

---

# Calypsi C compiler guide for the 65816

*Release 5.16*

Håkan Thörngren

Apr 14, 2026



---

## Contents

---

|     |                         |    |
|-----|-------------------------|----|
| 1   | Introduction            | 1  |
| 1.1 | Using this guide        | 1  |
| 1.2 | Copyright notice        | 1  |
| 1.3 | Disclaimer              | 2  |
| 2   | Installation            | 3  |
| 2.1 | macOS                   | 3  |
| 2.2 | Debian Linux            | 4  |
| 2.3 | Fedora Linux            | 4  |
| 2.4 | Arch and Manjaro Linux  | 5  |
| 2.5 | Windows                 | 5  |
| 2.6 | Using multiple versions | 6  |
| 3   | Tutorial                | 7  |
| 3.1 | Simulator               | 7  |
| 3.2 | C256 Foenix U           | 8  |
| 4   | Getting started         | 11 |
| 4.1 | C language              | 11 |
| 4.2 | File extensions         | 12 |
| 4.3 | Building applications   | 12 |
| 4.4 | Configuring             | 13 |
| 4.5 | Low level control       | 16 |
| 5   | Data storage            | 19 |
| 5.1 | 64K banks               | 19 |
| 5.2 | Ways to store data      | 20 |
| 5.3 | Address spaces          | 21 |
| 6   | Placing code and data   | 25 |
| 6.1 | Sections                | 25 |
| 6.2 | Data initialization     | 28 |
| 6.3 | Stack                   | 28 |
| 7   | C language              | 29 |
| 7.1 | Overview                | 29 |
| 7.2 | Extensions              | 29 |

|      |                                         |    |
|------|-----------------------------------------|----|
| 8    | Efficient coding                        | 31 |
| 8.1  | Data types . . . . .                    | 31 |
| 8.2  | Use prototypes . . . . .                | 32 |
| 8.3  | Use tables . . . . .                    | 33 |
| 9    | Target specifics                        | 35 |
| 9.1  | Board support . . . . .                 | 35 |
| 9.2  | Foenix C256 . . . . .                   | 36 |
| 9.3  | Foenix F256 . . . . .                   | 38 |
| 10   | Running the compiler                    | 41 |
| 10.1 | Basic invocation . . . . .              | 41 |
| 10.2 | Include search path . . . . .           | 42 |
| 10.3 | Compiler output . . . . .               | 42 |
| 10.4 | Diagnostics . . . . .                   | 44 |
| 10.5 | Command line options . . . . .          | 46 |
| 11   | Data representation                     | 55 |
| 11.1 | Basic types . . . . .                   | 55 |
| 11.2 | Structure types . . . . .               | 58 |
| 11.3 | Union types . . . . .                   | 59 |
| 11.4 | Enumeration types . . . . .             | 60 |
| 11.5 | Type qualifiers . . . . .               | 60 |
| 11.6 | Type definitions . . . . .              | 62 |
| 11.7 | Alignment . . . . .                     | 62 |
| 12   | Extended attributes                     | 63 |
| 12.1 | Overview . . . . .                      | 63 |
| 12.2 | Using attributes . . . . .              | 63 |
| 12.3 | Attribute reference . . . . .           | 64 |
| 13   | Language extensions                     | 69 |
| 13.1 | Overloaded functions . . . . .          | 69 |
| 13.2 | Statement expressions . . . . .         | 69 |
| 13.3 | Auto type . . . . .                     | 70 |
| 13.4 | Typeof . . . . .                        | 71 |
| 13.5 | Generics . . . . .                      | 71 |
| 14   | Pragma directives                       | 73 |
| 14.1 | Pragma directives reference . . . . .   | 73 |
| 15   | Intrinsic functions                     | 77 |
| 15.1 | Intrinsic functions reference . . . . . | 77 |
| 16   | Preprocessor                            | 83 |
| 16.1 | Overview . . . . .                      | 83 |
| 17   | Section reference                       | 89 |
| 17.1 | Section overview . . . . .              | 89 |
| 17.2 | Sections used by the compiler . . . . . | 90 |
| 18   | Runtime library                         | 97 |
| 18.1 | Design considerations . . . . .         | 97 |
| 18.2 | Using the library . . . . .             | 98 |
| 18.3 | Provided library files . . . . .        | 98 |
| 18.4 | Stubs interface . . . . .               | 99 |

|       |                                       |     |
|-------|---------------------------------------|-----|
| 18.5  | Examine use of library . . . . .      | 104 |
| 18.6  | Library on a diet . . . . .           | 105 |
| 18.7  | Override library functions . . . . .  | 106 |
| 18.8  | C startup . . . . .                   | 106 |
| 18.9  | Time and date . . . . .               | 107 |
| 18.10 | Rebuilding the C library . . . . .    | 108 |
| 19    | The optimizer . . . . .               | 109 |
| 19.1  | Overview . . . . .                    | 109 |
| 19.2  | Function inlining . . . . .           | 110 |
| 19.3  | Cross call . . . . .                  | 113 |
| 20    | Assembly language interface . . . . . | 115 |
| 20.1  | Intrinsic functions . . . . .         | 115 |
| 20.2  | Assembly functions . . . . .          | 115 |
| 20.3  | Calling convention . . . . .          | 118 |
| 20.4  | Inline assembler . . . . .            | 120 |
| 21    | Assembler . . . . .                   | 125 |
| 21.1  | Overview . . . . .                    | 125 |
| 21.2  | Syntax . . . . .                      | 125 |
| 21.3  | 65816 assembly . . . . .              | 126 |
| 21.4  | Sections . . . . .                    | 128 |
| 21.5  | Expressions . . . . .                 | 129 |
| 21.6  | Numeric constants . . . . .           | 130 |
| 21.7  | Location counter . . . . .            | 130 |
| 21.8  | Local labels . . . . .                | 130 |
| 21.9  | Directives . . . . .                  | 131 |
| 21.10 | Relocations . . . . .                 | 136 |
| 21.11 | Macro language . . . . .              | 139 |
| 22    | Running the assembler . . . . .       | 141 |
| 22.1  | Basic invocation . . . . .            | 141 |
| 22.2  | Include search path . . . . .         | 142 |
| 22.3  | Assembler output . . . . .            | 142 |
| 22.4  | Command line options . . . . .        | 146 |
| 23    | Linker . . . . .                      | 151 |
| 23.1  | Running the linker . . . . .          | 151 |
| 23.2  | Linking process . . . . .             | 152 |
| 23.3  | Simplified placement rules . . . . .  | 152 |
| 23.4  | ROM based systems . . . . .           | 153 |
| 23.5  | Host based systems . . . . .          | 154 |
| 23.6  | Objects . . . . .                     | 154 |
| 23.7  | Memory . . . . .                      | 156 |
| 23.8  | Linker rules file . . . . .           | 157 |
| 23.9  | Scatter . . . . .                     | 160 |
| 23.10 | Linker list files . . . . .           | 161 |
| 23.11 | Output formats . . . . .              | 165 |
| 23.12 | More on linker rules . . . . .        | 166 |
| 23.13 | Command line options . . . . .        | 166 |
| 24    | Librarian . . . . .                   | 173 |
| 24.1  | Creating a library . . . . .          | 173 |
| 24.2  | File format . . . . .                 | 173 |

|      |                                      |     |
|------|--------------------------------------|-----|
| 24.3 | Command line options . . . . .       | 173 |
| 25   | Debugger introduction . . . . .      | 175 |
| 25.1 | Command line debugger . . . . .      | 175 |
| 25.2 | Running the debugger . . . . .       | 176 |
| 25.3 | Command line options . . . . .       | 176 |
| 26   | Console command line . . . . .       | 181 |
| 26.1 | Command structure . . . . .          | 181 |
| 26.2 | A simple debugging session . . . . . | 182 |
| 27   | Remote debugging . . . . .           | 187 |
| 27.1 | Serial port . . . . .                | 187 |
| 28   | Visual Studio Code . . . . .         | 189 |
| 28.1 | Installation . . . . .               | 190 |
| 28.2 | Project setup . . . . .              | 191 |
| 28.3 | Building . . . . .                   | 191 |
| 28.4 | Running the debugger . . . . .       | 191 |
| 29   | Lua scripting . . . . .              | 195 |
| 29.1 | Running a script . . . . .           | 195 |
| 29.2 | Lua environments . . . . .           | 196 |
| 29.3 | Command interpreters . . . . .       | 196 |
| 29.4 | Console commands . . . . .           | 196 |
| 29.5 | MI commands . . . . .                | 197 |
| 29.6 | Adapted MI commands . . . . .        | 197 |
| 29.7 | Choosing a command set . . . . .     | 197 |
| 29.8 | Arrays and Lua . . . . .             | 197 |
| 29.9 | Debugger API . . . . .               | 197 |
| 30   | Lua data types reference . . . . .   | 199 |
| 30.1 | Address . . . . .                    | 199 |
| 30.2 | Adapted MI functions . . . . .       | 200 |
| 31   | Lua function reference . . . . .     | 201 |
| 31.1 | Execution . . . . .                  | 201 |
| 31.2 | Breakpoint . . . . .                 | 201 |
|      | Index . . . . .                      | 203 |

Welcome to the Calypsi C compiler tool chain for the 65816. This guide provides reference information to help you optimize your use of the compiler toolchain suite. It also offers coding technique suggestions and insights into design and implementation choices made for this C compiler.

## 1.1 Using this guide

This guide is aimed for those programming the 65816 and provides in depth reference information on how to use the the Calypsi C compiler tool chain.

