that the highest used address is within the flash address range, an assembler error message is generated in that case.

Currently the following pre defined dict files exist:

Table 3-1. Source Filenames

Filename	purpose	placement
<code>dict_minimum.inc</code>	All basic words beyond the primitives.	Usually included via <code>dict_appl.inc</code> and placed in the RWW section. If possible put it into the NRWW section however.
<code>dict_mcu.inc</code>	Microcontroller specific primitives.	It can be placed everywhere. Usually in <code>dict_appl_core.inc</code> .
<code>dict_core.inc</code>	All essential primitives.	Always near the inner interpreter (NRWW, <code>dict_appl_core.inc</code>).
<code>dict_compiler</code>	All Compiler words.	Usually part of the <code>dict_appl.inc</code> .
<code>dict_usart</code>	Usart Terminal IO words	Usually part of the <code>dict_appl.inc</code> .
<code>dict_vm</code>	Not-Yet proposed additional VM registers A and B	Usually part of the <code>dict_appl.inc</code> .
<code>dict_wl</code>	words dealing with word lists.	Usually part of the <code>dict_appl.inc</code>

The order in which the files are included defines the search order and there location within the flash

memory. Most words can be moved from one include file to another to optimize the flash usage.

There are additional files: `amforth.asm` and `macros.asm`. The first one is the master file and the only one the application needs to include. The file `macros.asm` contains some useful assembler macros that make the source code easier to read. `user.inc` contains the layout of the system user area.

3.2. Core system

The file `amforth.asm` is the core of `amforth`. Here is the startup code for the microcontroller, and the forth inner interpreter with the interrupt service routine. It includes the dictionary files.

3.2.1. Dictionary files

The dictionary files have two tasks: First they include the word definition files. Second, they determine each word's location in the resulting flash layout.

There are a few words left out from the dictionary lists. These words are either not always needed or are some variants of existing words or simply cannot be included in the core system due to size limitations in the NRWW section with smaller atmegas. They are usually included by the application specific include file(s).

3.2.1.1. `dict_core.inc`

The file `dict_core.inc` contains all words for the NRWW flash section, Since the word **I!** cannot write to this address range, no new words can be compiled to this section at runtime. Thus it is advisable to include as many words as possible in `dict_core.inc` if the amount of writable dictionary space is an issue. As a helper the file `dict_appl_core.inc` can be used to place application specific words in the core area.

3.2.1.2. `dict_minimum.inc`

A useful forth system needs in addition to the above at least the file `dict_minimum.inc`, which includes the forth interpreter words.

3.2.1.3. `dict_compiler.inc`

An almost complete forth system with a compiler gives the third include file: `dict_compiler.inc`.

3.2.1.4. `dict_appl.inc` and `dict_appl_core.inc`

Some words have their source files within the `core/words` directory but have to be included via the `dict_appl.inc` or `dict_appl_core.inc` files. These words may provide the hardware dependencies to access the amforth system. The serial line terminal is an example.

3.2.2. Device Settings

Every Atmega has its own specific settings. They are based on the official include files provided by Atmel and define the important settings for the serial IO port (which port and which parameters), the interrupt vectors and some macros.

Adapting another ATmega microcontroller is as easy as copy and edit an existing file from a similar type.

The last definition is a string with the device name in clear text. This string is used within the word **VER**.

3.3. Application Code

Every build of amforth needs an application. There are a few sample applications, which can be used either directly (AVR Butterfly) or serve as a source for inspiration (template application).

The structure is basically always the same. First the file `macros.asm` has to be included. After that some definitions need to be done: The size of the Forth buffers, the CPU frequency, initial terminal settings etc. Then the device specific part needs to be included and as the last step the amforth core is included.

For a comfortable development cycle the use of a build utility such as **make** or **ant** is recommended. The assembler needs a few settings and the proper order of the include directories.

Chapter 4. Architecture

4.1. Overview

amforth is a 16 bit Forth implementing the indirect threading model. The flash memory contains the whole dictionary. A few EEPROM cells are used to hold initial values and the dictionary pointers. The RAM contains buffers, variables and the stacks.

The compiler is a classic compiler without any optimization support.

amforth uses all of the CPU registers to hold vital data: The data stack pointer, the instruction pointer, the user pointer, and the Top-Of-Stack cell. The hardware stack is used as the return stack. Some registers are used for temporary data in primitives.

4.2. CPU -- Forth VM Mapping

The Forth VM has a few registers that need to be mapped to the microcontroller registers. The mapping has been extended over time and may cover all available registers. The actual coverage depends on the amount of additional packages. The default settings are shown in the table Register Mapping.

Table 4-1. Register Mapping

Forth Register	ATmega Register(s)
W: Working Register	R22:R23
IP: Instruction Pointer	XH:XL (R27:R26)
RSP: Return Stack Pointer	SPH:SPL
PSP: Parameter Stack Pointer	YH:YL (R29:R28)
UP: User Pointer	R4:R5
TOS: Top Of Stack	R24:R25
X: temporary (scratch pad) register	ZH:ZL (R31:R30)

Table 4-2. Extended Forth VM Register Mapping

Forth Register	ATmega Register(s)
A: Index and Scratch Register	R6:R7
B: Index and Scratch Register	R8:R9

In addition the register pair R0:R1 is used internally e.g. to hold the the result of multiply operations. The register pair R2:R3 is used as the zero value in many words. These registers must never be changed.

