

Manual for Tali Forth 2 for the 65c02

Scot W. Stevenson

Version BETA, 03. Nov 2018

Table of Contents

Dedication	2
Introduction	3
But why?	4
The Big Picture	4
Forth	5
Writing Your Own Forth	5
Overview of Tali Forth	7
Design Considerations	7
Deeper down the rabbit hole	8
User Guide	9
Installing Tali Forth	10
Downloading	10
Running	10
Installing on Your Own Hardware	10
Hardware Projects with Tali Forth 2	11
Running Tali Forth	12
Booting	12
Command-Line History	12
Words	12
Native Compiling	15
Underflow Detection	16
Restarting	16
Gotchas	16
Major Components	18
Blocks	18
The Block Editor	18
The Line-Based Editor ed	21
Assembler	24
Disassembler	24
Developer Guide	26
How Tali Forth Works	27
The Data Stack	27
Dictionary	29
Input	33
The Words create and does>	34
Control Flow	37
Native Compiling	39
cmove, cmove> and move	44

Developing	45
Adding New Words	45
Deeper Changes	45
Code Cheat Sheets	47
Future and Long-Term plans	49
Tutorials	50
Working with Blocks	51
The ed Line-Based Editor	52
First steps with ed	52
Saving Your Text	55
Getting Out of ed	55
Programming with ed	56
Further Information	56
Appendix	58
Reporting Problems	59
FAQ	60
Testing Tali Forth 2	62
User Tests	62
Cycle Tests	62
Old Tests	63
Thanks	64
References and Further Reading	65
Colophon	66

Tali Forth 2 is a bare-metal ANS(ish) Forth for the 65c02 8-bit MPU. It aims to be, roughly in order of importance, easy to try out (just run the included binary), simple (subroutine threading model), specific (for the 65c02 only), and standardized (ANS Forth).

Dedication

For the crew at 6502.org, who made this possible in more ways than one.

Introduction

But why?

Forth is well suited to resource-constrained situations. It doesn't need lots of memory and doesn't have much overhead. [\[CHM1\]](#)

— Charles H. Moore, redgate Hub 2009

The Big Picture

This section provides background information on Forth, the 6502 processor, and why anybody would want to combine the two. It can be skipped if you already know all those things.

The 6502 CPU

It is a well-established fact that humanity reached the apex of processor design with the 6502 in 1976.

Created by a team including Chuck Peddle and Bill Mensch, it was the engine that powered the 8-bit home computer revolution of the 1980s. [1: Rumor has it that there was another MPU called "Z80", but it ended up being a mere footnote.] The VIC-20, Commodore PET, Apple II, and Atari 800 all used the 6502, among others.




Figure 1. The 65c02 MPU. Photo: Anthony King, released in the public domain

More than 40 years later, the processor is still in production by the [Western Design Center](#). Apart from commercial uses, there is an active hobbyist scene centered on the website [6502.org](#). A number of people have built their own 8-bit computers based on this chip and the instructions there, including a [primer](#) by Garth Wilson. It is for these systems that Tali Forth 2 was created.

The most important variant of the 6502 produced today is the [65c02](#), a CMOS chip with some additional instructions. It is for this chip that Tali Forth 2 was written.

But why program in 8-bit assembler at all? The 65c02 is fun to work with because of its clean

instruction set architecture (ISA) This is not the place to explain the joys of assembler. The official handbook for the 65c02 is *Programming the 65816* [EnL].



Garth Wilson answers this question in greater detail as part of his 6502 primer at <http://wilsonminesco.com/6502primer/>.

Forth

If C gives you enough rope to hang yourself, Forth is a flamethrower crawling with cobras. [EW]

— Elliot Williams, *Forth: The Hacker's language*

Forth (Forth is the *enfant terrible* of programming languages. It was invented by Charles "Chuck" H. Moore in the 1960s to do work with radio astronomy, way before there were modern operating systems or programming languages.



A brief history of Forth can be found at <https://www.forth.com/resources/forth-programming-language>

As a language for people who actually need to get things done, it lets you run with scissors, play with fire, and cut corners until you've turned a square into a circle. Forth is not for the faint-hearted: It is trivial, for instance, to redefine `1` as `2` and `true` as `false`. Though you can do really, really clever things with few lines of code, the result can be hard for other people to understand, leading to the reputation of Forth being a "write-only language". However, Forth excels when you positively, absolutely have to get something done with hardware that is really too weak for the job.

It should be no surprise that NASA is one of the organizations that uses Forth. The *Cassini* mission to Saturn used a [Forth CPU](#), for instance. It is also perfect for small computers like the 8-bit 65c02. After a small boom in the 1980s, more powerful computers led to a decline of the language. The "Internet of Things" (IOT) with embedded small processors has led to a certain amount of [renewed interest](#) in the language. It helps that Forth is easy to implement: It is stack-based, uses Reverse Polish Notation (RPN) and a simple threaded interpreter model.

There is no way this document can provide an adequate introduction to Forth. There are quite a number of tutorials, however, such as *A Beginner's Guide to Forth* by J.V. Nobel [JVN] or the classic (but slightly dated) *Starting Forth* by Leo Brodie. [LB1] Gforth, one of the more powerful free Forths, comes with its own [tutorial](#).



Once you have understood the basics of the language, do yourself a favor and read *Thinking Forth* by Brodie [LB2] which deals with the philosophy of the language. Even if you never code a line of Forth in your life, exposure to Forth will change the way you think about programming, much like Lisp.

Writing Your Own Forth

Even if the 65c02 is great and Forth is brilliant, why got to the effort of writing a new, bare-metal

version of the languages? After almost 50 years, shouldn't there be a bunch of Forths around already?

FIG Forth

In fact, the classic Forth available for the whole group of 8-bit MPUs is FIG Forth. "FIG" stands for "Forth Interest Group". Ported to various architectures, it was original based on an incarnation for the 6502 written by Bill Ragsdale and Robert Selzer. There are PDFs of the [6502 version](#) from September 1980 freely available — there is a tradition of placing Forth in the public domain — and more than one hobbyist has revised it to his machine.

However, Forth has changed a lot in the past three decades. There is now a standardized version called [ANS Forth](#), which includes very basic changes such as how the `do` loop works. Learning the language with FIG Forth is like learning English with *The Canterbury Tales*.

A Modern Forth for the 65c02

Tali Forth was created to provide an easy to understand modern Forth written especially for the 65c02 that anybody can understand, adapt to their own use, and maybe actually work with. As part of that effort, the source code is heavily commented. And this document tries to explain the internals in more detail.

Overview of Tali Forth

Design Considerations

When creating a new Forth, there are a bunch of design decisions to be made.



Probably the best introduction to these questions is found in "Design Decisions in the Forth Kernel" at <http://www.bradrodriguez.com/papers/moving1.htm> by Brad Rodriguez.

Spoiler alert: Tali Forth is a subroutine-threaded (STC) variant with a 16-bit cell size and a dictionary that keeps headers and code separate. If you don't care and just want to use the program, skip ahead.

Characteristics of the 65c02

Since this is a bare-metal Forth, the most important consideration is the target processor. The 65c02 only has one full register, the accumulator A, as well as two secondary registers X and Y. All are 8-bit wide. There are 256 bytes that are more easily addressable on the Zero Page. A single hardware stack is used for subroutine jumps. The address bus is 16 bits wide for a maximum of 64 KiB of RAM and ROM.

For the default setup, we assume 32 KiB of each, but allow this to be changed so people can adapt Tali to their own hardware.

Cell Size

The 16-bit address bus suggests the cell size should be 16 bits as well. This is still easy enough to realize on a 8-bit MPU.

Threading Technique

A "thread" in Forth is simply a list of addresses of words to be executed. There are four basic threading techniques: [\[GK\]](#)

Indirect threading (ITC)

The oldest, original variant, used by FIG Forth. All other versions are modifications of this model.

Direct threading (DTC)

Includes more assembler code to speed things up, but slightly larger than ITC.

Token threading (TTC)

The reverse of DTC in that it is slower, but uses less space than the other Forths. Words are created as a table of tokens.

Subroutine threading (STC)

Converts the words to a simple series of `jsr` combinations. Easy to understand and less complex than the other variants, but uses more space and is slower.

Our lack of registers and the goal of creating a simple and easy to understand Forth makes subroutine threading the most attractive solution. We try to mitigate the pain caused by the 12 cycle cost of each and every `jsr-rts` combination by including a relatively high number of native words.

Register Use

The lack of registers — and any registers larger than 8 bit at that — becomes apparent when you realize that Forth classically uses at least four virtual registers:

Table 1. The classic Forth registers

Register	Name
W	Working Register
IP	Interpreter Pointer
DSP	Data Stack Pointer
RSP	Return Stack Pointer

On a modern processor like a RISC-V RV32I with 32 registers of 32 bit each, none of this would be a problem (in fact, we'd probably run out of ways to use the registers). On the 65c02, at least we get the RSP for free with the built-in stack pointer. This still leaves three registers. We cut that number down by one through subroutine threading, which gets rid of the IP. For the DSP, we use the 65c02's Zero Page indirect addressing mode with the X register. This leaves W, which we put on the Zero Page as well.

Data Stack Design

We'll go into greater detail on how the Data Stack works in a later chapter when we look at the internals. Briefly, the stack is realized on the Zero Page for speed. For stability, we provide underflow checks in the relevant words, but give the user the option of stripping it out for native compilation. There are no checks for overflow because those cases tend to be rare.

Dictionary Structure

Each Forth word consists of the actual code and the header that holds the meta-data. The headers are arranged as a simple single-linked list.

In contrast to Tali Forth 1, which kept the header and body of the words together, Tali Forth 2 keeps them separate. This lets us play various tricks with the code to make it more effective.

Deeper down the rabbit hole

This concludes our overview of the basic Tali Forth 2 structure. For those interested, a later chapter will provide far more detail.

User Guide

Installing Tali Forth

Downloading

Tali Forth 2 lives on GitHub at <https://github.com/scotws/TaliForth2>. This is where you will always find the current version. You can either clone the code with git or simply download it. To just test Tali Forth, all you need is the binary file `taliforth-py65mon.bin`.

Running

Downloading the py65mon Simulator

Tali was written to run out of the box on the py65mon simulator from <https://github.com/mnaberez/py65>. This is a Python program that should run on various operating systems. Py65mon is also required for the test suite.