You should have a working knowledge of:

- The 65816 architecture
- The C programming language
- The operating system of your computer
- The text editor or IDE of your choice

## 1.2 Copyright notice

Copyright Håkan Thörngren.

This document is under copyright and may not be reproduced without the written consent from Håkan Thörngren.

The Calypsi C compiler tool chain is licensed and must be used according to its terms.

The C library runtime is based on NuttX, which is under the Apache license. The license can be found in "Licenses/C Library License" within the installation directory. Assembly-written runtime parts are covered by the the Calypsi C compiler tool chain license, imposing no royalties or obligations when used with this product.

The Calypsi C compiler tool chain internally uses several open source components. Their licenses are in the "Licenses/Internal component licenses/" directory within the installation directory. These licenses do not impose obligations on any application built with the Calypsi C compiler tool chain.

The 64-bit floating-point library routines are implemented using the Berkeley SoftFloat library. Its license is found in Licenses/Berkeley SoftFloat License in the installation directory. This license applies only when using the `long double` type or the `--64bit-doubles` command-line option. It carries the obligation to include its copyright and license with any product that uses the Berkeley SoftFloat library.

---

Note: The runtime library does not contain any GPL licensed code and does not impose any royalty. Some components as outlined above do carry some obligations, e.g. to mention that they are included in a final product.

---

### 1.3 Disclaimer

What is described in this document is subject to change without notice and does not represent any commitments. While what is described in this document is believed to be accurate, Håkan Thörngren takes no responsibility for any errors or omissions.

In no event shall Håkan Thörngren, its employees, its contractors, or the authors of this document be liable for special, direct, indirect, or consequential damage, losses, costs, charges, claims, demands, claim for lost profits, fees, or expenses of any nature or kind.

This section contains installation information for supported platforms.

## 2.1 macOS

The macOS installer places the tools under `/usr/local`. Simply double-click the installation package to install the software.

After installation, files are in `/usr/local/calypsi-65816-5.16`, where 5.16 is the installed version. Each the Calypsi C compiler tool chain release creates a new directory, preventing overwrites of previous installations.

To simplify tool invocation, the installer script adds symbolic links from `/usr/local/bin` to the installation directory. Ensure `/usr/local/bin` is in your `PATH` environment variable to use the most recent installed version of the Calypsi C compiler tool chain. The installer also creates a symbolic link from `/usr/local/lib/calypsi-65816` to the installation directory, making it easy to refer to the installation without a version number.

### 2.1.1 Setting the PATH

To set the `PATH` environment variable on macOS, edit `/etc/paths` and add `/usr/local/bin`.

### 2.1.2 Uninstalling

To uninstall the Calypsi C compiler tool chain, run the uninstall script:

```
$ sudo /usr/local/lib/calypsi-65816-5.16/uninstall
```

This will remove the installation and the links in `/usr/local/bin`.

## 2.2 Debian Linux

The Debian package (.deb extension) is built and tested on Ubuntu 20.04. Installation on other Debian-based systems may be possible, but remains untested.

The package can be installed from the command line using:

```
$ sudo gdebi calypsi-65816-5.16.deb
```

---

Note: Download the Debian package to a local file before installation. `gdebi` does not fetch packages from a server.

---

If you do not have `gdebi` installed already, install it using:

```
$ sudo apt install gdebi-core
```

This installs the Calypsi C compiler tool chain in `/usr/local/lib/calypsi-65816-5.16`, where 5.16 is the installed version number.

To simplify tool invocation, the installer script adds symbolic links from `/usr/local/bin` to the installation directory. Ensure `/usr/local/bin` is in your `PATH` to use the latest the Calypsi tool chain. The installer also creates a symbolic link from `/usr/local/lib/calypsi-65816` to the installation directory. This allows referring to the installation without a version number.

It is also possible to install using `dpkg` in the following way:

```
$ sudo dpkg -i calypsi-65816-5.16.deb
```

This may fail if package dependencies are missing. The output will advise on necessary steps before retrying installation.

After installing any requested dependencies, retry the `sudo dpkg` installation command.

### 2.2.1 Uninstalling

To uninstall the Calypsi C compiler tool chain, use `dpkg` with the `-r` flag:

```
$ sudo dpkg -r calypsi-65816
```

## 2.3 Fedora Linux

The Fedora package (.rpm extension) is built and tested with Fedora version 40. Installation on other .rpm (Red Hat) based Linux systems may be possible, but remains untested.

The package file can be installed from the command line using:

```
$ sudo dnf install calypsi-65816-5.16-1.x86_64.rpm
```

---

Note: Download the package file to your local file system before installation.

---

This installs the Calypsi C compiler tool chain in `/usr/local/lib/calypsi-65816-5.16`, where 5.16 is the installed version number.

To simplify tool invocation, the installer script adds symbolic links from `/usr/local/bin` to the installation directory. Ensure `/usr/local/bin` is in your `PATH` to use the latest the Calypsi tool chain. The installer also creates a symbolic link from `/usr/local/lib/calypsi-65816` to the installation directory. This allows referring to the installation without a version number.

### 2.3.1 Uninstalling

To uninstall the Calypsi C compiler tool chain, use:

```
$ sudo dnf remove calypsi-65816-5.16-1.x86_64
```

## 2.4 Arch and Manjaro Linux

Manjaro is based on Arch Linux and uses `pacman`. Install the package file from the command line using:

```
$ sudo pacman -U calypsi-65816-5.16-x86_64.pkg.tar.zst
```

---

Note: Download the package file to your local file system before installation. `pacman -U` does not fetch packages from a server.

---

This installs the Calypsi C compiler tool chain in `/usr/local/lib/calypsi-65816-5.16`, where 5.16 is the installed version number.

To simplify tool invocation, the installer script adds symbolic links from `/usr/local/bin` to the installation directory. Ensure `/usr/local/bin` is in your `PATH` to use the latest the Calypsi tool chain. The installer also creates a symbolic link from `/usr/local/lib/calypsi-65816` to the installation directory. This allows referring to the installation without a version number.

### 2.4.1 Uninstalling

To uninstall the Calypsi C compiler tool chain, use `pacman` with the `-R` flag:

```
$ sudo pacman -R calypsi-65816
```

## 2.5 Windows

Extract the Windows product `calypsi-65816-5.16.zip` to a suitable location. It will expand into a folder `calypsi-65816-5.16`, allowing easy management of multiple versions.

### 2.5.1 Setting the PATH

To use the tools, add the bin folder inside the extracted archive to your system's search path. Search for "path" in the Start menu, then select "Edit the system environment variables - control panel".

### 2.5.2 Uninstalling

Uninstall by deleting the extracted archive via file explorer or command line.

## 2.6 Using multiple versions

On Linux, upgrading the software removes previous versions. On macOS, previously installed versions remain.

Sometimes it can be useful to have access to multiple versions of the software product. This can be done by copying the entire installation directory hierarchy to another location.

Use such copied installations by locally adjusting your PATH environment variable or by using an explicit path in the command line.

To remove a copied product, simply delete it. Do not run the `uninstall` script, as it will affect the normally installed copy.

---

Note: The software product is licensed under terms that govern copying and how it can be used.

---

This chapter shows how to build and run a simple “Hello World” application on different targets.

The command line examples offer a clear understanding of the process. Tools can also be used with an IDE, but setup varies.

Beyond the Calypsi C compiler tool chain, you will need Make to control builds and Git to obtain projects. An internet connection is also required to download example projects.

## 3.1 Simulator

The `db65816` contains a simulator mode which is enabled by default. It mimics the 65816 architecture and allows the application to run on your host machine.

From a terminal or command line, copy the example project from the internet using:

```
$ git clone https://github.com/hth313/Calypsi-65816-hello-world.git --recurse-submodules
```

Change to the newly cloned project directory:

```
$ cd Calypsi-65816-hello-world
```

Build the project with:

```
$ make
cc65816 --core=65816 --code-model=large --data-model=small --debug --list-file=obj/main-debug.lst -o
obj/main-debug.o src/main.c
ln65816 --debug -o hello.elf obj/main-debug.o linker-large-small.scm clib-lc-sd.a --list-file=hello-
debug.lst --cross-reference --rtattr printf=reduced --semi-hosted
```

This will create an output file `hello.elf`. To run the application you can load it into the debugger using:

```
$ db65816 hello.elf
Calypsi debugger for 65816
(db65816)
```

By default, the debugger enters interactive command mode. Run the application using the `run` command:

```
(db65816) run
running
Hello World!
program exited normally
(db65816)
```

When you are done with the debugger, you can leave it using the `quit` command:

```
(db65816) quit
$
```

You can also run the program directly from the command prompt, without entering interactive mode. You need to provide the debugger with additional command-line options to run and terminate automatically:

```
$ db65816 hello.elf -e run --terminate-on-program-exit
Calypsi debugger for 65816
run
running
Hello World!
program exited normally
$
```

This tells the debugger to execute the command `run` and that termination of the application also means termination of the debugger.

## 3.2 C256 Foenix U

This section shows how to build the “Hello World” application for a C256 Foenix U and run it on hardware.

This means you need to have access the actual hardware, but it may be possible to use a C256 emulator as well.

The the Calypsi C compiler tool chain generates ELF/DWARF by default, but this is not supported by the C256 Foenix kernel. To create an application suitable to run on the Foenix, the linker need output the application in the PGZ format. The provided `Makefile` is already configured to do this.