The registers from R10 to R13 are currently unused, but may be used for the VM extended registers X and Y sometimes. The registers R14 to R21 are used as temporary registers and can be used freely within one module as temp0 to temp7.

The forth core uses the **T** Flag in the machine status register SREG for signalling an interrupt. Any other code must not change that bit.

4.3. Core System

4.3.1. Threading Model

amforth implements the classic indirect threaded variant of forth.

4.3.2. Inner Interpreter

For the indirect threading model an inner interpreter is needed. The inner interpreter does the interrupt handling too.

4.3.2.1. EXECUTE

This operation reads the cell the IP currently points to and uses the value read as the destination of a branch. This EXECUTE is not the forth word EXECUTE. The forth EXECUTE sets the IP from the data stack TOS element.

4.3.2.2. NEXT

The NEXT routine is the core of the inner interpreter. It consists of 4 steps which are executed for every forth word.

The first step in NEXT is to check whether an interrupt needs to be handled. It is done by looking at the **T** flag in the machine status register. If it is set, the code jumps to the interrupt handling part. If the flag is cleared the following normal NEXT routine runs.

The next step is to read the cell the IP points to and stores this value in the W register. For a COLON word W contains the address of the code field.

The 3rd step is to increase the IP register by 1.

The 4th step is to read the content of the cell the W register points to. The value is stored in the scratch pad register X. The data in X is the address of the machine code to be executed in the last step.

This last step finally jumps to the machine code pointed to by the X scratch pad register.

4.3.2.3. DO COLON

DO COLON (aka NEST) first pushes the IP (which points to the next word to be executed when the current word is done) to the return stack. It then increments W by one flash cell, so that it points to the body of the (colon) word, and sets IP to point to that value. Then it continues with NEXT, which begins executing the words in the body of the (parent) colon word.

4.3.2.4. EXIT

The code for EXIT (aka UNNEST) is the forth word **EXIT** in the dictionary. It reads the IP from the return stack and jumps to NEXT. The return stack pointer is incremented by 2 (1 flash cell).

4.3.2.5. DO_DOES

This code is the runtime part of the forth word **DOES>**. It pushes the current address of the MCU IP register onto the returnstack and jumps to DO_DOES. DO_DOES gets that address back, saves the current IP and sets the forth IP to the address it got from the stack. Finally it continues with NEXT.

4.3.3. Interpreter

The interpreter is a line based command interpreter. It based upon **REFILL** to acquire the next line of characters, located at a position **SOURCE** points to. While processing the line, the pointer **>IN** is adjusted accordingly. Both words **REFILL** and **SOURCE** are USER based deferred words which allows to use any input source on a thread specific level. The interpreter itself does not use any static buffers or variables (**>IN** is a USER variable as well).

A given string is handled by **INTERPRET** which splits it into whitespace delimited words. Every word is processed using a list of recognizers. Processing ends either when the string end is reached or an exception occurs.

4.3.3.1. SOURCE and REFILL

SOURCE provides an addr/len string pair that does not change during processing. The task of **REFILL** is to fill the string buffer, **SOURCE** will point to when finished.

There is one default input source: The terminal input buffer. This buffer gets filled with **REFILL-TIB** that reads from the serial input buffers (**KEY**). **SOURCE** points to the Terminal Input Buffer itself. Another input source are plain strings, used by **EVALUATE**.

4.3.3.2. Recognizer

A recognizer gets the word address (API is not yet finalized) of the current word. If the word can be processed, the recognizer is responsible to do so. A word from the dictionary has to be either executed or compiled, a number as well. A recognizer must not change the word buffer content. Finally the recognizer returns a flag to the interpreter which signals success or not. This flag is consumed, any other stack change is kept.

Three recognizers are defined: Dictionary lookup **rec-find**, integer number conversion **rec-intnum** and a not-found dummy **rec-notfound**. The first two take care of the interpreter state and either leave more than the flag (e.g. the number entered) or compile the information to the dictionary.

The not-found recognizer prints the word and throws an exception -13 which can be caught.

The list of the recognizers is kept in the EEPROM, the maximum size of the entries is a compile time setting (currently 6 slot are available).

4.3.3.2.1. Example Recognizer

A recognizer gets the address of a counted word in memory and leaves at least the flag for the interpreter. If any data is to be left on the stack (e.g. numeric values) it has to be beneath the flag.

The small example illustrates the integration of the floating point library for amforth. It is based upon a conversion word **>float** which takes a string and tries to convert it into a float. The word **fliteral** compiles a floating point number into the dictionary.

```
: rec-float \ addr -- (f|) -1 | 0
  count >float
  if state @ if postpone fliteral then -1 else 0
  then ;
```

The recognizer first tries to convert the string to a number. If that fails, the flag from the **>float** is essentially duplicated and the recognizer is left. If the conversion succeeded, the floating point number is

on the data stack. The recognizer now checks whether the number needs to be compiled or not. In any case the success flag is returned.