To install py65mon on Linux, use the command

```
sudo pip install -U py65
```

If you don't have `pip` installed, you will have to add it first with something like `sudo apt-get install python-pip` (Ubuntu Linux). There is a `setup.py` script as part of the package.

Running the Binary

To start the emulator, run:

```
py65mon -m 65c02 -r taliforth-py65mon.bin
```

Note that the option `-m 65c02` is required, because Tali Forth makes extensive use of the additional commands of the CMOS version and will not run on a stock 6502 MPU.

Installing on Your Own Hardware

The Tali Forth project started out as a way to run Forth on my own 65c02 computer, the Übersquirrel. Though it soon developed a life of its own, a central aim of the project is to provide a working, modern Forth that people can install on their projects.



Figure 2. The functioning Übersquirrel Mark Zero prototype, August 2013. Photo by Scot W. Stevenson

The Platform Files

For this to work, you need to go to the `platform` folder and create your own kernel code to replace `platform-py65mon.asm`, the default kernel for use with the py65mon kernel. By convention, the name should start with `platform-`. See the `README.md` file in the the `platform` folder for details.

Hardware Projects with Tali Forth 2

This is a list of projects known to run Tali Forth 2. Please let me know if you want to have your project added to the list.

- **Steckschwein** (<https://steckschwein.de/>) by Thomas Woinke and Marko Lauke. A multi-board 8 MHz 65c02 system. Platform file: `platform-steckschwein.asm` (26. Oct 2018)
- **SamCo's SBC** (<https://github.com/SamCoVT/SBC>) by Sam Colwell. A single-board computer running at 4MHz. Platform file: `platform-sbc.asm` (29. Oct 2018)

Running Tali Forth

One doesn't write programs in Forth. Forth is the program.

— Charles Moore, Masterminds of Programming

Booting

Out of the box, Tali Forth boots a minimal kernel to connect to the `py65mon` simulator. By default, this stage ends with a line such as

```
Tali Forth 2 default kernel for py65mon (18. Feb 2018)
```

When you port Tali Forth to your own hardware, you'll have to include your own kernel (and probably should print out a different line).

Tali Forth itself boots next, and after setting up various internal things, compiles the high level words. This causes a slight delay, depending on the number and length of these words. As the last step, Forth should spit out a boot string like

```
Tali Forth 2 for the 65c02  
Version ALPHA 02. July 2018  
Copyright 2014-2018 Scot W. Stevenson  
Tali Forth 2 comes with absolutely NO WARRANTY  
Type 'bye' to exit
```

Because these are the last high-level commands Tali Forth executes, this functions as a primitive self-test. If you have modified the high level Forth words in either `forth_words.fs` or `user_words.fs`, the boot process might fail with a variant of the error message "unknown word". The built-in, native words should always work. For this `dump` is a built-in word — it is very useful for testing.

Command-Line History

Tali's command line includes a simple, eight-element history function. To access the previous entries, press `CONTROL-p`, to go forward to the next entry, press `CONTROL-n`.

Words

Tali Forth comes with the following Forth words out of the box:

```

see block_init_ramdrive evaluate thru load flush empty-buffers buffer update
block save-buffers block_words_deferred blockwrite blockread scr blk buffstatus
buffblocknum blkbuffer 2literal 2constant d.r d. ud.r ud. .r u.r action-of is defer@
defer! endcase endof of case while until repeat else then if .( ( drop dup swap ! @
over >r r> r@ nip rot -rot tuck , c@ c! +! execute emit type . u. ? false true
space 0 1 2 2dup ?dup + - abs dabs and or xor rshift lshift pick char [char] char+
chars cells cell+ here 1- 1+ 2* 2/ = <> < u< u> > 0= 0<> 0> 0< min max 2drop 2swap
2over 2! 2@ 2variable 2r@ 2r> 2>r invert negate dnegate c, bounds spaces bl
-trailing /string refill accept unused depth key allot create does> variable constant
value to s>d d>s d- d+ erase blank fill find-name ' ['] name>int int>name
name>string >body defer latestxt latestnt parse-name parse source source-id : ;
:noname
compile, [ ] 0branch branch literal sliteral ." s" postpone immediate compile-only
never-native always-native nc-limit uf-strip abort abort" do ?do i j loop +loop exit
unloop leave recurse quit begin again state evaluate base digit? number >number hex
decimal count m* um* * um/mod sm/rem fm/mod / /mod mod */mod */ \ move cmove> cmove
pad within >in <# # #s #> hold sign output input cr page at-xy marker words
wordsize aligned align bell dump .s disasm compare search environment? find word cold
bye

```



This list might be outdated. To get the current list, run `words` from inside Tali Forth.

Though the list might look unsorted, it actually reflects the priority in the dictionary, that is, which words are found first. For instance, the native words — those coded in assembler — always start with `drop` and end with `bye`. This is the last word that Tali will find in the dictionary. [2: If you're going to quit anyway, speed can't be that important] The words before `drop` are those that are defined in high-level Forth. For more information on individual the words, use the `see` command.

Note that the built-in words are lower case. While Tali is not case sensitive — `KASUMI` is the same word as `Kasumi` — newly defined words will be lowercased as they are created and entered into the dictionary. There is a slight speed advantage during lookup to using lowercase words (because Tali doesn't have to lowercase the entered text), so all of the tests for Tali are in lowercase.

The ANS Standard

Tali Forth is orientated on ANS Forth, the standard defined by the American National Standards Institute (ANSI). Tali also adopted some words from Gforth such as `bounds`. In practical terms, Tali aims to be a subset of Gforth: If a program runs on Tali, it should run on Gforth the same way or have a very good reason not to.

Tali-Specific Words

In addition, there are words that are specific to Tali Forth.

0branch

(`f—`) Branch if zero. Used internally for branching commands such as `if`. In modern Forths, this is usually replaced by `cs-pick` and `cs-roll`; Tali Forth might switch to these words in the future.

0

(`—0`) Push the number 0 on the Data Stack.

1

(`—0`) Push the number 1 on the Data Stack.

2

(`—0`) Push the number 2 on the Data Stack.

always-native

Mark last word in dictionary so that it is always natively compiled.

bell

Ring the terminal bell (ASCII 07).

block-read

(`addr blk#—`) This is a deferred word the user can change to point to their own routine for reading 1K blocks into memory from storage.

block-write

(`addr blk#—`) This is a deferred word the user can change to point to their own routine for writing 1K blocks from memory to storage.

block-ramdrive-init

(`u—`) Create a ram drive with the given number of blocks (numbered 0 to (u-1)) to allow use of the block words with no additional hardware. Because the blocks are only held in ram, they will be lost when the hardware is powered down or the simulator is stopped.

branch

Always take branch. See `0branch`.

compile-only

Mark last word in dictionary as compile-only.

digit?

(`char—u f | char f`) If character is a digit, convert and set flag to `true`, otherwise return the offending character and a `false` flag.

ed

Start the command-line editor. There is a whole chapter on this father down.

input

Return the address where the vector for the input routine is stored (not the vector itself). Used for input redirection for `emit` and others.

int>name

(`xt—nt`) Given the execution execution token (xt), return the name token (nt).

latestnt

(`—nt`) Return the last used name token. The Gforth version of this word is called `latest`.

nc-limit

(`—addr`) Return the address where the threshold value for native compiling native compiling is kept. To check the value of this parameter, use `nc-limit ?`. The default value is 20.

never-native

Mark most recent word so it is never natively compiled.

number

(`addr u—u | d`) Convert a string to a number. Gforth uses `s>number?` and returns a success flag as well.

output

(`—addr`) Return the address where the vector for the output routine is stored (not the vector itself). Used for output redirection for `emit` and others.

uf-strip

(`—addr`) Return the address where the flag is kept that decides if the underflow checks are removed during native compiling. To check the value of this flag, use `uf-strip ?`.

wordsize

(`nt—u`) Given the name token (nt) of a Forth word, return its size in bytes. Used to help tune native compiling.

Native Compiling

As the name says, subroutine threaded code encodes the words as a series of subroutine jumps. Because of the overhead caused by these jumps, this can make the code slow. Therefore, Tali Forth enables native compiling, where the machine code from the word itself is included instead of a subroutine jump. This is also called "inlining".

The parameter `nc-limit` sets the limit of how small words have to be to be natively compiled. To get the current value (usually 20), check the value of the system variable:

```
nc-limit ?
```

To set a new limit, save the maximal allowed number of bytes in the machine code like any other Forth variable:

```
40 nc-limit !
```

To complete turn off native compiling, set this value to zero.

Underflow Detection

When a word tries to access more words on the stack than it is holding, an "underflow" error occurs. Whereas Tali Forth 1 didn't check for these errors, this version does.

However, this slows the program down. Because of this, the user can turn off underflow detection for words that are natively compiled into new words. To do this, set the system variable `uf-strip` to `true`. Note this does not turn off underflow detection in the built-in words. Also, words with underflow detection that are not included in new words through native compiling will also retain their tests.

Restarting

Tali Forth has a non-standard word `cold` that resets the system. This doesn't erase any data in memory, but just moves the pointers back. When in doubt, you might be better off quitting and restarting completely.

Gotchas

Some things to look out for when using Tali Forth.

Cell Size

Tali has a 16-bit cell size.



Use `1 cells 8 *` to get the cell size in bits with any Forth.

This can trip up calculations when compared to the *de facto* standard Gforth with 64 bits. Take this example:

```
( Gforth )      decimal 1000 100 um* hex swap u. u. ( returns 186a0 0 ok )
( Tali Forth )  decimal 1000 100 um* hex swap u. u. ( returns 86a0 1 ok )
```

Tali has to use the upper cell of a double-celled number to correctly report the result, while Gforth doesn't. If the conversion from double to single is only via a `drop` instruction, this will produce different results.

There is a similar effect with the Gforth word `bounds`: Because of Tali's 16 bit address space, it wraps the upper address if we go beyond `$FFFF`:

```
( Gforth )      hex FFFF 2 bounds swap u. u. ( returns 10001 ffff ok )
( Tali )        hex FFFF 2 bounds swap u. u. ( returns      1 ffff ok )
```

Delimiters During Parsing

Both `parse-name` and `parse` skip white space - defined as ASCII characters from 00 to 32 (SPACE) inclusive - when the standard talks about "spaces". Otherwise, Tali would choke on TABs during compiling, and the `ed` editor couldn't be used to edit programs because of the Line Feed characters. This is covered in the standard, see the footnote at <https://forth-standard.org/standard/core/PARSE-NAME> by Anton Ertl, referencing <http://forth-standard.org/standard/usage#subsubsection.3.4.1.1> and <http://forth-standard.org/standard/file#subsection.11.3.5> .