If you have not previously cloned the project, you can do so using:

```
$ git clone https://github.com/hth313/Calypsi-65816-hello-world.git --recurse-submodules
```

Change directory to the just cloned project:


```
$ cd Calypsi-65816-hello-world
```

When you are standing in the project folder, you can build the application for the C256 Foenix using:

```
$ make hello.pgz
```

This will result in a program file `hello.pgz`. In order to run it you need to put it on the removable SD card storage media you have with your C256 Foenix U. You can also run it in the Foenix emulator. In that case you need to select Tools>SD Card from its menu and direct it to the directory where the `hello.pgz` file is.

Once you have the hardware or emulator set up you can load and run the `hello.pgz` file with BRUN from the BASIC prompt:





This chapter provides an overview of using the compiler and its related tools.

### 4.1 C language

C is a widely used programming language. It is well-suited for hardware-level programming and also functions as a powerful, generic high-level language. It offers powerful abstractions over target machines, enabling efficient application development and precise control.

This implementation uses the ISO/IEC 9899:1999 standard, commonly known as C99. In this guide, it is referred to as Standard C.

#### 4.1.1 Cross compiler

The compiler is a cross-compiler, running on a modern workstation but producing applications for more constrained target machines.

#### 4.1.2 Supported devices

This Calypsi C compiler tool chain supports the Western Design Center 65816.

## 4.2 File extensions

The following table shows the file extensions normally used with the Calypsi C compiler tool chain.

Table 4.1: File extensions

| Extension | Purpose                                          |
|-----------|--------------------------------------------------|
| .c        | C source                                         |
| .h        | C header source                                  |
| .s        | Assembler source                                 |
| .o        | ELF/DWARF object file                            |
| .a        | Library (collection of object files)             |
| .lst      | List file                                        |
| .scm      | Linker rules                                     |
| .elf      | ELF/DWARF output (executable file)               |
| .hex      | Intel-hex output                                 |
| .srec     | Motorola S-record output                         |
| .s19      | Motorola S-record output, 16-bit address records |
| .s28      | Motorola S-record output, 24-bit address records |
| .s37      | Motorola S-record output, 32-bit address records |
| .raw      | Raw output                                       |
| .pgz      | Foenix binary format                             |
| .prg      | Commodore 64 binary format                       |

## 4.3 Building applications

Applications can be built from source files and libraries. Source files, written in C or assembly language, are compiled into object files: C source files use `cc65816`, and assembly source files use `as65816`.

A library is a collection of object files produced by the `nlib` tool, combining them with an index into a single file. The C runtime library is provided as an example; third-party libraries are also supported.

The `ln65816` takes object files, libraries, and placement rules as input to construct the executable application.

### 4.3.1 Compiler

The compiler command-line interface processes a single source file to produce an object file:

```
$ cc65816 source.c
```

The object file produced will have the same base name as the input but with a `.o` file extension.

---

Note: You typically need command-line options to select the CPU core, runtime models, and other settings.

---

### 4.3.2 Assembler

C projects do not require knowledge of assembly language. C simplifies programming and enables portability across architectures.

However, for highly specific target code, deep-level control, or critical routines, the assembler is indispensable.

The assembler command-line interface is similar to the compiler, with the main difference being the file extension:

```
$ as65816 source.s
```

The produced object file has the same `.o` extension as the compiler.

---

Note: You may need to provide a `--core` option for the assembler to recognize the target exact machine instructions.

---

### 4.3.3 Linker

The linker combines object files and libraries to create an executable application.

A rules file is required by the linker to describe the memory system, including placement rules for code and data. Stack and heap sizes can also be specified in this file.

You can run the linker from the command line as:

```
$ ln65816 myfile1.o myfile2.o rules.scm
```

This simplified command produces `aout.elf`, an ELF binary.

There are many ways to tailor the output:

- Using hex output, in either Intel HEX or Motorola S-record file format
- RAW output, which is just plain binary output of a single memory area
- Foenix program files, which are segmented binary files with start address
- The `--debug` option includes DWARF debugging information in the ELF executable.

The linker can produce a list file with cross-reference information, showing memory usage, placement, and why certain library contents were included.

## 4.4 Configuring

Tune compiler code generation using various command-line options.

The most basic settings to consider are:

- The CPU core (`--core`) in use. This controls the exact instruction set.
- The data model, which affects how data are placed and accessed in memory.
- The size of the `double` floating point type.
- Optimization settings.

### 4.4.1 Core

The Calypsi C compiler tool chain supports the WDC65816 core.

### 4.4.2 Code model

The code model governs default function placement and function pointer size.

The available code models are Small, Compact and Large. If not specified the compiler uses the Large code model which allows addressing functions in the entire memory range.

The Small and Compact code models restricts the code to a single 64K bank. In the Small code this is the first bank (low 64K) while the Compact code model allows the use any bank.

Table 4.2: Code models

| Code model | Function pointer size | Size limit |
|------------|-----------------------|------------|
| Small      | 16 bits               | 64K bytes  |
| Compact    | 16 bits               | 64K bytes  |
| Large      | 24 bits               | 16M bytes  |

### 4.4.3 Data model

The data model dictates default global data placement and defines the default pointer type (a data pointer without specific address attributes).

The available data models are Small, Medium Large and Huge. If not specified the compiler uses the Small data model.

Generally, smaller data models result in smaller and faster applications. Larger data models are primarily used for very large applications where avoiding explicit memory attributes is desired.

#### Small

In the small data model all default data is located in the first 64K of memory. This allows for using efficient 16 bits pointers, also for functions provided in the C library runtime.

You can still have data outside the first 64K, but you have to use keywords on such objects and pointers to them. This makes it possible to use 24 bits addressing to reach such data.

---

Note: In the Small data model, the `near` attribute is disabled because only a single 64K data bank (bank 00) is used.

---

## Medium

The Medium data model uses a designated 64K area (other than bank 00) for static data by default. While default pointers to this data are 24 bits, you can achieve efficient 16-bit addressing using a typed `near` pointer.

The Medium data model is useful for static data, such as tables or string literals, that do not fit within the limitations of the Small data model.

## Large

In the Large data model all pointers and data have the `far` attribute by default. This means that you can have data that span the entire 16MB memory range without using any keywords. Data object size is limited to 64K minus one byte. For larger objects, you must use the `huge` attribute or switch to the Huge data model.

For fast access, you can still use the `tiny` or `near` attribute on such objects.

## Huge

In the Huge data model all limitations are removed and you can address the entire memory range using 24 bits addressing with no limitation on object sizes. This also means that the code size will grow as the compiler is forced to generate code that is less efficient.

Table 4.3: Default pointer behavior in different data models

| Data model | Default pointer size | Default attribute | Max object size | Total size limit |
|------------|----------------------|-------------------|-----------------|------------------|
| Small      | 16 bits              | Bank 00           | 64K bytes       | 64K bytes        |
| Medium     | 24 bits              | <code>near</code> | 64K bytes       | 64K bytes        |
| Large      | 24 bits              | <code>far</code>  | 16M bytes       | 16M bytes        |
| Huge       | 24 bits              | <code>huge</code> | 16M bytes       | 16M bytes        |

---

Note: The total size limit applies to static or global objects without explicit memory attributes. You can go beyond the size limit imposed by the default pointer by using explicit memory attributes, e.g `far`.

---



---

Note: For shorter addressing you can always use the `tiny` attribute which uses the direct page.

---



---

Note: In the Small, Medium and Large data models, the `huge` attribute is disabled by default. This is because `huge` is seldom needed, and enabling it requires `size_t` to be 32 bits instead of the more efficient 16 bits on the 65816. If enabled with these data models, you will also need to rebuild the C runtime library. See [Rebuilding the C library](#) for how to build your own C runtime library.

---

#### 4.4.4 Size of double

The double floating-point type uses IEEE 754 format and can be set to either 32 or 64 bits using `--32bit-doubles` or `--64bit-doubles`. It defaults to 32 bits if not specified.

The float type is always 32 bits, and the long double type is always 64 bits.

#### 4.4.5 Optimization

Select the optimization level using the `-O` command-line option, which accepts a numeric argument of 1 or 2.

The compiler applies a number of techniques to reduce the width of expressions, select efficient code sequences and dead code elimination, regardless of optimization settings.

The `-O` option enables further optimization to reduce the code memory footprint and typically increase execution speed.

### 4.5 Low level control

This section provides a brief overview of controlling access to specific memory and built-in functions, also known as intrinsics.

#### 4.5.1 Extended keywords

The 65816 has a direct page short address area which can be accessed more efficiently than other memory addresses. This area has a single-byte address size.

You can specify that something is allocated in the direct page using extension keywords, either as `__tiny` or `__attribute__((tiny))`.

Additional data space is available in the near address area which uses a 16 bits address (two bytes). In the Small data model that uses a default pointer of 16 bits, this data area must be bank 00. The hardware stack also resides in this bank. In the other larger data models, the near address area can be any fixed 64K bank and does not have to be in bank 00. However, it must be tied to a designated 64K bank and cannot change during execution. You can specify that something is allocated in the near area using extension keywords, either as `__near` or `__attribute__((near))`.

The next data attribute is the far address area which can cover the entire memory space. The only limitation for data in the far area is that a single object cannot exceed 65535 bytes (64K minus one). You can specify that something is allocated in the far area using extension keywords, either as `__far` or `__attribute__((far))`.

Finally the huge address area covers the entire memory space and has no limitation on object size. You can specify that something is allocated in the far area using extension keywords, either as `__huge` or `__attribute__((huge))`.