Future versions may change the API from counted strings at `addr` to `addr/len` information.

4.3.4. Stacks

4.3.4.1. Data Stack

The data stack uses the CPU register pair YH:YL as its data pointer. The Top-Of-Stack element (TOS) is in a register pair. Compared to a straight forward implementation this approach saves code space and gives higher execution speed (approx 10-20%). Saving even more stack elements does not really provide a greater benefit (much more code and only little speed enhancements).

The data stack starts at a configurable distance below the return stack (RAMEND) and grows downward.

4.3.4.2. Return Stack

The Return Stack is the hardware stack of the controller. It is managed with push/pop assembler instructions. The default return stack starts at RAMEND und grows downward.

4.3.5. Interrupts

Amforth routes the low level interrupts into the forth inner interpreter. The inner interpreter switches the execution to a predefined word if an interrupt occurs. When that word finishes execution, the interrupted word is continued. The interrupt handlers are completely normal forth colon words without any stack effect.

4.3.5.1. Example Interrupt Handling

The example illustrates the basic usage of interrupts. The code implements a very basic timer functionality.

We use the timer/counter 0 to generate an interrupt any time the overflow condition is triggered. It is a 8bit counter, the input is the internal oscillator using a prescaler of 1024. The overflow value is set to 256. This gives an interrupt rate of $\text{frequency}/(1024*256)$. At 8Mhz this will be 32 interrupts per second.


```

variable tick

\ this code is executed as an interrupt
: timer-int-isr
  1 tick +!
;

: timer-init
  5 TCCR0 c! \ prescaler 1024, check data sheet
  ['] timer-int-isr TIMER0_OVFAddr int!
  0 tick !
;

\ turn on the timer, needs timer-init already in place
: +timer
  1 TIMSK c!
;

\ stops the timer
: -timer
  0 TIMSK c!
;

```

To ease debugging the code, the word `int-trap` can be used. This word simulates the interrupt in any respect but a real hardware cause. To test the routine from above enter the command

```
> TIMER0_OVFAddr int-trap tick @ .
```

, it should increment tick by 1

4.3.5.2. Implementation Details

The processing of interrupts takes place in two steps: The first one is the low level part. It is called whenever an interrupt occurs. The code is the same for all interrupts. It takes the number of the interrupt from its vector address and stores this in a RAM cell. Then the low level ISR sets the **T** flag in the status register of the controller and returns with `RET`.

The second step is taken from the inner interpreter. It checks the T-flag every time it is entered and, if it is set, it switches to interrupt handling at forth level. This approach has a penalty of 1 CPU cycle for checking and skipping the branch instruction to the isr forth code if no interrupt occurred.

If an interrupt is detected, the forth VM clears the T-flag and continues with the word **ISR-EXEC**. This word reads the currently active interrupt number and calls the associated execution token. When this

word is finished, the word **ISR-END** is called. This word clears the interrupt flag for the controller (RETI).

This interrupt processing has two advantages: There are no lost interrupts (the controller itself disables interrupts within interrupts and re-transmits newly discovered interrupts afterwards) and it is possible to use standard forth words to deal with any kind of interrupts.

Interrupts from some hardware sources (e.g. the usart) need to be cleared from the Interrupt Service Routine. If this is not done within the ISR, the interrupt is re-triggered immediatly after the ISR returned control.

The downside is a relatively long latency since the the forth VM has to be synchronized with the interrupt handling code in order to use normal colon words as ISR. This peanalty is usually small since only words in assembly can cause the delay.

4.3.6. Multitasking

amforth does not implement multitasking directly. It only provides the basic functions. Within IO words the deferred word **PAUSE** is called whenever possible. This word is initialized to do nothing (**NOOP**).

4.3.7. Exception Handling

amforth implements the **CATCH** and **THROW** exception handling. The outermost catch frame is located at the interpreter level in the word **QUIT** . If an exception with the value -1 or -2 is thrown, **QUIT** will print a message and re-start itself. Other values silently restart **QUIT** .

4.3.8. User Area

The User Area is a special RAM storage area. It contains the **USER** variables and the User deferred definitions. Access is based upon the value of the user pointer **UP**. It can be changed with the word **UP!** and read with **UP@** . The **UP** itself is stored in a register pair.

The size of the user area is determind by the size the system itself uses plus a configurable number at compile time. For self defined tasks this user supplied number can be changed for task local variables.

The first **USER** area is located at the first data address (usually **RAMSTART**).

Table 4-3. USER Area

Addressoffset (bytes)	Purpose
0	Multitasker Status
2	Multitasker Follower
4	RP0
6	SP0
8	SP (m
10	HANDLER (exception handling)
12	BASE (number conversion)
14	EMIT (character IO)
16	EMIT? (character IO)
18	KEY (character IO)
20	KEY? (character IO)
22	REFILL (deferred)
24	SOURCE (deferred)

The User Area is used to provide task local information. Without an active multitasker it contains the starting values for the stackpointers, the deferred words for terminal IO, the BASE variable and the exception handler.