Negative `allot`

The ANSI standard does not define what happens if there is an attempt to free more memory with `allot` by passing a negative value than is available. Tali will let the user free memory up the beginning of RAM assigned to the Dictionary (marked with `cp0` in the code), even though this can mean that the Dictionary itself is compromised. This is Forth, you're the boss.

However, any attempt to free more memory than that will set the beginning of RAM to `cp0`. Also, the Dictionary Pointer `dp` will point to the *last native word* of the Dictionary, which is usually `drop`. Because of this, the high level words defined during boot will *not* be available. There will be an error message to document this. Realistically, you'll probably want to restart with `cold` if any of this happens.

Major Components

Blocks

Tali supports the optional BLOCK word set. The 2012 Forth standard defines a block as 1024 bytes, and the buffers for them are the same size (as opposed to some older forths that had smaller buffers.) Tali currently comes with one buffer.

Before these words can be used, the user needs to write two routines: one for reading blocks into RAM and one for writing blocks out from RAM. Both of these should have the signature (`addr blk#—`). Once these have been written, they can be incorporated into the BLOCK word set by changing the deferred words `block-read` and `block-write`. That might look like:

```
' myblockreader IS BLOCK-READ
' myblockwriter IS BLOCK-WRITE
```

Once these two deferred words have been updated, you can use the block words.

If you would like to play with some blocks, but don't have any hardware or are running Tali in a simulator, fear not! Tali has a built-in RAM drive that can be accessed by running:

```
4 block-ramdrive-init
```

This reserves a chunk of ram with four blocks in it (numbered 0-3) which is enough to play around with. It also sets up the routines for reading and writing blocks in this ramdrive for you. If you want more blocks, you can change the number. Because they start at zero, the last valid block will always be one less than the number you provide.

Be careful about creating too many blocks as they are 1K each. It's also worth noting that running `block-ramdrive-init` again will create another ramdrive and the existing one will be inaccessible while still taking up space in RAM.

The Block Editor

If you are using blocks (see the block chapter), you can use the following code to create a very basic screen editor that allows you to replace a single line or an entire screen. Screens are 16 lines (numbered 0-15) of 64 characters each, for a total of 1K characters. Because newlines are not stored in the blocks (the remainder of each line is filled with spaces,) you should leave a space in the very last character of each line to separate the words in that line from the words in the next line.



This editor uses a word named `E` which may interfere with the use of hex. Once you have entered the editor words, you will need to use the hex value "0E" anywhere you want a single "E" value. This will not interfere with hex numbers that have additional digits. Alternately, you can rename the `E` word and update the `0` word which uses it.

```

( Simple Editor for screens /blocks )
decimal
( line provides the address, in the buffer, of the given line )
: line ( line# - c-addr)
  64 *      ( Convert line number to # characters offset )
  scr @ block ( Get the buffer address for that block )
  + ;      ( Add the offset )

: E ( line# - ) ( Erase the given line number with spaces )
  line 64 blank update ;

: 0 ( line# - ) ( Overwrite line with new text )
  dup E      ( Erase existing text on line )
  cr dup 2 u.r ." * " line 64 accept drop update ;

( Editor, continued )
: enter-screen ( scr# - )
  dup scr ! buffer drop
  16 0 do i o loop ;
: erase-screen ( scr# - )
  dup scr ! buffer 1024 blank update ;

```

To use this editor, first select a screen to work with by running list on it. If you are planning on using **load** to run some code later, it's worth noting that only screens above 0 can be LOADED. Screen 0 is reserved for comments describing what is on the other screens. It can be LISTed and edited, but cannot be LOADED.

```
1 list
```

Tali will show you the current (blank) contents of that screen.

```
Screen # 1
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
ok
```

To add some text to line 3, you might say

```
3 o
```

This will give you a prompt to enter the text to overwrite line 3. You can enter up to 64 characters. Once you have selected a screen with `list`, you can use just `L` to list it again.

To replace the contents of an entire screen, you can say something like:

```
2 enter-screen
```

This will prompt you, line by line, for the new contents to screen 2.

Once you have your screens the way you want them, you can type

```
flush
```

to flush your changes out to storage.

You can enter Forth code on these screens. At the moment, Tali only supports comments in parentheses inside of blocks, so make sure you put your comments (like this) rather than using `\` when entering Forth code. To load the Forth code on a screen, just type something like:

```
2 load
```

Because a screen only holds 16 lines, you may need to split your code across multiple screens. You

can load a series of screens (in order) using the `thru` command like so:

```
1 3 thru
```

The Line-Based Editor `ed`



This manual includes [a tutorial](#) for `ed`

`Ed` makes no response to most commands – there is no prompting or typing of messages like "ready". (This silence is preferred by experienced users, but sometimes a hangup for beginners.) [\[BWK\]](#)

— B. W. Kernighan, A Tutorial Introduction to the UNIX Text Editor

Tali Forth 2 currently ships with a clone of the `ed` line-based editor of Unix fame. It is invoked with `ed` and does not change the data stack. The formal name is `ed6502`.



`ed` is included in a very primitive form only and should be considered ALPHA.

Supported Commands

`ed` currently supports only a small number of the commands of the Unix version:

a

Add new lines below given line

d

Delete line

i

Add new lines above given line

q

Quit if no unsaved work

Q

Unconditional quit, unsaved work is lost

w

Write text to given memory location (eg `7000w`)

=

Print value of given parameter (eg `$=` gives number of last line)

For the parameters, these are currently available:

- .
Current line number
- ,
When alone: All lines, the same as 1, \$ or %
- \$
Last line
- %
All lines, the same as 1, \$ or , alone

An empty line (pressing the ENTER key) will advance by one line and print it. A simple number will print that line without the line number and make that line the new current line.

Future planned commands

There is no time frame for these additions.

- +
Advance by one line, print it and make it the new current line
- Go back by one line, print it and make it the new current line
- c
Change a line, possibly adding new lines
- e
Edit lines given as addr, u in text buffer
- f
Show current target address for writes
- j
Join two lines to a new line
- m
Move block of text to new line
- r
Read lines into text buffer
- s
Substitute one string on line with another
- !
Execute a shell command (Forth command in our case)

#

Comment, ignore rest of the line

The following parameter features:

;

Range from current line to end, same as `.,$`

Differences to Unix ed

(Apart from missing about 90 percent of the features, that is)

- The `w` (write) command takes its parameter before and not after the word. Where Unix `ed` uses the format `w <FILENAME>`, `ed6502` takes the address to write the text to as `7000w`.



At this point, it isn't totally clear what happens if the number base is set to hexadecimal via `hex`. Use at your own risk.

Using ed for programming

`Ed` can be used to write programs and then execute them with `evaluate`. For instance, a session to add a small string could look something like this:

```
ed
a
.( Shepard, is that ... You're alive?)
.
7000w ①
22 ②
q
```

① Address we save the command to

② Number of characters saved including final line feed

It is a common mistake to forget the `.` (dot) to end the input, and try to go immediately to saving the text. Then, we can run the program:

```
7000 22 evaluate
```

Note that `evaluate` will handle line feeds, carriage returns and other white space apart from simple spaces without problems.

Known Issues

Memory use

`Ed` currently uses memory without releasing it when done. For small, quick edits, this probably is

not a problem. However, if you know you are going to be using more memory, you probably will want to set a marker first.

```
marker pre-edit ①  
ed ②  
pre-edit ③
```

- ① Set marker at current value of [here](#)
- ② Edit normally
- ③ Call marker, releasing memory

This issue might be taken care of in a future release.

Address of Saved Text

Currently, `ed` returns the data stack just the way it found it. This means that you have to remember where you saved the text to with `w` and how long it was. A different option would be to return (`—addr u`), that is, the address and length of the text we saved. If nothing is saved, the program would return a zero length as TOS.

Developer Information

`Ed` will be massively rewritten under the hood once the code has been stabilized, has all features, and a testing suite. Currently, it's somewhat of a mess and contains some testing routines that will be removed in the final version.

The "buffer" of `ed` is a simple linked list of nodes, consisting of a pointer to the next entry, a pointer to the string address, and the length of that string. Each entry is two byte, making six bytes in total for each node. A value of 0000 in the pointer to the next address signals the end of the list. The buffer starts at the point of the `cp` (accessed with the Forth word [here](#)) and is only saved to the given location when the `w` command is given.

Assembler



Currently, there is no assembler included. The plan is to include a simple assembler based on Simpler Assembler Notation (SAN).

Disassembler

Tali Forth is currently shipped with a very primitive disassembler, which is started with `disasm (addr u—)`.

Format

The output format is in Simpler Assembler Notation (SAN) which adds the mode of an instruction to the main mnemonic, simplifying parsing of code. For instance,

```
lda #1
sta $1000,x
sta $80
lda ($80)
```

becomes

```
lda.# 1
sta.x $1000
sta.z $80
lda.di $80
```

A full discussion of SAN is beyond the scope of this document, see [this link](#) for an overview of the format (currently still under a different name).

Replacing the Disassembler

Tali was designed to make it easy for the user to swap in a different disassembler. The code is in a separate file, `disassembler.asm`. To replace this by your version, move the file to a different name, make sure that it accepts (`addr u`) on the Forth Data Stack as parameters, and start your code after the label `disassembler:.`

Developer Guide

How Tali Forth Works

Our intent was to create a pleasant computing environment for ourselves, and our hope was that others liked it. [DMR]

— Dennis M. Ritchie, Reflections on Software Research

The Data Stack

Tali Forth uses the lowest part of the top half of the Zero Page for the Data Stack (DS). This leaves the lower half of the Zero Page for any kernel stuff the user might require. The DS grows towards the initial user variables. See the file `definitions.asm` for details. Because of the danger of underflow, it is recommended that the user kernel's variables are kept closer to `$0100` than to `$007F`.

The X register is used as the Data Stack Pointer (DSP). It points to the least significant byte of the current top element of the stack ("Top of the Stack", TOS).



In the first versions of Tali Forth 1, the DSP pointed to the next *free* element of the stack. The new system makes detecting underflow easier and follows the convention in Liara Forth.