Apart from different capability on how much and large objects that can be stored in a given address area, the compiler code that the compiler generates to access the different data areas differs. Accessing objects in larger, more general address spaces requires more instructions and executes slower than in smaller ones. Thus, it makes sense to plan where data objects are located and use the smallest possible one. You can of course put different objects in different address spaces, taking advantage of the performance of the tiny area for small frequently access objects while large and seldom used objects can be allocated in the larger address spaces.

The function attribute `__saveds` can be used on API and interrupt functions that may be called from another run-time context. It will set up the direct page and data bank when the function is called and restore the previous direct page and data bank on return. This is typically useful for entry points to an operating system that may be called by applications that operate with another direct page or data bank.

---

Note: For `saveds` API calls, you should generally prevent arguments from being passed on the direct page. Use either a single register argument or the `simple_call` calling convention to force arguments onto the stack. Direct page differences when using `saveds` mean accessing caller-placed arguments there will not work.

---

---

Note: When using `saveds` for API functions you most likely want to use the `saveds` attribute on C written interrupt functions as well, as interrupts may be triggered at any time.

---

## 4.5.2 Intrinsics

Intrinsics functions appear as ordinary calls but are special compiler constructs, emitting specific instruction sequences. To enable them, include the `calypsi/intrinsics65816.h` header file.

## 4.5.3 Assembly code

You can write functions in assembly language by following the C calling convention. These functions can be called from C in the same way as any C function.



This chapter discusses controlling data placement and its impact on application efficiency.

The 65816 has an address range of 16MB code and data space. Even though the address space 16MB, it can be seen as a sequence of 64K address ranges. While address calculations can cross 64K bank boundaries, this is not the most efficient way to utilize the 65816. Additionally, a single function must fit within a single 64K bank.

## 5.1 64K banks

There are 256 such 64K banks on the 65816 numbered 00 to FF hexadecimal. The first bank 00 is special in that the system stack and the direct page must be located in it. Certain vectors, such as the reset and interrupt vectors also reside in bank 00.

---

Note: Do not confuse the term bank here with a banked memory system where it often is used to refer to an address range which acts as a window to a larger memory area. A bank in the 65816 refers to a 64K memory area.

---

### 5.1.1 Direct page

The direct page is 256 byte large and can be located anywhere in the first 64K of memory. The compiler reserves 32 bytes from the direct page for pseudo registers, forming the tiny area which uses the `tiny` attribute. The 65816 has several efficient addressing modes related to this area, providing shorter instructions compared to other areas.

The DP (direct page) register is 16 bits and points to the start of the direct page. This register is set up before the `main()` function is called and is expected to remain fixed while the application is running.

### 5.1.2 Near bank

The near bank is the bank pointed to by the data bank register. This is a single 64K bank that is in the middle of addressing efficiency.

The DB or data bank register is an 8 bit register that points to the active near bank. This register is set up before the `main()` function is called and is expected to remain fixed for the duration of the execution.

In the Small data model the near bank must be the same bank as the CPU stack, which is always in bank 00. This is because the default pointer is 16 bits wide in the Small data model and it must be possible to point to a data object either on the stack or the default static storage area.

## 5.2 Ways to store data

Data can be allocated as auto, static, or on the heap. Variable scope and the need for allocation during runtime determine placement.

As a general rule, use auto-allocated variables whenever possible. This offers the compiler maximum flexibility for resource allocation closest to the core, which typically results in the most efficient data access.

### 5.2.1 Auto variables

Auto variables include function parameters and local variables not defined with `static`. The compiler attempts to allocate these in processor registers or direct page pseudo registers; otherwise, the stack is used.

Auto variables are allocated only when used. Their registers can be reused for other auto variables or temporary data once no longer needed. This reuse also applies to stack locations.

Auto variables can have multiple live ranges; a variable with the same name might be used in distinct parts of a function. Internally, these are treated as different variables and may be allocated to various locations, potentially not existing between live ranges.

All auto variables associated with a function are deallocated upon function exit.

---

Note: If you take the address of an auto variable, you can pass its pointer to other functions, which is useful for temporary storage. However, avoid using such pointers outside their scope, as auto variables are deallocated on function exit. Be aware that taking an address allocates the variable on the stack, potentially increasing access cost compared to other auto variables.

---

### 5.2.2 Static variables

Global, module, or function static variables are allocated in global memory, occupying space for the application's duration.

Their visibility varies. A global variable is universally visible but requires `extern` for use. A module static variable, declared with `static` at file scope, is visible within one compilation unit. A static local variable within a function's scope is only visible there, retaining its value across function calls. Use `static` to differentiate it from an auto variable.

### 5.2.3 Dynamically allocated

A dynamically allocated variable is retrieved from a heap using the malloc function. This is useful when the required data size is unknown at program startup.

---

Note: Dynamically allocated variables are a potential problem in memory constrained systems if the program is left running for a long time due to heap fragmentation.

---

## 5.3 Address spaces

The compiler provides multiple address spaces, which are addressable memory areas with specific properties:

- Address width for pointers
- Width of the associated integral index type
- Different instruction sequences for accessing various address spaces
- An extension keyword or type attribute name
- Section names tied to the address space for linking control

Address space attributes are always active in the compiler.

The Calypsi C compiler tool chain provides five address spaces, tiny (direct page), near (64K), far, far24 and huge. The latter three can address the full 16MB range where far and far24 limit object size to 64K minus one.

All unqualified pointers are either 16 or 24 bits wide, depending on the selected data model.

### 5.3.1 Tiny address space

The tiny address space is 256 bytes of memory located somewhere in the first 64K of memory. It has an address range `0x00-0xff`. This corresponds to the direct page on the 65816.

A tiny pointer is 16 bits wide and occupies two bytes when stored in memory. While a single byte would suffice, the compiler uses two bytes for efficiency.

Due to its 256-byte size and sharing with pseudo registers, the tiny area has limited storage. However, careful use can lead to a more efficient application.

### 5.3.2 Near address space

The near address space is a single 64K bank. This address space is not available in the Small data model. The bank can be located to any 64K bank and is fixed to that bank.

A near memory pointer is 16 bits and occupies two bytes when stored in memory.

### 5.3.3 Far address space

The far address space covers the entire 16MB address range. The only limitation on far objects is that they must be at most 64K (minus one byte) large.

A far memory pointer is 24 bits and occupies four bytes when stored in memory.

### 5.3.4 Far24 address space

The far24 address space cover the entire 16MB address range. The only limitation on far24 objects are that they must be at most 64K (minus one byte) large.

A far24 memory pointer is 24 bits and occupies three bytes when stored in memory.

---

Note: The only difference between far24 and far is that far24 is stored in memory without 32 bits padding. This is useful when hardware or API definitions require unpadding pointer storage. The generated code to store a 24 bits pointer is somewhat larger compared to a 32 bits pointer on the 65816.

---

### 5.3.5 Huge address space

The huge address space covers the entire 16MB address range. The maximum size of an object is 16MB (minus one byte).

A huge memory pointer is 24 bits and occupies four bytes when stored in memory.

### 5.3.6 Summary

The following table summarizes the available address spaces.

Table 5.1: Address spaces

| Memory type | Keyword              | Address range                    | Pointer size | Index type           |
|-------------|----------------------|----------------------------------|--------------|----------------------|
| Direct page | <code>__tiny</code>  | <code>0x00-0xff</code>           | 16 bits      | <code>int8_t</code>  |
| Near        | <code>__near</code>  | <code>0x0000-0xffff</code>       | 16 bits      | <code>int16_t</code> |
| Far         | <code>__far</code>   | <code>0x000000-0xffffffff</code> | 32 bits      | <code>int16_t</code> |
| Far24       | <code>__far24</code> | <code>0x000000-0xffffffff</code> | 24 bits      | <code>int16_t</code> |
| Huge        | <code>__huge</code>  | <code>0x000000-0xffffffff</code> | 32 bits      | <code>int32_t</code> |

### 5.3.7 Syntax

An address space attribute keyword such as `__tiny` is a type qualifier. Syntactically it works the same as other C language defined type qualifier, e.g. `const` and `volatile`.

The following declaration defines four variables in the Direct page address space:

```
__tiny int a, b;
int __tiny c, d;
```

The `__tiny` type qualifier applies to the closest type, `int` in this example. In C, the order between type qualifiers and types does not matter; they convey the same meaning.

### 5.3.8 Pointers

A pointer in C points to something in memory. Both the pointer itself and what it points to have types. As an example, the type `char *` is a pointer to a `char`.

Pointer types are easier to understand if you read them from right to left. The `*` is a pointer, so `char *` reads from right to left as “pointer to char”. This order of reading is especially useful when you mix in type qualifiers in pointer types, as it makes it a lot easier to read and understand what the type means.

```
int __attribute__((tiny)) * p1;
long * __attribute__((tiny)) p2;
```

Here, `p1` is a pointer stored in default memory that points to an `int` in Direct page memory. `p2` is a pointer stored in Direct page memory that points to a `long` in default memory.

### 5.3.9 Structures

You can place a structure in a specified address space. This means that all its members are in that address space. You cannot override individual structure members using an address space keyword. It is however possible to have members of the structure that point to a different address space.

```
struct tag {
    int __tiny * p;
    int value;
};

struct tag __tiny myTag;
```

This is however not allowed:

```
struct tag {
    int * __tiny p;      /* incorrect */
    int __tiny value;   /* incorrect */
};
```



---

### Placing code and data

---

When compiling for modern computers, code and data placement often requires minimal attention, as tools handle it effectively for same-machine targets.