The multitasker uses the first 2 cells to store the status and the link to the next entry in the task list. In that situation the user area is/can be seen as the task control block.

Beginning with release 3.7 the USER area has been split into two parts. The first one called system user area contains all the variables described above. The second one is the application user area that contains all variables defined with the USER command. The default application user area is empty and by default of size zero.

4.3.9. Word Lists and Environment Queries

Wordlists and environment queries are implemented using the same structure. They are based upon the simple linked list built with **create**. The wordlist identifier is a EEPROM address that holds the starting point address for the wordlist search.

Environment queries are normal colon words. They are called within **environment?** and leave their results at the data stack.

find uses an array of wordlist identifiers to search for the word. This list can be accessed with **get-order** as well.

4.4. Memory Layout

4.4.1. Flash

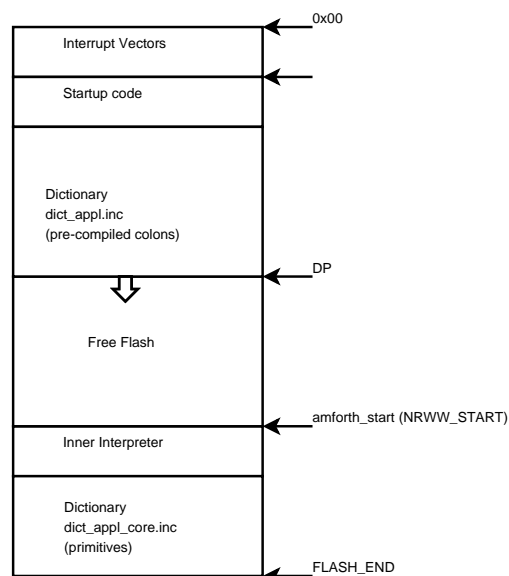
The flash memory is divided into 4 sections. The first section, starting at address 0, contains the interrupt vector table for the low level interrupt handling and a character string with the name of the controller in plain text.

The 2nd section contains the low level interrupt handling routines. The interrupt handler is very closely tied to the inner interpreter. It is located near the first section to use the faster relative jump instructions.

The 3rd section is the first part of the dictionary. Nearly all colon words are located here. New words are appended to this section. This section is filled with FFFF cells when flashing the controller initially. The current write pointer is the **DP** pointer.

The last section is identical to the boot loader section of the ATmegs. It is also known as the NRWW area. Here is the heart of amforth: The inner interpreter and most of the words coded in assembly language.

Figure 4-1. FLASH Structure Overview



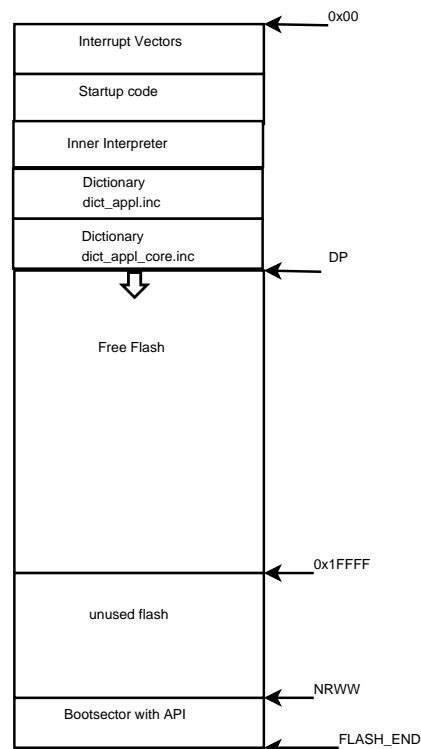
The reason for this split is a technical one: to work with a dictionary in flash the controller needs to write to the flash. The ATmega architecture provides a mechanism called self-programming by using a special

instruction and a rather complex algorithm. This instruction only works in the boot loader/NRWW section. amforth uses this instruction in the word I!. Due to the fact that the self programming is a lot more than only a simple instruction, amforth needs most of the forth core system to achieve it. A side effect is that amforth cannot co-exist with classic bootloaders. If a particular boot loader provides an API to enable applications to call the flash write operation, amforth can be restructured to use it. Currently only very few and seldom used bootloaders exist that enable this feature.

Atmegas can have more than 64 KB Flash. This requires more than a 16 bit address, which is more than the cell size. For one type of those bigger atmegas there will be an solution with 16 bit cell size: Atmega128 Controllers. They can use the whole address range with an interpretation trick: The flash addresses are in fact not byte addresses but word addresses. Since amforth does not deal with bytes but cells it is possible to use the whole address range with a 16 bit cell. The Atmegas with 128 KBytes Flash operate slightly slower since the address interpretation needs more code to access the flash (both read and write). The source code uses assembly macros to hide the differences.

An alternative approach to place the elements in the flash shows picture . Here all code goes into the RWW section. This layout definitely needs a routine in the NRWW section that provides a cell level flash write functionality. The usual bootloaders do not have such an runtime accessible API, only the DFU bootloader from atmel found on some USB enabled controllers does.