Initially, the DSP points to `$78`, not `$7F` as might be expected. This provides a few bytes as a "floodplain" for underflow. The initial value of the DSP is defined as `dsp0` in the code.

Single Cell Values

Since the cell size is 16 bits, each stack entry consists of two bytes. They are stored little endian (least significant byte first). Therefore, the DSP points to the LSB of the current TOS. [3: Try reading that last sentence to a friend who isn't into computers. Aren't abbreviations fun?]

Because the DSP points to the current top of the stack, the byte it points to after boot — `dsp0` — will never be accessed: The DSP is decremented first with two `dex` instructions, and then the new value is placed on the stack. This means that the initial byte is garbage and can be considered part of the floodplain.

Snapshot of the Data Stack with one entry as TOS. The DSP has been increased by one and the value written.

```

+-----+
|          ... |
+-          -+
|          |   ...
+- (empty)  -+
|          |   FE,X
+-          -+
... |          |   FF,X
+-----+
$0076 |          LSB| 00,X <-- DSP (X Register)
+-    TOS    -+
$0077 |          MSB| 01,X
+-----+
$0078 | (garbage) | 02,X <-- DSP0
+-----+
$0079 |          | 03,X
+ (floodplain) +
$007A |          | 04,X
+-----+

```

Note that the 65c02 system stack — used as the Return Stack (RS) by Tali — pushes the MSB on first and then the LSB (preserving little endian), so the basic structure is the same for both stacks.

Because of this stack design, the second entry ("next on stack", NOS) starts at **02,X** and the third entry ("third on stack", 3OS) at **04,X**.

Underflow Detection

Most native words come with built-in underflow detection. This works by comparing the Data Stack Pointer (X) to values that it must be smaller than (because the stack grows towards 0000). For instance, to make sure we have one element on the stack, we write

```

        cpx #dsp0-1
        bmi _ok
        jmp underflow
_ok:
        (...)
```

For the most common cases, this gives us:

Table 2. DSP values for underflow testing

Test for	Pointer offset
one element	dsp0-1
two elements	dsp0-3

Test for	Pointer offset
three elements	dsp0-5
four elements	dsp0-7

Underflow detection adds seven bytes to the words that have it. However, it increases the stability of the program. There is an option for stripping it out during for user-defined words (see below).

Double Cell Values

The double cell is stored on top of the single cell. Note this places the sign bit at the beginning of the byte below the DSP.

```

+-----+
|           |
+=====+
|           LSB| $0,x  <-- DSP (X Register)
++ Top Cell +-
|S|           MSB| $1,x
+-----+
|           LSB| $2,x
+- Bottom Cell +-
|           MSB| $3,x
+=====+

```

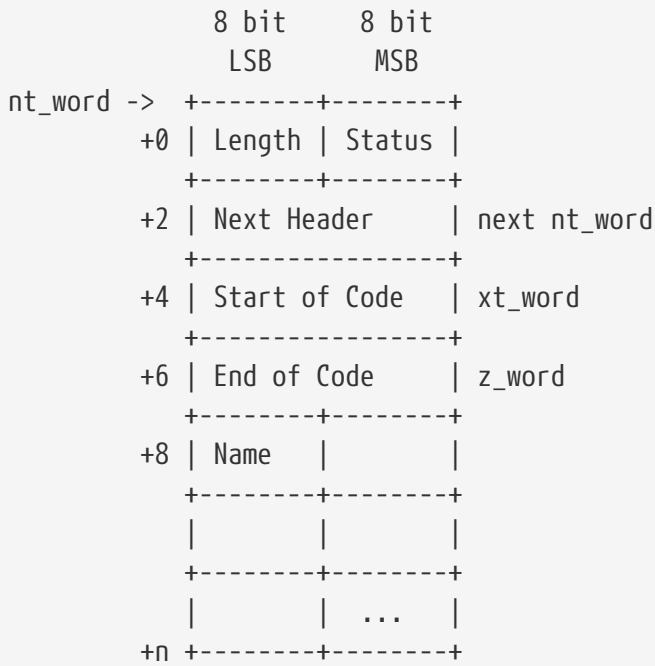
Tali Forth does not check for overflow, which in normal operation is too rare to justify the computing expense.

Dictionary

Tali Forth follows the traditional model of a Forth dictionary — a linked list of words terminated with a zero pointer. The headers and code are kept separate to allow various tricks in the code.

Elements of the Header

Each header is at least eight bytes long:



Each word has a **name token** (nt, nt_word in the code) that points to the first byte of the header. This is the length of the word's name string, which is limited to 255 characters.

The second byte in the header (index 1) is the status byte. It is created by the flags defined in the file **definitions.asm**:

Table 3. Header flags

Flag	Function
CO	Compile Only
IM	Immediate Word
NN	Never Native Compile
AN	Always Native Compile
UF	Underflow detection
HC	Has Code Field Area (CFA)

Note there are currently two bits unused.

CO

This word may only be used inside definitions of new words.

IM

Word is executed immediately during definitions of new words.

NN

Word is never inlined. Usually means that the return address from a subroutine jump is required for processing.

AN

Word must always be inlined.

HC

Consider first three bytes of the word's code the Code Field Area (CFA) of the word. Used by words defined with `create` so `>body` returns the correct value.

The status byte is followed by the **pointer to the next header** in the linked list, which makes it the name token of the next word. A 0000 in this position signals the end of the linked list, which by convention is the word `bye` for the native code words.

This is followed by the current word's **execution token** (`xt`, `xt_word`) that points to the start of the actual code. Some words that have the same functionality point to the same code block.



Because Tali uses a subroutine threaded model (STC), the classic Forth distinction between the Code Field Area (CFA) and the Parameter Field Area (PFA, also Data Field Area) is meaningless — it's all "payload".

The next pointer is for the **end of the code** (`z_word`) to enable native compilation of the word (if allowed and requested).

The **name string** starts at the eighth byte. The string is *not* zero-terminated. Tali Forth lowercases names as they are copied into the dictionary and also lowercases during lookup, so `quarian` is the same word as `QUARIAN`. If the name in the dictionary is directly modified, it is important to ensure that only lowercase letters are used, or else Tali will not be able to find that word.

Structure of the Header List

Tali Forth distinguishes between three different word sources: The **native words** that are hard-coded in the file `native_words.asm`, the **Forth words** from `forth_words.asm` which are defined as high-level words and then generated at run-time when Tali Forth starts up, and **user words** in the file `user_words.asm`.

Tali has an unusually high number of native words in an attempt to make the Forth as fast as possible on the 65c02 and compensate for the disadvantages of the subroutine threading model (STC). The first word on that list — the one that is checked first — is always `drop`, the last one — the one checked for last — is always `bye`. The words which are (or are assumed to be) used more than others come first. Since humans are slow, words that are used more interactively like `words` always come later.

The list of Forth words ends with the intro strings. This functions as a primitive form of a self-test: If you see the welcome message, compilation of the Forth words worked.

```
$0000 +-----+ ram_start, zpage, user0
      | User variables |
      +-----+
      | ^ Data Stack | <-- dsp
      | |           |
```

\$0078	+-----+	dsp0, stack
	(Reserved for	
	kernel)	
\$0100	+=====+	
	^ Return Stack	<-- rsp
\$0200	+-----+	rsp0, buffer, buffer0
	v Input Buffer	
\$0300	+-----+	cp0
	v Dictionary	
	(RAM)	
	~~~~~	<-- cp
	+-----+	
	ACCEPT history	
\$7FFF	#####	ram_end
\$8000		forth, code0
	Tali Forth	
	(24 KiB)	
\$E000	+-----+	kernel_putc, kernel_getc
	Kernel	
\$F000	+-----+	
	I/O addresses	
	+-----+	
	Kernel	
\$FFFA	+-----+	
	65c02 vectors	
\$FFFF	+-----+	

Note that some of these values are hard-coded into the test suite; see the file `definitions.txt` for details.

# Input

Tali Forth follows the ANS Forth input model with `refill` instead of older forms. There are four possible input sources:

- The keyboard ("user input device")
- A character string in memory
- A block file
- A text file

To check which one is being used, we first call `blk` which gives us the number of a mass storage block being used, or 0 for the "user input device" (keyboard). In the second case, we use `source-id` to find out where input is coming from:

Table 4. Non-block input sources

Value	Source
0	keyboard
-1	string in memory
<i>n</i>	file-id

Since Tali currently doesn't support blocks, we can skip the `blk` instruction and go right to `source-id`.

## Booting

The initial commands after reboot flow into each other: `cold` to `abort` to `quit`. This is the same as with pre-ANS Forths. However, `quit` now calls `refill` to get the input. `refill` does different things based on which of the four input sources (see above) is active:

### Keyboard entry

This is the default. Get line of input via `accept` and return `true` even if the input string was empty.

### evaluate string

Return a `false` flag

### Input from a buffer

*Not implemented at this time*

### Input from a file

*Not implemented at this time*

## The Command Line Interface (CLI)

Tali Forth accepts input lines of up to 256 characters. The address of the current input buffer is stored in `cib`. The length of the current buffer is stored in `ciblen`—this is the address that `>in`

returns. `source` by default returns `cib` and `ciblen` as the address and length of the input buffer.

## The Word `evaluate`

The word `evaluate` is used to execute commands that are in a string. A simple example:

```
s" 1 2 + ." evaluate
```

Tali Forth uses `evaluate` to load high-level Forth words from the file `forth_words.asc` and, if present, any extra, user-defined words from `user_words.asc`.

## The Words `create` and `does>`

The tandem of words `create` and `does>` is the most complex, but also most powerful part of Forth. Understanding how it works in Tali Forth is important if you want to be able to modify the code. In this text, we walk through the generation process for a subroutine threaded code (STC) such as Tali Forth.



For a more general explanation, see Brad Rodriguez' series of articles at <http://www.bradrodriguez.com/papers/moving3.htm> There is a discussion of this walkthrough at <http://forum.6502.org/viewtopic.php?f=9&t=3153>

We start with the following standard example, a high-level Forth version of the word `constant`.