Cross-compilation is more complex, requiring knowledge of the target system. Embedded systems are even more challenging, with widely varying memory systems necessitating precise control over code and data placement.

Effective embedded system tools provide control over code and data placement. This is achieved using building blocks that offer placement flexibility and control. The compiler assists by providing abstractions through type qualifiers, initializers, and runtime model settings, which simplify configuring the memory system.

### 6.1 Sections

A section is a unit of code or data. The compiler will place code and data into different sections which are saved as separate sections in the object file.

The compiler places data objects and functions into suitable sections following specific rules. For data objects, it considers if the object is `const`, its address space, and explicit initializers.

Table 6.1: Sections

| Section name    | Type   | Memory kind | Description                                                         |
|-----------------|--------|-------------|---------------------------------------------------------------------|
| code            | text   | ROM         | Executable code                                                     |
| compactcode     | text   | ROM         | Executable code                                                     |
| farcode         | text   | ROM         | Executable code                                                     |
| ztiny           | bss    | RAM         | Zero initialized (bss) direct page data                             |
| tiny            | data   | RAM         | Initialized direct page data                                        |
| zdata           | bss    | RAM         | Zero initialized (bss) data, default memory, Small data model only  |
| data            | data   | RAM         | Initialized near data, default memory, Small data model only        |
| cdata           | rodata | ROM         | Constant near data, default memory, Small data model only           |
| znear           | bss    | RAM         | Zero initialized (bss) near data, not available in Small data model |
| near            | data   | RAM         | Initialized near data, not available in Small data model            |
| cnear           | rodata | ROM         | Constant near data, not available in Small data model               |
| zfar            | bss    | RAM         | Zero initialized (bss) far data                                     |
| far             | data   | RAM         | Initialized far data                                                |
| cfar            | rodata | ROM         | Constant far data                                                   |
| zhuge           | bss    | RAM         | Zero initialized (bss) huge data                                    |
| huge            | data   | RAM         | Initialized huge data                                               |
| chuge           | rodata | ROM         | Constant huge data                                                  |
| switch          | rodata | ROM         | Switch tables                                                       |
| itiny           | rodata | ROM         | Tiny data initializers                                              |
| idata           | rodata | ROM         | Data initializers, default memory, Small data model only            |
| inear           | rodata | ROM         | Near data initializers, not available in Small data model           |
| ifar            | rodata | ROM         | Far data initializers                                               |
| ihuge           | rodata | ROM         | huge data initializers                                              |
| data_init_table | rodata | ROM         | Data initializer table                                              |
| registers       | noinit | RAM         | Pseudo registers in direct page                                     |
| reset           | text   | ROM         | Reset vector, when used                                             |
| heap            | noinit | RAM         | Heap memory, for <code>malloc()</code>                              |
| stack           | noinit | RAM         | CPU stack, anywhere in bank 0                                       |

The sections `itiny`, `inear`, `ifar`, `ihuge` and `data_init_table` in the table above are linker generated.

Note: It is assumed that there is no ROM in direct page. Constants placed in `tiny` address space are handled in the same way as non-constant data objects in the `tiny` address space. The only difference is that they have the `const` attribute in the C type system.

Note: The table assigns sections to ROM and RAM for ROM-based applications that start on power-up. For hosted systems loading from storage to RAM, this distinction is not applicable. Always consider ROM-marked sections as read-only.

Section types are derived from the UNIX world; see [Section kinds](#) in the assembler reference for details.

---

Note: When generating a ROM-executable application, the linker separates an initialized data section into two. Initializers go into a new `i-`prefixed section (placed in ROM); the data section is placed in RAM. The C runtime startup module copies initializers from ROM to RAM before giving control to the `main()` function.

---

### 6.1.1 Linking process

The linker reads object files with sections and a placement rules file. This file describes available memory areas and their permitted sections, optionally with additional placement constraints.

Memories containing actual value bits (program code or initialized data) are emitted in the executable.

An executable can serve as input for ROM or FLASH programmers, be loaded by the target machine, or be used by a debugger that connects to or simulates the target. Various executable formats are available to suit different execution environments.

---

Note: Although code and data are entirely separated into different sections, placing these sections together in the same memory may be desirable. This can be achieved at the link stage.

---

### 6.1.2 Placing sections in memory

The linker rules file uses file extension `.scm`, which is actually a source file in the Scheme language. However, you do not need to learn Scheme in order to use it, as it just a simple description of available memories and their associated properties. Sections are then bound to the memories. You can also define the size of the stack and heap in this file.

A simple example of such linker rules file is:

```
(define memories
  '((memory banks (address (#x10000 . #x1ffff)) (type ANY))
    (memory LoRAM (address (#x2100 . #x7fff)) (type ANY))
  )))
```

The Scheme interpreter evaluates the rules file, which defines a single global `memories` object. This object lists `memory` and `block` objects. Sections can be bound to memories either by specifying the memory type or by listing the sections within each memory.

The example simplifies memory system description. The linker deduces where compiler and linker-generated sections are placed; this requires specifying the memory type. To prevent automatic section placement in a memory, omit its type specification.

---

Note: You can specify a memory type and manually place custom sections. This is useful for sections requiring specific placement.

---

---

Note: Since the linker rules file is a Scheme program, it can theoretically construct the `memories` object via functions or macros. After evaluation, the Scheme interpreter is expected to leave a `memories` object in its global environment. However, a simple file like the example is usually sufficient.

---

## 6.2 Data initialization

A startup routine is performed before the `main()` function is called. This startup routine takes care of preparing the execution environment for the C program by initializing the stack pointer, initialize global variables and prepare system resources such as file streams and the heap, if used by the application.

### 6.2.1 Static data

Static data that needs to be initialized comes in two forms, zero initialized and those that have non-zero initializers. Zero initialization is done by filling such memory ranges with zero. Explicit initialization is done by copying an initializer section in ROM to its corresponding RAM area.

This process is table driven using a section named `data_init_table` which should be placed in the same memory as the `cdata` section, which holds constants.

---

Note: In a hosted environment, the linker provides the `--hosted` command-line option. This assumes the program image loads directly into RAM from external media. Variable initialization is done in-place via the load action; therefore, separate initializers are not required.

The `--target` option also has the effect of enabling running in a hosted environment.

---

## 6.3 Stack

The stack stores function return addresses, saved registers, and local variables. It tracks the current execution state, enabling reentrant and recursive functions.

### 6.3.1 Stack size considerations

The stack size is defined in the linker rules file. The stack size should be set low as possible, but not any lower. If given too much space you are wasting RAM space not being used for anything. However, if set too small it will cause corruption of other data.

Exact stack use can be hard to predict and it is better to err on the side of some safety. Also during development of your application you may want to use a bit oversized stack if possible, to allow you to focus on other issues. In the end, you are most likely constrained by some fixed amount of RAM and need to balance how much you can use for the stack with the optional heap and statically allocated data.

This chapter details the C language implementation, supporting the ISO/IEC 9899:1999 standard (C99).

### 7.1 Overview

For those accustomed to older C versions, C99 introduces several useful features:

- Declarations and statements can be mixed in the same block scope
- A declaration can be placed in the initialization expression of a `for` loop
- The `bool` datatype, available as `_Bool` or as `bool` after the `stdbool.h` header file has been included
- The `long long` data type
- C++ style comments
- Use of incomplete array at end of a structure
- Variable length arrays
- Initialization of compound objects by the use of designated initializers improve readability

### 7.2 Extensions

Language extensions are always enabled, encompassing various qualifier keywords, intrinsic functions, and additional library features.



The Calypsi C compiler tool chain is designed for coding in memory-constrained environments, typically embedded systems. This chapter explores proper data type usage and considerations for both smaller and larger targets.

Even if your current coding targets a specific system, the code may evolve and migrate to different system sizes. This chapter discusses trade-offs between such environments.

## 8.1 Data types

Data type selection significantly impacts overall program size and execution speed. Here are some rules of thumb:

- Generally, use the smallest possible types.
- Floating-point operations are typically inefficient, consuming space and often executing slowly.
- Use variables with the narrowest possible scope.
- Avoid taking a variable's address. If necessary, limit its use as much as possible. The compiler cannot allocate such variables in registers (registers have no addresses), which prevents the optimizer from fully understanding how function calls affect the variable.

### 8.1.1 Integer types

While smallest types are recommended, for larger targets, `int` and `unsigned int` can be more efficient, often being the sweet spot.

The `stdint.h` header provides integer types useful for selecting appropriate sizes, optimizing efficiency, and ensuring portability.

A type like `uint_fast16_t` defines an unsigned integer of at least 16 bits, optimized for speed. An 8-bit target will use a 16-bit type, while a 32-bit target might choose a 32-bit type if its instruction set handles 32-bit operations more efficiently than 16-bit operations.

### 8.1.2 Floating point types

Floating-point operations are often implemented using assembly or C library routines, which consume code space and incur execution overhead.

Alternatively, use fixed-point integers (Q notation). For instance, a Q24.8 number has 24 integer bits and 8 fractional bits, allowing manipulation with standard integer operations. This requires scaling values, which is straightforward.

### 8.1.3 Implicit conversions

C's conversion rules aim to convert integers to `int` and `float` to `double`. This can introduce extra code. If this is a concern, you must understand these rules or examine the compiler's generated code.

---

Note: While studying generated assembly code may seem daunting, a basic understanding of its mechanics can be highly beneficial. The compiler can generate a list file (`-l` or `--list-file`) showing C code intermixed with the resulting assembly.