Figure 4-2. Alternative FLASH Structure



The unused flash area beyond 0x1FFFF is not directly accessible for amforth. It could be used as a block device.

4.4.1.1. Flash Write

The word performing the actual flash write operation is **!!** (i-store). This word takes the value and the address of a single cell to be written to flash from the data stack. The address is a word address, not a byte address!

The flash write strategy follows Atmel's appnotes. The first step is turning off all interrupts. Then the affected flash page is read into the flash page buffer. While doing the copying a check is performed whether a flash erase cycle is needed. The flash erase can be avoided if no bit is turned from 0 to 1. Only if a bit is switched from 0 to 1 must a flash page erase operation be done. In the fourth step the new flash data is written and the flash is set back to normal operation and the interrupt flag is restored. The whole process takes a few milliseconds.

This write strategy ensures that the flash has minimal flash erase cycles while extending the dictionary. In addition it keeps the forth system simple since it does not need to deal with page sizes or RAM based buffers for dictionary operations.

4.4.2. EEPROM

The built-in EEPROM contains vital dictionary pointer and other persistent data. They need only a few EEPROM cells. The remaining space is available for user programs. The easiest way to use EEPROM is the use of forth VALUES. Their intended design pattern (read often, write seldom) is like that for the typical EEPROM usage.

Another use for EEPROM cells is to hold execution tokens. The default system uses this for the turnkey vector. This is an EEPROM variable that reads and executes the XT at runtime. It is based on the DEFER/IS standard. To define a deferred word in the EEPROM use the Edefer definition word. The standard word IS is used to put a new XT into it.

Low level space management is done through the EDP variable. This is not a forth value but a EEPROM based variable. To read the current value an @e operation must be used, changes are written back with !e. It contains the highest EEPROM address currently allocated. The name is based on the DP variable, which points to the highest dictionary address.

4.4.3. RAM

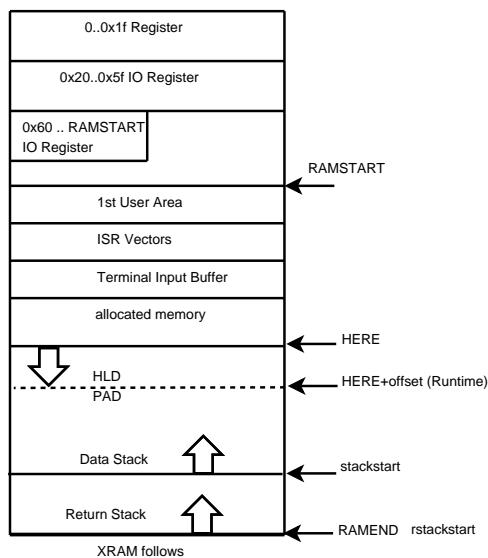
The RAM address space is divided into three sections: the first 32 addresses are the CPU registers. Above come the IO registers and extended IO registers and finally the RAM itself.

amforth needs very little RAM space for its internal data structures. The biggest part are the buffers for the terminal IO. In general RAM is managed with the words **VARIABLE** and **ALLOT**.

Forth defines a few transient buffer regions for various purposes. The most important is PAD, the scratch buffer. It is located 100 bytes above the current **HERE** and goes to upper addresses. The Pictured Numeric Output is just at PAD and grows downward. The word **WORD** uses the area above **HERE** as its buffer to store the just recognized word from **SOURCE**.

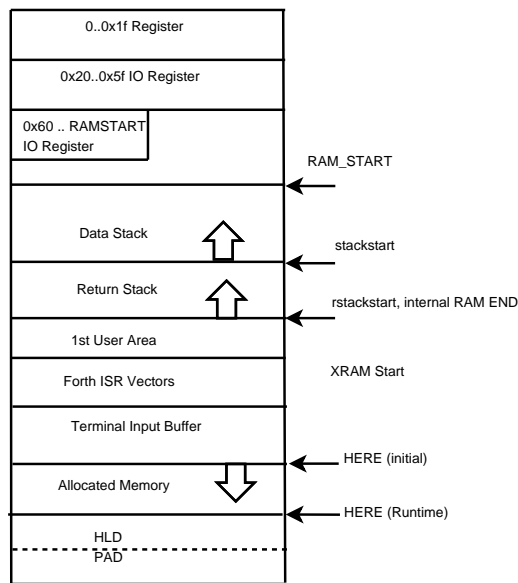
Figure 4-3 shows a ram layout that can be used on systems without external RAM. All elements are located within the internal memory pool.

Figure 4-3. RAM Structure Overview



Another layout, that makes the external RAM easily available is shown in Figure 4-4. Here are the stacks at the beginning of the internal RAM and the data space region. All other buffers grow directly into the external data space. From an application point of view there is not difference but a speed penalty when working with external RAM instead of internal.

Figure 4-4. RAM Structure Overview



With amforth all three sections can be accessed using their RAM addresses. That makes it quite easy to work with words like `C@`. The word `!` implements a LSB byte order: The lower part of the cell is stored at the lower address.