```
: constant ( "name" -- ) create , does> @ ;
```

We examine this in three phases or "sequences", following Rodriguez (based on [DB]).

### Sequence 1: Compiling the Word `constant`

`constant` is a defining word, one that makes new words. In pseudocode, ignoring any compilation to native 65c02 assembler, the above compiles to:

```
      jsr CREATE
      jsr COMMA
      jsr (DOES>)      ; from DOES>
a:    jsr DODOES      ; from DOES>
b:    jsr FETCH
      rts
```

To make things easier to explain later, we've added the labels `a` and `b` in the listing.



This example uses the traditional word (`does>`), which in Tali Forth 2 is actually an internal routine that does not appear as a separate word. This version is easier to explain.

`does>` is an immediate word that adds not one, but two subroutine jumps, one to `(does>)` and one to `dodoes`, which is a pre-defined system routine like `dovar`. We'll discuss those later.

In Tali Forth, a number of words such as `defer` are "hand-compiled", that is, instead of using forth such as

```
: defer create ['] abort , does> @ execute ;
```

we write an optimized assembler version ourselves (see the actual `defer` code). In these cases, we need to use `(does>)` and `dodoes` instead of `does>` as well.

## Sequence 2: Executing the Word `constant`

Now when we execute

```
42 constant life
```

This pushes the `rts` of the calling routine — call it "main" — to the 65c02's stack (the Return Stack, as Forth calls it), which now looks like this:

```
(1) rts ; to main routine
```

Without going into detail, the first two subroutine jumps of `constant` give us this word:

```
(Header "LIFE")
jsr DOVAR ; in CFA, from LIFE's CREATE
4200 ; in PFA (little-endian)
```

Next, we `jsr` to `(does>)`. The address that this pushes on the Return Stack is the instruction of `constant` we had labeled `a`.

```
(2) rts to CONSTANT ("a")
(1) rts to main routine
```

Now the tricks start. `(does>)` takes this address off the stack and uses it to replace the `dovar jsr` target in the CFA of our freshly created `life` word. We now have this:

```
(Header "LIFE")
jsr a ; in CFA, modified by (DOES>)
c: 4200 ; in PFA (little-endian)
```

Note we added a label `c`. Now, when `(does>)` reaches its own `rts`, it finds the `rts` to the main routine on its stack. This is Good Thing™, because it aborts the execution of the rest of `constant`, and we

don't want to do `dodoes` or `fetch` now. We're back at the main routine.

### Sequence 3: Executing `life`

Now we execute the word `life` from our "main" program. In a STC Forth such as Tali Forth, this executes a subroutine jump.

```
jsr LIFE
```

The first thing this call does is push the return address to the main routine on the 65c02's stack:

```
(1) rts to main
```

The CFA of `life` executes a subroutine jump to label `a` in `constant`. This pushes the `rts` of `life` on the 65c02's stack:

```
(2) rts to LIFE ("c")  
(1) rts to main
```

This `jsr` to `a` lands us at the subroutine jump to `dodoes`, so the return address to `constant` gets pushed on the stack as well. We had given this instruction the label `b`. After all of this, we have three addresses on the 65c02's stack:

```
(3) RTS to CONSTANT ("b")  
(2) RTS to LIFE ("c")  
(1) RTS to main
```

`dodoes` pops address `b` off the 65c02's stack and puts it in a nice safe place on Zero Page, which we'll call `z`. More on that in a moment. First, `dodoes` pops the `rts` to `life`. This is `c`, the address of the PFA or `life`, where we stored the payload of this constant. Basically, `dodoes` performs a `dovar` here, and pushes `c` on the Data Stack. Now all we have left on the 65c02's stack is the `rts` to the main routine.

```
[1] RTS to main
```

This is where `z` comes in, the location in Zero Page where we stored address `b` of `constant`. Remember, this is where the PFA of `constant` begins, the `fetch` command we had originally codes after `does>` in the very first definition. The really clever part: We perform an indirect `jmp` — not a `jsr!` — to this address.

```
jmp (z)
```

Now the little payload program of `constant` is executed, the subroutine jump to `fetch`. Since we just put the PFA (`c`) on the Data Stack, `fetch` replaces this by 42, which is what we were aiming for all

along. And since `constant` ends with a `rts`, we pull the last remaining address off the 65c02's stack, which is the return address to the main routine where we started. And that's all.

Put together, this is what we have to code:

`does>`

Compiles a subroutine jump to (`does>`), then compiles a subroutine jump to `dodoes`.

(`does>`)

Pops the stack (address of subroutine jump to `dodoes` in `constant`, increase this by one, replace the original `dovar` jump target in `life`).

`dodoes`

Pop stack (PFA of `constant`), increase address by one, store on Zero Page; pop stack (PFA of `life`), increase by one, store on Data Stack; `jmp` to address we stored in Zero Page.

Remember we have to increase the addresses by one because of the way `jsr` stores the return address for `rts` on the stack on the 65c02: It points to the third byte of the `jsr` instruction itself, not the actual return address. This can be annoying, because it requires a sequence like:

```
inc z
bne +
inc z+1
*
(...)
```

Note that with most words in Tali Forth, as any STC Forth, the distinction between PFA and CFA is meaningless or at least blurred, because we go native anyway. It is only with words generated by `create` and `does>` where this really makes sense.

## Control Flow

### Branches

For `if` and `then`, we need to compile something called a "conditional forward branch", traditionally called `0branch`.



Many Forths now use the words `cs-pick` and `cs-roll` instead of the `branch` variants, see <http://lars.nocrew.org/forth2012/rationale.html#rat:tools:CS-PICK> Tali Forth will probably switch to this variant in the future.

At run-time, if the value on the Data Stack is false (flag is zero), the branch is taken ("branch on zero", therefore the name). Except that we don't have the target of that branch yet — it will later be added by `then`. For this to work, we remember the address after the `0branch` instruction during the compilation of `if`. This is put on the Data Stack, so that `then` knows where to compile its address in the second step. Until then, a dummy value is compiled after `0branch` to reserve the space we need.





This section and the next one are based on a discussion at <http://forum.6502.org/viewtopic.php?f=9&t=3176> see there for more details. Another take on this subject that handles things a bit differently is at <http://blogs.msdn.com/b/ashleyf/archive/2011/02/06/loopy-do-i-loop.aspx>

In Forth, this can be realized by

```
: if postpone 0branch here 0 , ; immediate
```

and

```
: then here swap ! ; immediate
```

Note **then** doesn't actually compile anything at the location in memory where it is at. It's job is simply to help **if** out of the mess it has created. If we have an **else**, we have to add an unconditional **branch** and manipulate the address that **if** left on the Data Stack. The Forth for this is:

```
: else postpone branch here 0 , here rot ! ; immediate
```

Note that **then** has no idea what has just happened, and just like before compiles its address where the value on the top of the Data Stack told it to — except that this value now comes from **else**, not **if**.

## Loops

Loops are more complicated, because we have **do**, **?do**, **loop**, **+loop**, **unloop**, and **leave** to take care of. These can call up to three addresses: One for the normal looping action (**loop** and **+loop**), one to skip over the loop at the beginning (**?do**) and one to skip out of the loop (**leave**).

Based on a suggestion by Garth Wilson, we begin each loop in run-time by saving the address after the whole loop construct to the Return Stack. That way, **leave** and **?do** know where to jump to when called, and we don't interfere with any **if-then** structures. On top of that address, we place the limit and start values for the loop.

The key to staying sane while designing these constructs is to first make a list of what we want to happen at compile time and what at run time. Let's start with a simple **do-loop**.

### **do** at compile-time:

- Remember current address (in other words, **here**) on the Return Stack (!) so we can later compile the code for the post-loop address to the Return Stack
- Compile some dummy values to reserve the space for said code
- Compile the run-time code; we'll call that fragment (**do**)
- Push the current address (the new **here**) to the Data Stack so **loop** knows where the loop contents begin

### do at run-time:

- Take limit and start off Data Stack and push them to the Return Stack

Since `loop` is just a special case of `+loop` with an index of one, we can get away with considering them at the same time.

### loop at compile time:

- Compile the run-time part (`+loop`)
- Consume the address that is on top of the Data Stack as the jump target for normal looping and compile it
- Compile `unloop` for when we're done with the loop, getting rid of the limit/start and post-loop addresses on the Return Stack
- Get the address on the top of the Return Stack which points to the dummy code compiled by `do`
- At that address, compile the code that pushes the address after the list construct to the Return Stack at run-time

### loop at run-time (which is `+loop`)

- Add loop step to count
- Loop again if we haven't crossed the limit, otherwise continue after loop

At one glance, we can see that the complicated stuff happens at compile-time. This is good, because we only have to do that once for each loop.

In Tali Forth, these routines are coded in assembler. With this setup, `unloop` becomes simple (six `pla` instructions — four for the limit/count of `do`, two for the address pushed to the stack just before it) and `leave` even simpler (four `pla` instructions for the address).

## Native Compiling

In a pure subroutine threaded code, higher-level words are merely a series of subroutine jumps. For instance, the Forth word `[char]`, formally defined in high-level Forth as

```
: [char] char postpone literal ; immediate
```

in assembler is simply

```
jsr xt_char  
jsr xt_literal
```

as an immediate, compile-only word. There are two problems with this method: First, it is slow, because each `jsr-rts` pair consumes four bytes and 12 cycles as overhead. Second, for smaller words, the jumps use far more bytes than the actual code. Take for instance `drop`, which in its naive form is simply

```
inx
inx
```

for two bytes and four cycles. If we jump to this word as is assumed with pure subroutine threaded Forth, we add four bytes and 12 cycles — double the space and three times the time required by the actual working code.

(In practice, it's even worse, because `drop` checks for underflow. The actual assembler code is

```
      cpx #dsp0-1
      bmi +
      jmp underflow
*
      inx
      inx
```

for eleven bytes. We'll discuss the underflow checks further below.)

To get rid of this problem, Tali Forth supports **native compiling** (also known as inlining). The system variable `nc-limit` sets the threshold up to which a word will be included not as a subroutine jump, but in machine language. Let's start with an example where `nc-limit` is set to zero, that is, all words are compiled as subroutine jumps. Take a simple word such as

```
: aaa 0 drop ;
```

and check the actual code with `see`

```
see aaa
  nt: 7CD  xt: 7D8
  size (decimal): 6

07D8 20 52 99 20 6B 88  ok
```

(The actual addresses might be different, this is from the ALPHA release). Our word `aaa` consists of two subroutine jumps, one to zero and one to `drop`. Now, if we increase the threshold to 20, we get different code, as this console session shows:

```

20 nc-limit ! ok
: bbb 0 drop ; ok
see bbb
  nt: 7DF xt: 7EA
  size (decimal): 17

07EA CA CA 74 00 74 01 E0 77 30 05 A9 0B 4C C7 AC E8
07FA E8 ok

```

Even though the definition of `bbb` is the same as `aaa`, we have totally different code: The number 0001 is pushed to the Data Stack (the first six bytes), then we check for underflow (the next nine bytes), and finally we `drop` by moving X register, the Data Stack Pointer. Our word is definitely longer, but have just saved 12 cycles.

To experiment with various parameters for native compiling, the Forth word `words&sizes` is included in `user_words.fs` (but commented out by default). The Forth is:

```

: words&sizes ( -- )
  latestnt
  begin
    dup
  0<> while
    dup name>string type space
    dup wordsize u. cr
    2 + @
  repeat
  drop ;

```

An alternative is `see` which also displays the length of a word. One way or another, changing `nc-limit` should show differences in the Forth words.

## Return Stack Special Cases

There are a few words that cause problems with subroutine threaded code (STC): Those that access the Return Stack such as `r>`, `>r`, `r@`, `2r>`, and `2>r`. We first have to remove the return address on the top of the stack, only to replace it again before we return to the caller. This mechanism would normally prevent the word from being natively compiled at all, because we'd try to remove a return address that doesn't exist.

This becomes clearer when we examine the code for `>r` (comments removed):

```

xt_r_from:
    pla
    sta tmplos
    ply

    ; --- CUT FOR NATIVE CODING ---

    dex
    dex
    pla
    sta 0,x
    pla
    sta 1,x

    ; --- CUT FOR NATIVE CODING ---

    phy
    lda tmplos
    pha

z_r_from:    rts

```

The first three and last three instructions are purely for housekeeping with subroutine threaded code. To enable this routine to be included as native code, they are removed when native compiling is enabled by the word `compile`. This leaves us with just the six actual instructions in the center of the routine to be compiled into the new word.

## Underflow Stripping

As described above, every underflow check adds seven bytes to the word being coded. Stripping this check by setting the `uf-strip` system variable to `true` simply removes these seven bytes from new natively compiled words.

It is possible, of course, to have lice and fleas at the same time. For instance, this is the code for `>r`:

```

xt_to_r:
        pla
        sta tmplos
        ply

        ; --- CUT HERE FOR NATIVE CODING ---

        cpx #dsp0-1
        bmi +
        jmp underflow
*
        lda 1,x
        pha
        lda 0,x
        pha

        inx
        inx

        ; --- CUT HERE FOR NATIVE CODING ---

        phy
        lda tmplos
        pha

z_to_r:  rts

```

This word has *both* native compile stripping and underflow detection. However, both can be removed from newly native code words, leaving only the eight byte core of the word to be compiled.

## Enabling Native Compiling on New Words

By default, user-defined words are flagged with the Never-Native (NN) flag. While the words used in the definition of the new word might have been natively compiled into the new word, this new word will always be compiled with a JSR when used in future new words. To override this behavior and allow a user-defined word to be natively compiled, the user can use the `always-native` word just after the definition has been completed (with a semicolon). An example of doing this might be:

```
: double dup + ; always-native
```

Please note adding the `always-native` flag to a word overrides the `never-native` flag and it also causes the word to be natively compiled regardless of the setting of `nc_limit`.



Do not apply always-native to a word that has any kind of control structures in it, such as `if`, `case` or any kind of loop. If these words ever get native compiled, the JMP instructions used in the control structures are copied verbatim, causing them to jump back into the original words.



When adding your own words in assembly, if a word has a `jmp` instruction in it, it should have the NN (Never Native) flag set in the headers.asm file and should never have the AN (Always Native) flag set.

## `cmove`, `cmove>` and `move`

The three moving words `cmove`, `cmove>` and `move` show subtle differences that can trip up new users and are reflected by different code under the hood. `cmove` and `cmove>` are the traditional Forth words that work on characters (which in the case of Tali Forth are bytes), whereas `move` is a more modern word that works on address units (which in our case is also bytes).

If the source and destination regions show no overlap, all three words work the same. However, if there is overlap, `cmove` and `cmove>` demonstrate a behavior called "propagation" or "clobbering" : Some of the characters are overwritten. `move`: does not show this behavior. This example shows the difference:

```
create testbuf char a c, char b c, char c c, char d c, ( ok )
testbuf 4 type ( abcd ok )
testbuf dup char+ 3 cmove ( ok )
testbuf 4 type ( aaaa ok )
```

Note the propagation in the result. `move`, however, doesn't propagate. The last two lines would be:

```
testbuf dup char+ 3 move ( ok )
testbuf 4 type ( aabc ok )
```

In practice, `move` is usually what you want to use.

# Developing

After spending an entire weekend wrestling with blocks files, stacks, and the like, I was horrified and convinced that I had made a mistake. Who in their right mind would want to program in this godforsaken language! [DH]

— Doug Hoffman, Some notes on Forth from a novice user

## Adding New Words

The simplest way to add new words to Tali Forth is to include them in the file `forth_code/user_words.fs`. This is the suggested place to put them for personal use.

To add words to the permanent set, it is best to start a pull request on the GitHub page of Tali Forth. How to setup and use `git` and GitHub is beyond the scope of this document — we'll just point out it they are not as complicated as they look, and the make experimenting a lot easier.

During development, Tali Forth tends to follow a sequence of steps for new words:

- If it is an ANS Forth word, first review the standard online. In some cases, there is a reference implementation that can be used.
- Otherwise, check other sources for a high-level realization of the word, for instance Jonesforth or Gforth. A direct copy is usually not possible (or legally allowed, given different licenses), but studying the code provides hints for a Tali Forth version.
- Write the word in Forth in the interpreter. After it has been tested tested interactively, add a high-level version to the file `forth_code/forth_words.fs`.
- Add automatic tests for the new word to the test suite. Ideally, there will be test code included in the ANS Forth specification. If not, document what the test does.
- In a further step, if appropriate, convert the word to assembler. This requires adding an entry to `headers.asm` and the code itself to `native_words.asm`. In this first step, it will usually be a simple 1:1 sequence of `jsr` subroutine jumps to the existing native Forth words.
- If appropriate, rewrite all or some of the subroutine jumps in direct assembler. Because we have the automatic tests in place, we can be confident that the assembly version is correct as well.

However, if you are contributing code, feel free to happily ignore this sequence and just submit whatever you have.

## Deeper Changes

Tali Forth was not only placed in the public domain to honor the tradition of giving the code away freely. It is also to let people play around with it and adapt it to their own machines. This is also the reason it is (perversely) over-commented.

To work on the internals of Tali Forth, you will need the Ophis assembler.



## The Ophis Assembler

Michael Martin's Ophis Cross-Assembler can be downloaded from <http://michaelmartin.github.io/Ophis/>. It uses a slightly different format than other assemblers, but is in Python and therefore will run on pretty much any operating system. To install Ophis on Windows, use the link provided above. For Linux:

```
git clone https://github.com/michaelmartin/Ophis
cd Ophis/src
sudo python setup.py install
```

Switch to the folder where the Tali code lives, and run the Makefile with a simple `make` command. This also updates the file listings in the `docs` folder.

Ophis has some quirks. For instance, you cannot use math symbols in label names, because it will try to perform those operations. Use underscores instead.

### General Notes

- The X register is used as the Data Stack Pointer (DSP) and should only be used if there is no other alternative.
- The Y register, however, is free to be changed by subroutines. This also means it should not be expected to survive subroutines unchanged.
- Naively coded words generally should have exactly one point of entry — the `xt_word` link — and exactly one point of exit at `z_word`. In many cases, this requires a branch to an internal label `_done` right before `z_word`.
- Because of the way native compiling works, the trick of combining `jsr-rts` pairs to a single `jmp` instruction (usually) doesn't work.

### Coding Style

Until there is a tool for Ophis assembly code that formats the source file the way `gofmt` does for Go (`golang`), the following format is suggested.

- Tabs are **eight characters long** and converted to spaces.
- Opcodes are indented by **two tabs**.
- Function-like routines are followed by a one-tab indented "function doc string" based on the Python 3 format: Three quotation marks at the start, three at the end in their own line, unless it is a one-liner. This should make it easier to automatically extract the docs for them at some point.
- The native words have a special comment format with lines that start with `##` that allows the automatic generation of word lists by a tool in the `tools` folder, see there for details.
- Assembler mnemonics are lower case. I get enough uppercase insanity writing German, thank you very much.
- Hex numbers are, however, upper case, such as `$FFFE`.



The Ophis assembler interprets numbers with a leading zero as octal. This can be an annoying source of errors.

- Numbers in mnemonics are a stripped-down as possible to reduce visual clutter: use `lda 0,x` instead of `lda $00,x`.
- Comments are included like popcorn to help readers who are new both to Forth and 6502 assembler.

## Code Cheat Sheets

Programming computers can be crazy-making. [\[LB2\]](#)

— Leo Brodie, *Thinking Forth*

### The Stack Drawing

This is your friend and should probably go on your wall or something.

```

+-----+
|           ... |
+-         -+
|           |   ...
+- (empty) -+
|           | FE,X
+-         -+
... |           | FF,X
+=====+
$0076 |           LSB| 00,X  <-- DSP (X Register)
+-   TOS   -+
$0077 |           MSB| 01,X
+=====+
$0078 | (garbage) | 02,X  <-- DSP0
+-----+
$0079 |           | 03,X
+ (floodplain) +
$007A |           | 04,X
+-----+

```

### Coding Idioms

The first modern FORTH was coded in FORTRAN. Shortly thereafter it was recoded in assembler. Much later it was coded in FORTH. [\[CHM2\]](#)

— Charles Moore, *The Evolution of FORTH*

While coding a Forth, there are certain assembler fragments that get repeated over and over again. These could be included as macros, but that can make the code harder to read for somebody only

familiar with basic assembly.

Some of these fragments could be written in other variants, such as the "push value" version, which could increment the DSP twice before storing a value. We try to keep these in the same sequence (a "dialect" or "code mannerism" if you will) so we have the option of adding code analysis tools later.

- **drop** cell of top of the Data Stack

```
inx
inx
```

- **push** a value to the Data Stack. Remember the Data Stack Pointer (DSP, the X register of the 65c02) points to the LSB of the TOS value.