---

## 8.2 Use prototypes

Functions can be declared using two styles. Traditional C, which lacks prototypes, is not recommended as it hinders compiler error detection. Prototypes offer safer, more efficient code generation by avoiding certain type promotion rules that introduce extra code. Traditional style is supported only for compatibility with older code. In this style, a declaration does not specify parameters:

```
/* Traditional style */
int foo(); /* declaration */

int foo(c, i) /* definition */
    char c;
    int i;
{
    return i - c;
}
```

In this case, the function takes a `char` parameter, which is promoted to `int` during the function call, introducing additional code.

Using prototypes with Standard C you specify parameters in the declaration:

```
/* Standard C prototype style */
int foo(char c, int i); /* declaration */

int foo(char c, int i) /* definition */
{
    return i - c;
}
```

In this case the `char` parameter is not promoted to `int` and is passed as a `char` type.

## 8.3 Use tables

Consider the trade-off between calculating a value and using a table for knowledge or operation encoding. For example, a sine operation in Q (fixed-point) notation can be approximated with a table lookup instead of calculation, making it very fast and compact.



This chapter describes the built-in target-specific support.

The support is enabled using the `--target` option which has the following effects:

- Select a suitable default `--core`.
- Enable preprocessor macros to identify the target and core used. This can be used for conditional compilation, in both the compiler and the assembler.
- Additional target specific header files can be included using `#include <file.h>`.
- Enable certain language extensions to suit the target, if applicable. This will also affect certain aspects of the code generator to take advantage of extensions and capabilities of the target.
- When using `--target` the linker will assume a RAM based system where a program is loaded into memory and data sections are initialized by loading the data in place. This is a sensible default, but can be controlled using the `--rom-code` and `--copy-initialize` options.

The installation comes with header files, pre-built target specific board support link libraries and linker control files.

## 9.1 Board support

Enabling board support in the linker requires two conditions: First, specify the appropriate `--target` option. Second, use a supplied linker control file whose name corresponds to the board support.

For example, this means that to select to link with a Foenix C256U board support you need to specify:

```
$ ln65816 --target=c256 c256-u-plain.scm [object-files]
```

---

**Note:** If you need to alter the supplied linker rules file or create your own, name it with the same prefix name as the board support one. The linker will first look in the current directory for the `.scm` file, then in the installation directory.

---

The following table gives an overview of the board support definition.

Table 9.1: Board support definitions

| --target | name      | kernel/OS | .scm prefix |
|----------|-----------|-----------|-------------|
| Foenix   | C256U/FMX | Kernel    | c256-       |

## 9.2 Foenix C256

The C256 family includes the FMX, U, and U+ new retro-style computers with a 65816.

### 9.2.1 Board support

The board support package is for the standard C256 kernel.

The board support includes C header files, a link library and linker control files from the [Foenix C256 support](#)<sup>1</sup> project. The current variant is built from commit hash `b3713b3c0fa27aac0640d7bd2f45ee7941d24e4a`.

### 9.2.2 Header files

The supplied header files are:

Table 9.2: Foenix C256 header files

| include path                  | description                   |
|-------------------------------|-------------------------------|
| #include <foenix/gavin.h>     | Gavin chip                    |
| #include <foenix/vicky.h>     | Vicky graphics chip           |
| #include <foenix/dma.h>       | DMA controller                |
| #include <foenix/interrupt.h> | Interrupt controller          |
| #include <foenix/network.h>   | Network interface             |
| #include <foenix/system.h>    | C256 system control registers |
| #include <foenix/timer.h>     | Timer definitions             |

The C256 library, which includes startup code, exit handling, and I/O adaption between the C library and the C256 kernel, comes pre-built. Your application will automatically be linked with it if you specify `--target=c256` and use a linker rules file starting with the `c256-` prefix:

```
$ ln65816 --target=c256 c256-u-plain.scm [object-files]
```

The installation provides several library variants to accommodate different compiler settings, such as double size. The linker automatically selects a compatible variant for your application.

---

<sup>1</sup> <https://github.com/hth313/Calypsi-65816-Foenix>

### 9.2.3 Linker rules

Linker rules use the `c256-` prefix. Naming your custom linker rules file with this prefix, or specifying `--target=c256`, ensures automatic linking with the supplied pre-built C256 kernel library.

`c256-u-plain.scm`

Suitable for use with the C256U.

```
;;; Example linker rules for the for the C256 Foenix U

(define memories
  '( (memory Banks (address (#x10000 . #x1fffff)) (type ANY))
      (memory LoMem (address (#x2100 . #x7fff)) (type ANY))
      (memory Vector (address (#xffe4 . #xffff)))

      (memory palettes (address (#xaf2000 . #xaf3fff))
        (section (palette0 #xaf2000) (palette1 #xaf2400))
        (section (palette2 #xaf2800) (palette3 #xaf2c00))
        (section (palette4 #xaf3000) (palette5 #xaf3400))
        (section (palette6 #xaf3800) (palette7 #xaf3c00)))
      (memory vram (address (#xb00000 . #xcfffff))
        (section vram))
  ))
```

`c256-u_-plain.scm`

Suitable for use with the C256U+. This variant has 4MB system RAM compared to 2MB in the C256U. If your application does not need the extra memory, consider using the `c256-u-plain.scm` variant to run on both C256U and C256U+.

```
;;; Example linker rules for the for the C256 Foenix U+

(define memories
  '( (memory Banks (address (#x10000 . #x2fffff)) (type ANY))
      (memory LoMem (address (#x2100 . #x7fff)) (type ANY))
      (memory Vector (address (#xffe4 . #xffff)))

      (memory palettes (address (#xaf2000 . #xaf3fff))
        (section (palette0 #xaf2000) (palette1 #xaf2400))
        (section (palette2 #xaf2800) (palette3 #xaf2c00))
        (section (palette4 #xaf3000) (palette5 #xaf3400))
        (section (palette6 #xaf3800) (palette7 #xaf3c00)))
      (memory vram (address (#xb00000 . #xcfffff))
        (section vram))
  ))
```

c256-fmx-plain.scm

Suitable for use with the C256-FMX.

```
;;; Example linker rules for the for the C256 Foenix FMX

(define memories
  '((memory Banks (address (#x10000 . #x2fffff)) (type ANY))
    (memory LoMem (address (#x2100 . #x7fff)) (type ANY))
    (memory Vector (address (#xffe4 . #xffff)))

    (memory palettes (address (#xaf2000 . #xaf3fff))
      (section (palette0 #xaf2000) (palette1 #xaf2400))
      (section (palette2 #xaf2800) (palette3 #xaf2c00))
      (section (palette4 #xaf3000) (palette5 #xaf3400))
      (section (palette6 #xaf3800) (palette7 #xaf3c00)))
    (memory vram (address (#xb00000 . #xefffff))
      (section vram))
  ))
```

## 9.3 Foenix F256

The F256 family includes the F256 and F256K.

### 9.3.1 Board support

The board support package provides a flat memory load for the 65816 on the F256 series computers.

The board support includes C header files, a link library and linker control files from the [F256 support](#)<sup>2</sup> project. The current variant is built from commit hash 8b4cec98c8281e2988c97481209502a03a7044d4.

### 9.3.2 Header files

The supplied header files are:

Table 9.3: Foenix F256 header files

| include path                | description                                  |
|-----------------------------|----------------------------------------------|
| #include <f256.h>           | Main header to include all F256 header files |
| #include <f256/dma.h>       | DMA controller                               |
| #include <f256/interrupt.h> | Interrupt controller                         |
| #include <f256/rts.h>       | Real time clock                              |
| #include <f256/sdc.h>       | SD card                                      |
| #include <f256/sid.h>       | SID                                          |
| #include <f256/system.h>    | F256 system control registers                |
| #include <f256/uart.h>      | UART, serial port                            |
| #include <f256/via.h>       | 6522 VIA                                     |
| #include <f256/vicky.h>     | Vicky graphics chip                          |

<sup>2</sup> <https://github.com/hth313/Calypsi-65816-F256>

The F256 library includes startup code; other specifics are absent due to the current lack of a ready kernel. Your application will automatically be linked with it if you specify `--target=f256` or use a linker rules file starting with the `f256-` prefix:

```
$ ln65816 --target=f256 f256-plain.scm [object-files]
```

The installation provides several library variants to accommodate different compiler settings, such as double size. The linker automatically selects a compatible variant for your application.

### 9.3.3 Linker rules

Linker rules use the `f256-` prefix. Naming your custom linker rules file with this prefix, or specifying `--target=f256`, ensures automatic linking with the supplied pre-built F256 board support library.

`f256-plain.scm`

This file defines the basic flat memory map.

### 9.3.4 I/O and Near addressing

The 65816 can address 64K using shorter 16-bit addressing. In the Small data model, bank 0 is set up as the Near addressing area. In the Medium data model, a single bank is set aside for this purpose. However, in the Large data model, the Near bank is set to point to the I/O bank. The C startup module in this board support configures this.



---

## Running the compiler

---

The compiler is run from the command line, or from an IDE that interfaces with the compiler using the command line.

### 10.1 Basic invocation

The compiler is invoked in the following way:

```
$ cc65816 [options] sourcefile [options]
```

As an example, to compile a source file with debugging information and a list file, you can use:

```
$ cc65816 --debug source.c -l
```

Command line options are optional arguments that tune the behavior of the compiler. They always start with a dash character. There are two variants, single letter options (-l to instruct the compiler to create a list file) starts with a single dash. The other variant is long descriptive options that starts with two dashes.