For the RAM there is the word **Rdefer** which defines a deferred word, placed in RAM. As a special case there is the word **Udefer**, which sets up a deferred word in the user area. To put an XT into them the word **IS** is used. This word is smart enough to distinguish between the various Xdefer definitions.

Chapter 5. Implementation

5.1. ANS Words

amforth is close to the ANS94 Forth standard. The main difference comes from the fact that the AVR ATmegs use a Havard architecture (separate code and data address space) that amforth does not hide. amforth gives full and unmodified access to the whole address space.

amforth implements most or all words from the ANS word sets CORE, CORE EXT, EXCEPTION and DOUBLE NUMBERS. A loadable floating point library that contains the basic routines is available. Words from the word sets LOCALS and FILE-ACCESS are dropped completely. The others are partially implemented.

5.1.1. Core and Core EXT

All words from the CORE word set are available. CORE EXT drops the words C", COMPILE, CONVERT, EXPECT, SPAN, and ROLL.

The following words have non-standard behavior

Loop counters are checked on signed compares.

5.1.2. Block

amforth has limited block support with I2C/TWI serial eeprom chips with 2 byte addresses.

5.1.3. Double Number

Double cell numbers work as expected. Not all words are implemented. Entering them directly using the dot- notation work for dots at the end of the number, not if the dot is somewhere within it.

5.1.4. Exception

Exceptions are fully supported. The words **ABORT** and **ABORT"** use them internally.

The **THROW** codes -1, -2 and -13 work as specified.

The implementation is based upon a variable **HANDLER** which holds the current return stack pointer position. This variable is a **USER** variable.

5.1.5. Facility

The basic system uses the **KEY?** and **EMIT?** words as deferred words in the **USER** area.

The word **MS** is implemented with the word **1MS** which busy waits almost exactly 1 millisecond. The calculation is based upon the frequency specified at compile time.

The words **TIME&DATE** , **EKEY** , **EKEY>CHAR** are not implemented.

To control a VT100 terminal the words **AT-XY** and **PAGE** are written in forth code. They emit the ANSI control codes according to the VT100 terminal codes.

5.1.6. File Access

amforth does not have filesystem support. It does not contain any words from this word set.

5.1.7. Floating Point

amforth has a loadable floating point library. It contains the basic words to deal with single precision floats. The floats are managed on the standard data stack. After loading the library floats can be entered directly at the command prompt. Some speed sensitive words are available as assembly code as well.

5.1.8. Locals

amforth does not currently support locals.

5.1.9. Memory Allocation

amforth does not support the words from the memory allocation word set.

5.1.10. Programming Tools

Variants of the words **.S** , **?** and **DUMP** are implemented or can easily be done. The word **SEE** won't be supported since amforth uses a optimization strategy to strip forth headers whenever possible. The other reason for dropping **SEE** is that amforth is OpenSource software. If your vendor does not disclose the full source, let me know. He violates the GPL. Nevertheless a word **XT2NFA** exists that leads to the name field.

STATE works as specified.

The word **WORDS** does not sort the word list and does not take care of screen sizes.

The words **;CODE** and **ASSEMBLER** are not supported. amforth has a loadable assembler which can be used with the words **CODE** and **END-CODE** .

The control stack commands **CS-ROLL** , **CS-PICK** and **AHEAD** are not implemented. The compiler words operate with the more traditional **MARK** / **RESOLVE** word pairs.

FORGET is not implemented since it would be nearly impossible to reset the search order word list with reasonable efforts. The better way is using **MARKER** from the library.

An **EDITOR** is not implemented.

[IF] , **[ELSE]** and **[THEN]** are not implemented.

5.1.11. Word Lists and Search Order

Amforth supports the ANS Search Order Wordlist. A word list consist of a linked list of words in the dictionary. There are no limits on the number of word lists defined. Only the length of the active search order is limited: There can be up to 8 entries at any given moment. This limit can be changed at compile time in the application definition file.

Internally the wordlist identifier is the address where the wordlist start address is stored in the EEPROM. Creating a new word list means to allocate a new EEPROM cell. Since the ANS standard does not give named wordlist there is library code available that uses the old fashioned vocabulary.

5.1.12. Strings

SLITERAL , **CMOVE>** , **CMOVE** , **COMPARE>** , and **/STRING** are implemented.

-**TRAILING** , **BLANK** , and **SEARCH** are not implemented.

5.2. Amforth

5.2.1. COLD

The startup code is in the file `cold.asm`. It gets called directly from the address 0 vector.

This assembly part of the startup code creates the basic runtime environment to start the virtual forth machine. It sets up the stack pointers and the user pointer and places the forth instruction pointer on the word **WARM**. Then it boots the forth virtual machine by jumping to the inner interpreter.

The start addresses of the stacks are placed to the user area for later use as well.

5.2.1.1. WARM

The word **WARM** is the high level part of the forth VM initialization. When called from within forth it is the equivalent to a **RESET**. **WARM** initializes the **PAUSE** deferred word to do nothing, calls the application defined **TURNKEY** action and finally hands over to **QUIT**.