```
dex
dex
lda <LSB>      ; or pla, jsr kernel_getc, etc.
sta 0,x
lda <MSB>      ; or pla, jsr kernel_getc, etc.
sta 1,x
```

- **pop** a value off the Data Stack

```
lda 0,x
sta <LSB>      ; or pha, jsr kernel_putc, etc
lda 1,x
sta <MSB>      ; or pha, jsr kernel_putc, etc
inx
inx
```

## vim Shortcuts

One option for these is to add abbreviations to your favorite editor, which should of course be vim, because vim is cool. There are examples farther down. They all assume that auto-indent is on and we are two tabs into the code, and use # at the end of the abbreviation to keep them separate from the normal words. My `~/.vimrc` file contains the following lines for work on `.asm` files:

```
ab drop# inx<tab><tab>; drop<cr>inx<cr><left>
ab push# dex<tab><tab>; push<cr>dex<cr>lda $<LSB><cr>sta $00,x<cr>lda $<MSB><cr>sta
$01,x<cr><up><up><u>
ab pop# lda $00,x<tab><tab>; pop<cr>sta $<LSB><cr>lda $01,x<cr>sta
$<MSB><cr>inx<cr>inx<cr><up><up><up><up>>
```

# Future and Long-Term plans



This section is missing. See the GitHub page for further details.

# Tutorials

# Working with Blocks



This section is currently missing.

# The `ed` Line-Based Editor

While TECO was known for its complex syntax, `ed` must have been the most user-hostile editor ever created.[\[PHS\]](#)

— Peter H. Saulus, *The Daemon, the Gnu and the Penguin*

Tali Forth 2 comes with two editors, a traditional block-based editor of the type common with Forth, and the line-based editor `ed`, formally known as `ed6502`. This second editor is included because I like line-based editors. More to the point, the saved text uses less space than the block editor, where every block, regardless of how much text is in it, uses 1024 bytes. In contrast, `ed` uses one byte per character plus one end-of-line character per line. We'll see an example of this later.

The original `ed` was created by Ken Thompson and Dennis Ritchie along with the Unix operating system, sometime about 1971. It is terse, small, robust, and has a reputation for being completely unhelpful. An error is just signaled with a question mark (?). There isn't even a prompt unless it is explicitly turned on.



Newer versions of `ed` allow an option to display an explanation of the last error, but Tali doesn't have space for that. Error messages are for wimps. And who really needs a prompt anyway?

Commands in `ed` are single-letter commands like `a` or `p`. They can be prefixed with a combination of letters and special characters to designate the line numbers the command is supposed to work on. For example, `1,4d` deletes lines one to four.

## First steps with `ed`

Like its big brother `vi` (or its newer incarnation `vim`), `ed` has various modes, except that `ed` is so small it only has two. We start out in the *command mode* in which we accept, well, commands. Using `a` or `i` switches to *input mode* where all of the characters are added to the buffer. The first important thing is about how to get out of command mode: You type `.` (the period or dot) at the beginning of the line as the only character to return to command mode. A typical `ed` session will look something like this: [4: All quotes in the `ed` tutorial are taken from the *Mass Effect* games by BioWare/EA. As stated already, they hold the rights to all characters and whatnot.]

```
ed      ①
a      ②
After time adrift among open stars
Along tides of light
And through shoals of dust
I will return to where I began.
.      ③
      ④
```

① Start the editor from Tali Forth. It doesn't take anything on the stack.

②

Switch to insert mode and type the text.

- ③ The dot alone on the line signals the end of the text. We return to command mode.
- ④ The cursor moves down to the next line, without printing any confirmation. This is where you continue typing.

When you first use `ed`, you'll spend lots of time print what you've written and trying to figure out what the line numbers are. The commands for this are `p` (print without line numbers) and `n` (print with line numbers). The first special character prefix we'll learn for this is `%` (the percent symbol, alternatively a comma) works as well. This makes the command that follows it apply to the whole text.

```
%p      ①
After time adrift among open stars
Along tides of light
And through shoals of dust
I will return to where I began.
                ②
```

- ① This could also be `,p`
- ② Note again we return to an empty line.

The `%n` (or `,n`) command is usually more helpful because it gives you line numbers:

```
,n      ①
1      After time adrift among open stars
2      Along tides of light
3      And through shoals of dust
4      I will return to where I began.
```

- ① This could also be `%n`

Line numbers are indented automatically by one tab. Note we start counting with 1, not 0, because this is an editor for real humans.

Just entering the command directly without a line number will print the *current line*, which `ed` adjusts depending on what you are doing. After `a` it is the last line.



To find out which is the current line, type the `=` (equal sign) command.

This session could continue as such:

```
n
4      I will return to where I began.
```

The `d` (delete) command removes a line. Let's explicitly remove the second line:



2d

Again, `ed` provides absolutely no feedback on what just happened. We need to call `%n` (or `,n`) again if we are unsure:

```
%n
1   After time adrift among open stars
2   And through shoals of dust
3   I will return to where I began.
```

Note that lines three and four have moved up — they are now lines two and three.



To avoid confusion, when you have to delete a large number of lines, start at the bottom and move upwards towards the beginning of the text.

We can also use comma-separated numbers to indicate a range of lines (say, `1,2d`). As you probably will have guessed, or the `,` (or `%`) prefix can be used to delete the complete text. Be careful — in the real version of `ed`, you can undo changes with the `u` command. Tali's version currently doesn't support this option. If you delete it, it's gone.



The undo (`u`) function may be added to a future version of `ed` if space allows, at least for the last change.

Now, let's say we want to put back the second line. We can do this again with `a`, to add text *after* the first line. Note there is currently no way to paste the line we have just deleted. If we can't remember it, we're in trouble.

```
1a   ①
     And then, I, uh, did something
     .   ②
     .   ③
```

- ① Add text *after* the first line.
- ② The dot takes us out again.
- ③ Still no feedback.

Displaying our whole text with `%n` again, we get:

```
%n
1   After time adrift among open stars
2   And then, I, uh, did something
3   And through shoals of dust
4   I will return to where I began.
```

Lines three and four are numbered again as they were.

Instead of using `1a`, we could have used `2i` to insert the new line *before* line number two. Most long-term users of `ed` (like, all three of them) develop a preference for `a` or `i`. This is easy because `ed` accepts `0a` as a way to add new lines *before* the first line. In most other circumstances, line `0` is illegal. There is also the `$` prefix for the last line.



The combination `$=` will print the number of the last line.

## Saving Your Text

The only way to currently save text with `ed` on Tali is to write the buffer to a location in memory.

```
7000w ①  
128    ②
```

- ① The address in memory comes immediately before the `w` command with no space.
- ② `ed` returns the number of characters written, including the end-of-line characters. Yes, this is actually feedback of sorts. But don't get cocky!



Changing the number base hasn't been tested yet, so stick to decimal numbers for the time being when saving text.

The `w` command was originally created for files. Tali doesn't have files, just addresses. This means that you can write anything anywhere, at the risk of completely destroying your system. Really, really don't write anything to `0000`, which will overwrite the zero page of the `65c02`.

## Getting Out of `ed`

We can leave `ed` at any time with `Q` - note this is the capital letter "q". Any unsaved (unwritten, rather) text will be lost. The lowercase `q` will refuse to quit if there is still unwritten text. When in doubt, use `q`.

To access your text from the Forth command line, you can use standard Forth words like `type` with the address chosen and the length of the text provided after the `w` command.

```
7000 128 cr type ①  
After time adrift among open stars  
And then I, uh, did something  
And through the shoals of dust  
I will return to where I began.  
ok ②
```

- ① Place the `cr` word before the `type` word to prevent the first line of the text being placed right after the Forth command.
- ② We're back to the helpful Forth interpreter.



In future, `ed` might provide the address and length of the saved text on the data stack when quitting. The stack signature of `ed` would then change to `(--addr u)`. This would make further processing of the text easier.

You can also use `dump` to show how compact `ed` stores the text:

```
7000 128 dump
1B58 41 66 74 65 72 20 74 69 6D 65 20 61 64 72 69 66 After ti me adrif
1B68 74 20 61 6D 6F 6E 67 20 6F 70 65 6E 20 73 74 61 t among open sta
1B78 72 73 0A 41 6E 64 20 74 68 65 6E 20 49 2C 20 75 rs.And t hen I, u ①
1B88 68 2C 20 64 69 64 20 73 6F 6D 65 74 68 69 6E 67 h, did s omething
1B98 0A 41 6E 64 20 74 68 72 6F 75 67 68 20 74 68 65 .And thr ough the
1BA8 20 73 68 6F 61 6C 73 20 6F 66 20 64 75 73 74 0A shoals of dust.
1BB8 49 20 77 69 6C 6C 20 72 65 74 75 72 6E 20 74 6F I will r eturn to
1BC8 20 77 68 65 72 65 20 49 20 62 65 67 61 6E 2E 0A where I began..
1BD8 ok
```

- ① The dot in the text part of the hexdump at address \$157A is not the period at the end of the line, but the way `dump` displays the non-printable \$0A character. This control character marks the end of the line.

Note this text uses 128 bytes, in the block editor it would use one block of 1024 bytes.

## Programming with `ed`

You can use `ed` to write and save programs. Fire it up as usual:

```
ed
a
: myloop ( -- )      ①
    101 1 do i . loop ②
;
myloop
.
7000w
48
q
```

- ① Type normally as you would with any other editor.
- ② Any indentation has to be provided by hand. There is no auto-indent.

Running `7000 48 evaluate` will now print the numbers from 1 to 100.

## Further Information

This tutorial will be expanded as new commands become available. In the meantime, there are other sources:

- [https://en.wikipedia.org/wiki/Ed_\(text_editor\)](https://en.wikipedia.org/wiki/Ed_(text_editor)) Background and history
- <https://www.gnu.org/software/ed/ed.html> The official GNU ed page
- [https://www.gnu.org/software/ed/manual/ed_manual.html](https://www.gnu.org/software/ed/manual/ed_manual.html) The official GNU ed manual
- <https://sanctum.geek.nz/arabesque/actually-using-ed/> Small tutorial of Unix ed
- [http://www.psue.uni-hannover.de/wise2017_2018/material/ed.pdf](http://www.psue.uni-hannover.de/wise2017_2018/material/ed.pdf) A tutorial by B. W. Kernighan (yes, *that* Kernighan).

# Appendix

# Reporting Problems

The best way to point out a bug or make any other form of a comment is on Tali Forth's page on GitHub at <https://github.com/scotws/TaliForth2> There, you can "open an issue", which allows other people who might have the same problem to help even when the author is not available.

# FAQ

## Why does Tali Forth take so long to start up?

After the default kernel string is printed, you'll notice a short pause that didn't occur with Tali Forth 1. This is because Tali Forth 2 has more words defined in high-level Forth (see [forth_code/forth-words.fs](#)). The pause happens because they are being compiled on the fly.

## What happened to Tali Forth 1 anyway?

Tali Forth 1, informally just Tali Forth, was my first Forth. As such, it is fondly remembered as a learning experience. You can still find it online at GitHub at <https://github.com/scotws/TaliForth>. When Tali Forth 2 entered BETA, Tali Forth was discontinued. It does not receive bug fixes. In fact, new bugs are not even documented.