Some options require arguments. These follow the option, separated by a space or an equals sign (=). For single-argument options, the separator is optional:

```
$ cc65816 -Iinclude source.c -D VERBOSE=2 --list-file=tiny-source.lst
```

In this case the -I option adds `include` as a directory to scan for header files. The symbol `VERBOSE` is defined in the preprocessor with value 2. A list file with a specific name is also produced.

The order in which the options appear normally do matter, except for the -I option that adds directories to search for header files. Such directories are searched in the order in which they appear on the command line.

To display the version of the compiler, use `--version`:

```
$ cc65816 --version
Calypsi ISO C compiler for 65816 version 5.16
```

Command line options are described in [Command line options](#).

## 10.2 Include search path

Header files are included by surrounding the filename with either double quotes or angled brackets:

```
#include "myheader.h"
#include <stdio.h>
```

Angled brackets are typically for system header files, while double quotes are for application-specific header files.

The search order for include files is as follows:

1. Relative to compilation directory (double quote include only)
2. In directories specified in the `-I` option in the order they appear on the command line
3. The system directory, which is the installation directory

Header files specified in angle brackets are not searched relative to the compilation directory. Otherwise, they behave identically. The system header file directory is located within the installation directory.

---

Note: More precisely, the system header file directory is relative to the `cc65816` executable. If an installation is moved, it will still find the correct system header file directory.

---

## 10.3 Compiler output

The compiler outputs one or two files, an object file that can be linked with other object files and libraries by `ln65816` to produce an executable file, and optionally a list file.

### 10.3.1 Object file

The object file is in ELF format, optionally including DWARF debugging information if `--debug` (or `-g`) is specified.

It contains:

- A symbol table
- Relocatable sections for code and data
- Relocations, allowing linker to modify address values after section placement
- DWARF-formatted source-level debugging information (if `--debug` was specified)
- Vendor-specific data, including runtime attributes and additional the Calypsi C compiler tool chain-specific information not covered by ELF/DWARF, used by the linker and debugger, with section types `0x8000000a` and `0x8000000b`.

---

Note: DWARF version 5 is used. The implementation covers the needs of the Calypsi C compiler tool chain and includes vendor extensions.

---

### 10.3.2 List file

The list file is a text file meant to be shown with a fixed width font. It contains a header that shows compiler, version, time when it was created and the command line used. The C source file follows intermixed with generated assembly code. Finally there is a summary of the code and data sizes.

The following simple summation function:

```
int sum(int a, int b, int c) {
    if (a < b) {
        return a + c;
    } else {
        return b + c;
    }
}
```

Compiling with the -l option produces a list file:

```
#####
#
# Calypsi ISO C compiler for 65816                      version 5.16 #
#                                                         14/Apr/2026 16:42:18 #
# Command line: sum.c -l                                #
#                                                         #
#####

\ 000000          .rtmodel version,"1"
\ 000000          .rtmodel codeModel,"large"
\ 000000          .rtmodel dataModel,"small"
\ 000000          .rtmodel core,"65816"
\ 000000          .rtmodel huge,"0"
\ 000000          .rtmodel target,"none-specified"
\ 000000          .extern _Dp
\ 000000          .extern _Vfp
0001          int sum(int a, int b, int c) {
\ 000000          .section farcode,text
\ 000000          .public sum
\ 000000 5a      sum:      phy
\ 000001 5a      phy
\ 000002 aa      tax
\ 000003 a5..    lda      dp:.tiny _Dp
\ 000005 8301    sta      1,s
\ 000007 a5..    lda      dp:.tiny (_Dp+2)
\ 000009 8303    sta      3,s
0002          if (a < b) {
\ 00000b 8a      txa
\ 00000c 38      sec
\ 00000d e301    sbc      1,s
\ 00000f 5003    bvc      `?L9`
\ 000011 490080 eor      ##-32768
\ 000014 1006    `?L9`:   bpl      `?L4`
0003          return a + c;
\ 000016 8a      txa
\ 000017 18      clc
\ 000018 6303    adc      3,s
\ 00001a 8005    bra      `?L3`
\ 00001c          `?L4`:
0004          } else {
```

(continues on next page)

(continued from previous page)

```
0005          return b + c;
    \ 00001c a303          lda    3,s
    \ 00001e 18           clc
    \ 00001f 6301          adc    1,s
    \ 000021      `?L3`:
0006          }
0007      }
    \ 000021 7a           ply
    \ 000022 7a           ply
    \ 000023 6b           rtl
```

```
#####
#                               #
# Memory sizes (decimal) #
#                               #
#####
```

Executable (Text): 36 bytes

### 10.3.3 Assembly source

The compiler can generate assembly source output, which is similar to a list file, except that this output file is actual source code that can be given to the assembler. This can be used to generate a starting point for an assembly source file. By providing suitable declarations and simplified C functions you can study what the compiler does and build upon it.

## 10.4 Diagnostics

Compiler diagnostic messages identify problems that either must or should be addressed.

### 10.4.1 Errors

An error identifies a problem in the code that prevents the compiler from producing a valid output file.

### 10.4.2 Warnings

A warning identifies a potential problem in the code that you probably want to take a look at and possibly rectify. The compiler is still able to produce a valid output file.

### 10.4.3 Internal errors

An internal error indicates an unexpected problem within the compiler product.

### 10.4.4 Controlling warnings

Diagnostic can be controlled using the `-W` option that takes a wide range of possible alternatives. It is beyond the scope of this guide to list them all. A complete reference can be found at <https://clang.llvm.org/docs/DiagnosticsReference.html> as the parser and diagnostics engine is the same as in the Clang project.

Consider the following code:

```
int test(int x)
{
    if (x = 6) {
        return 10;
    } else {
        return 0;
    }
}
```

In the code, an assignment is used in a comparison, a common C error. While legally testing a non-zero assignment expression, this often indicates an intended use of `==`. The compiler cannot know your intent, but flags a potential problem. Compiling this code results in warnings:

```
w.c:3:10: warning: using the result of an assignment as a condition without parentheses [-Wparentheses]
    if (x = 6) {
        ~^~
w.c:3:10: note: place parentheses around the assignment to silence this warning
    if (x = 6) {
        ^
        (  )
w.c:3:10: note: use '==' to turn this assignment into an equality comparison
    if (x = 6) {
        ^
        ==
```

The warnings point out that there may be a problem in the code. The name of the warning that triggered is also given (`-Wparentheses`). Finally, it makes suggestions on how you can change the code.

If you think the warning is of no use to you and you do not want to see this particular warning again, you can add the option `-Wno-parentheses` to the command line to silence it.

Adding an extra pair of parentheses around the assignment is a common way to indicate an intentional assignment within a test expression:

```
int test(int x)
{
    if ((x = 6)) {
        return 10;
    } else {
        return 0;
    }
}
```

Finally, if the intention with the code was to use the equality operator, you should change it to use the correct operator:

```
int test(int x)
{
    if (x == 6) {
        return 10;
    } else {
        return 0;
    }
}
```

## 10.5 Command line options

This section covers the cc65816 command-line options in detail.

### 10.5.1 Options overview

If you run the compiler from the command line without any arguments it will complain that the input source file is missing and then shows a short form help:

```
$ cc65816
Missing: FILE

Usage: cc65816 [--version] [-o|--output-file OUTPUT-FILE] [-l]
      [--list-file LIST-FILE] [-I DIRECTORY] [-D IDENTIFIER]
      [-U IDENTIFIER] [-g|--debug] [--32bit-doubles] |
      [--64bit-doubles]) [--char-is-signed] | [--char-is-unsigned])
      [-E] [--print-macro-definitions] [--align-functions ALIGNMENT]
      [-O LEVEL] [--space] | [--speed]) [--no-cross-call]
      [--no-interprocedural-cross-jump] [--no-inline] |
      [--always-inline]) [--only-marked-as-inline] [--strong-inline]
      [--inline-on-matching-custom-text-section] [--rtattr NAME=VALUE]
      [--weak-symbols] [--no-vector-sections] [-c]
      [--assembly-source ASSEMBLY-FILE] [-S] [--force-switch STRATEGY]
      [-W ARG] [--dependencies] [-M ARG] [--pedantic-errors]
      [--include-system DIRECTORY] [--include-system-after DIRECTORY]
      [--core CORE] [--target TARGET] [--code-model NAME]
      [--data-model NAME] [--pascal-strings] [--enable-huge-attribute]
      [--no-ppu-mul] FILE
  use 'cc65816 --help' for detailed help
```

For more detailed help, use the `--help` option:

```
$ cc65816 --help
Calypsi ISO C compiler for 65816 version 5.16

Usage: cc65816 [--version] [-o|--output-file OUTPUT-FILE] [-l]
      [--list-file LIST-FILE] [-I DIRECTORY] [-D IDENTIFIER]
      [-U IDENTIFIER] [-g|--debug] [--32bit-doubles] |
      [--64bit-doubles]) [--char-is-signed] | [--char-is-unsigned])
      [-E] [--print-macro-definitions] [--align-functions ALIGNMENT]
      [-O LEVEL] [--space] | [--speed]) [--no-cross-call]
      [--no-interprocedural-cross-jump] [--no-inline] |
      [--always-inline]) [--only-marked-as-inline] [--strong-inline]
      [--inline-on-matching-custom-text-section] [--rtattr NAME=VALUE]
```

(continues on next page)