5.2.1.2. TURNKEY

The **turnkey** is a EEPROM deferred word that points to an application specific startup word.

Its main task is to initialize the character IO to enable the forth interpreter to interact with the command prompt. The examples shipped with amforth do this by "opening" the serial port, switching to decimal number conversion and setting up the character IO deferred words (**KEY**, **EMIT** etc).

5.2.1.3. QUIT

QUIT repeats the initialization of the stack pointers by reading them from the user area and enters the traditional **ACCEPT -- INTERPRET** loop that never ends.

5.2.2. MCU Access

amforth provides wrapper words for the microcontroller instructions **SLEEP** and **WDR** (watch dog reset). To work properly, the MCU needs more configuration. amforth itself does not call these words.

Microcontrollers supporting the JTAG interface can be programmed to turn off JTAG at runtime. Similiar the watch dog timer can be disabled. Since both actions require strict timing they need to be implemented as primitives: **-JTAG** and **-WDT** .

5.2.3. Assembler

Lubos Pekny has written an assembler for amforth. To support it, amforth provides the two words **CODE** and **END-CODE**. The first creates a dictionary entry and sets the code field to the data field address. The interpreter will thus jump directly into the data field assuming some machine code there. The word **END-CODE** places a **JUMP NEXT** into the data field. This finishes the machine instruction execution and jumps back to the forth interpreter.

5.2.4. Memory

Atmega microcontroller have three different types of memory. RAM, EEPROM and Flash. The words **@** and **!** work on the RAM address space (which includes IO Ports and the CPU register), the words **@e** and **!e** operate on the EEPROM and **@i** and **!i** deal with the flash memory. All these words transfer one cell (2 bytes) between the memory and the data stack. The address is always the native address of the target storage: byte-based for EEPROM and RAM, word-based for flash. Therefore the flash addresses 64KWords or 128 KBytes address space.

External RAM shares the normal RAM address space after initialization (which can be done in the turnkey action). It is accessible without further changes.

For RAM only there is the special word pair **c@** / **c!** which operate with the lower half of a stack cell. The upper byte is either ignored or set to 0 (zero).

All other types of external memory need special handling, which may be masked with the block word set.

5.2.5. Input Output

amforth uses character terminal IO. A serial console is used. All IO is based upon the standard words **EMIT** / **EMIT?** and **KEY** / **KEY?** . Additionally the word **/KEY** is used to signal the sender to stop. All these words are deferred words in the USER area and can be changed with the **IS** command.

The predefined words use an interrupt driven IO with a buffer for input and output. They do not implement a handshake procedure (XON/XOFF or CTS/RTS). The default terminal device is selected at compile time.

These basic words include a call to the **PAUSE** command to enable the use of multitasking.

Other IO depend on the hardware connected to the microcontroller. Code exists to use LCD and TV devices. CAN, USB or I2C are possible as well. Another use of the redirect feature is the following: consider some input data in external EEPROM (or SD-Cards). To read it, the words **KEY** and **KEY?** can be redirected to fetch the data from them.

5.2.6. Strings

Strings can be stored in two areas: RAM and FLASH. It is not possible to distinguish between the storage areas based on the addresses found on the data stack, it's up to the developer to keep track.

Strings are stored as counted strings with a 16 bit counter value (1 flash cell) Strings in flash are compressed: two consecutive characters (bytes) are placed into one flash cell. The standard word **S''** copies the string from the RAM into flash using the word **S**, .

Chapter 6. Library

Amforth does not have a formal library concept. Amforth has a lot of forth code that can be seen as a library of words and commands.

6.1. Hardware Access

In the `device/` subdirectory are the controller specific register definitions. They are taken directly from the appnotes from Atmel. The register names are all uppercase. It is recommended to extract only the needed definitions since the whole list occupy a lot of flash memory.

Some commonly used lowlevel words can be included with the `dict_mcu.inc` include file at compile time.

6.2. Software Modules

6.2.1. Multitasking

The Library contains a cooperative multitasker in the file `multitask.frt`. It defines a command **multitaskpause** which can assigned to **pause** : `' multitaskpause is pause`

The multitasker has the following commands

onlytask (--)

Initialize the task system. The current task is placed as the only task in the task list.

alsotask (tid --)

Append a newly created task to the task list. A running multitasker is temporarily stopped. Make sure that the status of the task is sleep.

task (dstacksize rstacksize -- tid)

Allocate RAM for the task control block (aka user area) and the two stacks. Initializes the whole user area to direct IO to the serial line. The task has still no code associated and is not inserted to the task list.

task-sleep (tid --)

Let the (other) task sleep. The task switcher skips the task on the next round. When a task executes this command for itself, the task continues until the next call of **pause** .

task-awake (tid --)

The task is put into runnable mode. It is not activated immediately.

activate (tid --)

Skip all of the remaining code in the current colon word and continue the skipped code as task when the task list entry is reached by the multitasker.