```
Booting Kernel for the Überquirrel Mark Zero
Scot W. Stevenson <scot.stevenson@gmail.com>
Kernel Version Alpha 002 (25. Feb 2014)
```

```
Tali Forth for the 65c02
Version ALPHA 003 (15. Apr 2014)
Tali Forth comes with absolutely no warranty
Type 'bye' to exit
```

```
words LOOP DO REPEAT WHILE UNTIL ELSE THEN IF WORDS DROP 2DROP ! @ >R R> R@ OVE
R 2OVER DUP ?DUP 2DUP SWAP 2SWAP NIP TUCK ROT DEPTH 1- 1+ FALSE TRUE BRANCH (BRA
NCH) ØBRANCH (ØBRANCH) BEGIN AGAIN (DO) I J LEAVE UNLOOP (LOOP) ABS DABS + - * /
*/ */MOD UM* M* UM/MOD UD/MOD SM/REM FM/MOD MOD /MOD M+ AND OR XOR INVERT NEGAT
E DNEGATE MAX MIN LSHIFT RSHIFT S>D D>S D+ D- <# # #S HOLD SIGN #> U. U.R UD. .
.R D. D.R 2 1 0 < = > 0= 0< COUNT >IN TYPE EMIT CHAR [CHAR] TOUPPER IN-PORT OUT-
PORT KEY BASE HERE PAD UNUSED ERASE FILL CELL+ CELLS CHAR+ CHARS C, C! C@ CMOVE
CMOVE> MOVE ALIGN ALIGNED ALLOT , ' >BODY >NAME EXECUTE EVALUATE STATE COMPILE,
: ; ] [ POSTPONE IMMEDIATE COMPILE-ONLY NATIVE-COMPILE CREATE DOES> VARIABLE 2VA
RIABLE CONSTANT +! SOURCE DECIMAL HEX BINARY DIGIT>NUMBER NUMBER >NUMBER BL CR S
PACE SPACES ." (".) S" (S") .( ( \ /STRING -TRAILING FIND ACCEPT PARSE PARSE-NAM
E .S SEE DUMP QUIT ABORT LITERAL (LITERAL) BYE ok
```

Figure 3. Screenshot of the Tali Forth 1 boot screen, version Alpha 3, April 2014

## What does the word "FEHLT" mean in unfinished code?

That's German for "missing". During development, I use it as a marker where stuff has to be added, sort of like another level of "TODO". Since there is no English word like that, it's very easy to find with `grep`.

## Who's "Tali"?

I like the name, and we're probably not going to have any more kids I can give it to. If it sounds vaguely familiar, you're probably thinking of Tali'Zorah vas Normandy a character in the *Mass Effect* universe created by BioWare. This software has absolutely nothing to do with neither the game nor the companies and neither do I, expect that I've played the whole series and enjoyed it. [5: Though I do wish they would tell us what happened to the quarian ark in *Andromeda*.]

### **And who is "Liara"?**

Liara Forth is another STC Forth for the big sibling of the 6502, the 65816. Tali Forth 1 came first, then I wrote Liara with that knowledge and learned even more, and now Tali 2 is such much better for the experience. And yes, it's another *Mass Effect* character.



# Testing Tali Forth 2

Tali Forth 2 comes with a test suite in the `tests` folder. It is based on the official ANS test code by John Hayes and was first adapted for Tali Forth by Sam Colwell.

To run the complete test, type `make test` from the main folder (this assumes a Unix-type system). Alternatively, switch to the test folder and start the `talitest.py` program with Python3. The tests should take only a very few minutes to run and produce a lot of output, including, at the end, a list of words that didn't work. A detailed list of results is saved to the file `results.txt`.

## User Tests

A special test file named `user.fs` is available for users to add their own tests. The results of this will be found just before the cycle tests near the end of `results.txt`. To run only this set of tests, you can use the command:

```
./talitest.py -t user
```

in the tests folder.

## Cycle Tests

The last set of tests, found in `cycles.fs`, determines cycle counts for the various built-in words. Users who are adding words may want to add cycle tests as well and there are instructions for doing that in that file. The cycle tests only work with the simulator and will not work on real hardware.

The cycle tests time (in 65C02 clock cycles) from the `jsr` that calls a word to the `rts` that returns from the word, including the `jsr` and `rts`. These cycle counts are the number of cycles if the word was used directly in interpreted mode. Some words will use more or fewer cycles depending on their input, so the cycles reported are for the input provided in the `cycles.fs` file.

The cycle tests work with some help from the `py65mon` simulator and extensions to it in `talitest.py`. Accesses to special addresses in the 65C02 memory map are used to start, stop, and read back the cycle counter in the simulator. A special word named `cycle_test` is created near the top of `cycles.fs` to help with this. It accepts the `xt` of the word you want to test (you can get the `xt` of any word by using the word `'`) and runs that word with the special memory accesses before and after, printing out the results.

## Cycle Tests and Native Compiling

Because Tali Forth 2 has native compiling capability, small words used in a word declaration will have their assembly code compiled directly into the word being defined, rather than using a `jsr`. This means that small words will not have the overhead of a `jsr` and `rts` when they are compiled into other words.

A perfect example of that is the built-in word `ALIGN`. This word has no assembly instructions (except

for an `rts`), but the cycle testing shows it takes 12 cycles. This is the number of cycles to run the word by itself, and it's the number of cycles to run a `jsr` instruction followed immediately by an `rts` instruction.

When this word is compiled into another word, however, Tali will use native compiling and will put the (empty) body of this word into the word being compiled rather than using a `jsr`. This results in 0 extra cycles for the word being defined. Twelve cycles will be saved for each small word that is natively compiled into a new definition. See the section on Native Compiling for more information.

## Old Tests



During early development, testing was done by hand with a list of words that has since been placed in the `old` folder. These tests might be still useful if you are in the very early stages of developing your own Forth.

# Thanks

Tali Forth would never have been possible without the help of a very large number of people, very few of whom I have actually met.

First, there is the crew at [6502.org](https://6502.org) who not only helped me build my own actual, working 6502 computer, but also introduced me to Forth. Tali would not exist without their inspiration, support, and feedback.

Special thanks go out to Mike Barry and Lee Pivonka, who both suggested vast improvements to the code in size, structure, and speed. And then there is Sam Colwell who contributed the invaluable test suite and a whole lot of code.

Thank you, everybody.

# References and Further Reading

- [FB] *Masterminds of Programming*, Federico Biancuzzi, O'Reilly Media 1st edition, 2009.
- [CHM1] "Charles H. Moore: Geek of the Week", redgate Hub 2009 <https://www.red-gate.com/simple-talk/opinion/geek-of-the-week/chuck-moore-geek>
- [CHM2] "The Evolution of FORTH, an Unusual Language", Charles H. Moore, *Byte* 1980, [https://wiki.forth-ev.de/doku.php/projects:the_evolution_of_forth](https://wiki.forth-ev.de/doku.php/projects:the_evolution_of_forth)
- [CnR] *Forth Programmer's Handbook*, Edward K. Conklin and Elizabeth Rather, 3rd edition 2010
- [DB] *Forth Encyclopedia*, Mitch Derick and Linda Baker, Mountain View Press 1982
- [DH] "Some notes on Forth from a novice user", Douglas Hoffman, Feb 1988 [https://wiki.forth-ev.de/doku.php/projects:some_notes_on_forth_from_a_novice_user](https://wiki.forth-ev.de/doku.php/projects:some_notes_on_forth_from_a_novice_user)
- [DMR] "Reflections on Software Research", Dennis M. Ritchie, Turing Award Lecture in *Communications of the ACM* August 1984 Volume 27 Number 8 <http://www.valleytalk.org/wp-content/uploads/2011/10/p758-ritchie.pdf>
- [EnL] *Programming the 65816, including the 6502, 65C02 and 65802*, David Eyes and Ron Lichty (Currently not available from the WDC website)
- [EW] "Forth: The Hacker's Language", Elliot Williams, <https://hackaday.com/2017/01/27/forth-the-hackers-language/>
- [GK] "Forth System Comparisons", Guy Kelly, in *Forth Dimensions* V13N6, March/April 1992 <http://www.forth.org/fd/FD-V13N6.pdf>{<http://www.forth.org/fd/FD-V13N6.pdf>}
- [JN] *A Beginner's Guide to Forth*, J.V. Nobel, <http://galileo.phys.virginia.edu/classes/551.jvn.fall01/primer.htm>
- [BWK] *A Tutorial Introduction to the UNIX Text Editor*, B. W. Kernighan, [http://www.psue.uni-hannover.de/wise2017_2018/material/ed.pdf](http://www.psue.uni-hannover.de/wise2017_2018/material/ed.pdf)
- [LB1] *Starting Forth*, Leo Brodie, new edition 2003, <https://www.forth.com/starting-forth/>{<https://www.forth.com/starting-forth/>}
- [LB2] *Thinking Forth*, Leo Brodie, 1984, <http://thinking-forth.sourceforge.net/#21CENTURY>
- [LL] *6502 Assembly Language Programming*, Lance A. Leventhal, OSBORNE/McGRAW-HILL 1979
- [PHS] "The Daemon, the Gnu and the Penguin", Peter H. Saulus, 22. April 2005, <http://www.groklaw.net/article.php?story=20050422235450910>

# Colophon

The Tali Forth 2 Manual was written with the [vim](#) editor in [AsciiDoc](#) format, formatted to HTML with AsciiDoctor, and version controlled with [Git](#), all under [Ubuntu](#) Linux 16.04 LTS.