It is possible to use a timer interrupt to call the command **pause** and turn the cooperative multitasker into a preemptive one. The latency is in the worst case that of the longest running uninterruptable forth commands: **1ms** , **!i** and **!i** . For a preemptive task switcher a lot more tools like semaphores may be needed.

6.2.2. TWI / I2C

The file `twi.frt` contains the basic words to operate with the hardware TWI module of the microcontroller. The file `twi-eeeprom.frt` uses these words to implement a native block buffer access for I2C EEPROMs with 2byte addresses.

The word **+twi** initializes the TWI hardware module with the supplied parameters. **-twi** turns the module off. The start-stop conditions are sent with the **twi.start** and **twi.stop** words. Data is transferred with the three words **twi.tx** for transmitting (sending) a byte, **twi.rx** for reading a byte (and sending an ACK signal) and **twi.rxn** for reading a byte and sending a NACK signal.

The command **twi.status** fetches the TWI status register, the command **twi.status?** compares the status with a predefined value and throws the exception -14 if they do not match.

The command **twi.scan** scans the whole (7 bit) address range and prints the address of any device found.

6.2.3. I2C EEPROM

I2C EEPROMs can be used in various ways. The file `twi-eeeprom.frt` defines words to access the EEPROM at byte address level and at block level. A page is the native block size of the eeprom device, that is stored in the **VALUE twi.ee-b/blk**. The hardware (i2c-) device address is stored in the value **twi.ee-addr**. Currently EEPROM devices with 2byte addresses are supported.

Byte level access is done with the words **twi.ee-c!** and **twi.ee-c@**. They transfer one byte from/to the eeprom address given. The stack diagram is exactly the same as for the RAM **c@ / c!**. Every store operation performs an full EEPROM erase/write cycle.

To transfer more bytes the block level words can be used. The transfer a whole EEPROM page to/from RAM. The first page is at address 0, page 1 starts at address **twi.ee-b/blk**.

Chapter 7. Tools

7.1. Host

There are a few number of tools on the host side (PC) that are specifically written to support amforth. They are written in script languages like perl and python and should work on all major operating systems. They are not needed to use amforth but may be useful.

7.1.1. Partdescription Converter

The **pd2amforth.pl** script reads a part description file in XML format (comes with the Atmel Studio package) and produces the controller specific `devices/controllername/*` files.

7.1.2. Documentation

The tool **makerefcard** reads the assembly files from the `words` subdirectory and creates a reference card. The resulting LaTeX file needs to be processed with **latex** to generate a nice looking overview of all words available in the amforth core system.

The command **make-htmlwords** creates the linked overview of all words on the amforth homepage.

7.1.3. Uploader

To transfer forth code to the microcontroller some precautions need to be taken. During a flash write operation all interrupts are turned off. This may lead to lost characters on the serial line. One solution is to send very slowly and hope that the receiver gets all characters. The program **ascii-xfr** can do the job:

```
ascii-xfr -s -c $delay_char -l $delay_line $file > $tty
```

This works but the upload of longer files needs a very long time: `$delay_char` can be 1 or 2 ms, `$delay_line` around 800 ms.

Another solution is **amforth-upload.py**. It was initially created by user *pix* (<http://pix.test.at/>). His algorithm checks for the echo of every character sent to the controller. At line ends the uploader waits for the ok prompt to continue with the next line.

This algorithm works very fast without the risk of lost characters. An extension of this script provides limited library support. In the source files a command

#include filename

is used to upload the content of `filename` instead of the two words. The sources will only work with this uploader utility, others will trigger the "word-not-found" exception on the microcontroller unless they recognize the `#include` syntax (similar to the `c` preprocessor).

Chapter 8. Roadmap

8.1. More ANS94 Words

Amforth supports the ANS94 Core Wordset. Other word sets are more or less completely supported as well. Some deprecated words (like **expect** and **query** are left out however.

Support for Blocks may be useful. It is not trivial to implement a standard 1KB block buffer on an Atmega with only 1KB RAM. It can be useful to deploy block sizes smaller than 1KB to match the native block sizes of the attached storage devices: serial EEPROM have e.g. 64 bytes, SD-Cards have (usually) 512 bytes. Some rather simple code can be used from the library for I2C/TWI EEPROM modules with native block sizes.

8.2. More Controller Types

Amforth can be used on any atmega with at least 8KB flash memory. A system with at least 16Kb is recommended. There are pre-generated driver files from the XML part description files. Atxmega's may work with a future version.

8.3. Support

Amforth is not a commercial software. I hesitate to call it a product. Since you get all the source code for the system, you should be able to solve all problems yourself. On the other side I'm more than interested in any use of amforth and want to know what you're doing with it. If you find anything strange or faulty don't hesitate to mail it to the mailing list (<mailto:amforth-devel@lists.sourceforge.net>) .

8.4. Contributors

amforth would not be the system it now is without the feedback and help from its users. I would like to thank all of them. The following people made an outstanding work to improve amforth (in no particular order): Milan Horkel, Ullrich Hoffmann, Michael Kalus, Karl Lunt, Bruce Wolk, Lubos Pekny, Erich Wälde. But there are many more that helped by simply asking how to do some tasks.