(continued from previous page)

```

    [--weak-symbols] [--no-vector-sections] [-c]
    [--assembly-source ASSEMBLY-FILE] [-S] [--force-switch STRATEGY]
    [-W ARG] [--dependencies] [-M ARG] [--pedantic-errors]
    [--include-system DIRECTORY] [--include-system-after DIRECTORY]
    [--core CORE] [--target TARGET] [--code-model NAME]
    [--data-model NAME] [--pascal-strings] [--enable-huge-attribute]
    [--no-ppu-mul] FILE
    use 'cc65816 --help' for detailed help

```

## Available options:

```

--version                Display version number
-o,--output-file OUTPUT-FILE
                        Name of output file
-l                        Generate a list file, named by appending '.lst' to
                        input file
--list-file LIST-FILE    Generate list file, using given name
-I DIRECTORY             Include directory
-D IDENTIFIER            Predefine a macro
-U IDENTIFIER            #undef a predefined macro
-g,--debug              Produce debugging information
--32bit-doubles          Make the 'double' data type 32-bits (this is the
                        default)
--64bit-doubles          Make the 'double' data type 64-bits
--char-is-signed         Treat 'char' as a signed type
--char-is-unsigned       Treat 'char' as an unsigned type (this is the
                        default)
-E                      Stop after preprocessing
--print-macro-definitions
                        Print macro definitions in -E mode in addition to
                        normal output
--align-functions ALIGNMENT
                        Align the start of functions to given alignment
-O LEVEL                Enable optimizer, -O0, -O1 or -O2
--space                 Optimize for space (this is the default, use together
                        with -O)
--speed                 Optimize for speed (use together with -O)
--no-cross-call          Disable cross call optimization
--no-interprocedural-cross-jump
                        Disable interprocedural cross jump optimization
--no-inline             Disable all function inlining
--always-inline          Always inline functions
--only-marked-as-inline  Only consider inlining functions marked as 'inline'
--strong-inline          Always (try to) inline functions marked as 'inline'
--inline-on-matching-custom-text-section
                        Only inline custom text section functions when the
                        section names match
--rtattr NAME=VALUE      Define a runtime attribute (identifier or quoted
                        string value accepted)
--weak-symbols           Make all public symbols entries weak
--no-vector-sections     Discard auto generated interrupt vector sections
-c                      Generate object file and do not try to link
--assembly-source ASSEMBLY-FILE
                        Generate an assembly source file, using given name
-S                      Generate an assembly source file, named by appending
                        '.s' to input file
--force-switch STRATEGY  Switch strategy, one of 'if-else', 'jump-table' or

```

(continues on next page)

(continued from previous page)

|                                  |                                                                                              |
|----------------------------------|----------------------------------------------------------------------------------------------|
|                                  | 'value-table'                                                                                |
| -W ARG                           | Warning control                                                                              |
| --dependencies                   | Generate dependencies to stdout                                                              |
| -M ARG                           | Dependency file control                                                                      |
| --pedantic-errors                | Error on language extensions                                                                 |
| --include-system DIRECTORY       | Add system include directory                                                                 |
| --include-system-after DIRECTORY | Add system include directory after                                                           |
| --core CORE                      | Core, one of '65816' (defaults to '65816')                                                   |
| --target TARGET                  | Target system, one of 'C256', 'F256' or 'SNES'<br>(defaults to embedded/ROM use, if omitted) |
| --code-model NAME                | Code model, one of 'small', 'compact' or 'large'<br>(defaults to 'large')                    |
| --data-model NAME                | Data model, one of 'small', 'medium', 'large' or<br>'huge' (defaults to 'small')             |
| --pascal-strings                 | Generate Pascal strings on leading \p character                                              |
| --enable-huge-attribute          | Enable the huge data memory attribute, also makes<br>size_t 32 bits wide                     |
| --no-ppu-mul                     | Do not use the PPU multiplier (SNES target only)                                             |
| -h,--help                        | Show this help text                                                                          |

## 10.5.2 Options in detail

### --version

Displays the name and version of the compiler.

### --output-file, -o

Specify the output object file name. If omitted, the output file is derived from the input file name (excluding path) with a .o extension, written to the current directory by default.

This option can also alter the output file name and provide a directory path. The specified directory must already exist.

### -l

Generate a list file. The name used is the name of the input file (ignoring any directory path) with a .lst file extension. See also --list-file.

### --list-file

Generate a list file. The name of the list file is given as argument to this option. See also -l to generate a list file based on the source filename.

**-I**

Add a directory to the current include search path. This option can be used multiple times on a command line. The order in which they appear specifies the search order between the directories.

The system include directory is always added last to the search order list.

**--include-system**

Add a directory to the current system include search path before the provided system include directories. This option can be used multiple times on a command line. The order in which they appear specifies the search order between the directories.

**--include-system-after**

Add a directory to the current system include search path after the provided system include directories. This option can be used multiple times on a command line. The order in which they appear specifies the search order between the directories.

**-D**

Define a macro. This takes an argument with the symbol name and optionally an assignment value `-Dsymbol[=value]`. If no value is given, the macro is given the value 1.

**-U**

Undefine a macro. This takes an argument with the symbol name to be undefined.

**--debug**

Generate DWARF symbolic debugging information. In order to get debugging information all the way to the debugger, the linker must also be given this option.

When debugging information is enabled the `__CALYPSI_DEBUG__` macro is also defined and set to 1.

**-g**

Synonym for `--debug`.

**-S**

Generate an assembly source file as output. The name used is the name of the input file (ignoring any directory path) with a `.s` file extension. No object file is generated. See also `--assembly-source`.

**--assembly-source**

Generate an assembly source file as output. The name of the output file is given as an argument to this option. See also `-S` to generate an assembly source file based on the C source filename. No object file is generated.

**--32bit-doubles**

The double floating point data type has 32 bits precision.

**--64bit-doubles**

The double floating point data type has 64 bits precision.

**--char-is-signed**

Makes the `char` data type signed. The default is unsigned.

**--char-is-unsigned**

Makes the `char` data type unsigned. This is also the default.

**-E**

Stop after running the preprocessor. The output from the preprocessor is written to standard output (`stdout`).

**--print-macro-definitions**

This is used together with the `-E` option to also write out all macro definitions.

**-O**

Runs the optimizer, expecting an argument of 0, 1, or 2. `-O0` (the default) performs no additional optimization passes but still produces good code. `-O1` enables additional passes to further improve code, and `-O2` enables all optimizer passes.

**--space**

Optimize for space, this is the default. Use `-O` to enable optimizations.

**--speed**

Optimize for speed. Use `-O` to enable optimizations. This option controls decisions about tradeoffs and will balance towards making the code run faster at the expense of additional code space. See also `--no-cross-call` below.

**--no-cross-call**

Disables the cross-call optimizer, normally enabled at `-O2`. The cross-call optimizer extracts common code sequences into small subroutines, which can significantly reduce code space, so it is also enabled by default with `--speed`.

**--no-inline**

Disable all function inlining which is normally enabled at `-O1`.

**--always-inline**

Enable function inlining regardless of optimization level setting.

**--only-marked-as-inline**

Only considers functions marked with the `inline` keyword for inlining. This disregards small functions and single-use static functions for inlining.

**--strong-inline**

Regard functions marked with the `inline` keyword as a strong hint that they should be inlined. The compiler may still decide not to do it.

**--rtattr NAME=VALUE**

This defines a runtime attribute which is written to the object file. This is used to specify runtime attribute value that can be used in the linker for checking object file consistency.

**--weak-symbols**

Makes all public symbols in the object file weak. This is normally not needed, but can provide a default library function implementation that is used if no replacement is provided.

`--no-vector-sections`

Do not generate vector table entries for interrupt functions. This is useful when writing interrupt functions that are going to be installed using some other mechanism.

`-c`

Generate an object file without linking. This is a no-operation, as the compiler always generates an object file. This option is provided due to its common use with many compilers.

`--force-switch`

Forces the compiler to use a specific strategy for generating switch tables. The compiler normally uses heuristics, but this option overrides them to force a particular variant.

Available variants are: series of “if-else” tests, a jump table, or a table with value and label pairs. “If-else” is often best for very small switches. A jump table is typically fastest for a decent number of close entries. Value and label pairs work well for larger tables with spread-out values, using a binary search lookup that scales fairly well.

In big O notation: “if-else” is  $O(n)$ , a jump table is  $O(1)$ , and value and label pairs are  $O(\log(n))$ .

`-MD`

Write a dependency file that contains user and system headers.

`--dependencies`

Like `-MD`, but also implies `-E` and writes to `stdout` by default.

`-MF<file>`

Write dependency file output from `-MMD`, `-MD`, `-MM`, or `-M` to `<file>`.

`-MG`

Add missing headers to the dependency file.

`-MJ<arg>`

Write a compilation database entry for each input.

`-MM`

Like `-MMD`, but also implies `-E` and writes to `stdout` by default.

`-MMD`

Write a dependency file containing user headers.

`-MP`

Create phony target for each dependency (other than main file).

`-MQ<arg>`

Specify name of main file output to quote in dependency file.

`-MT<arg>`

Specify name of main file output in dependency file.

`-MV`

Use NMake/Jom format for the dependency file.

`-W`

Control warnings, see [Controlling warnings](#).

`--pedantic-errors`

Generates errors when language extensions are used. The supported language includes relaxations to the strict C standard for increased flexibility; use this option to disable such extensions.

`--align-functions`

Aligns the start of all functions to the provided alignment.

`--core`

This option is mostly available for symmetry with other tool chains from Calypsi and future expansion. At the moment only the 65816 core is recognized.

### `--target`

This tells the compiler that the code is intended for a certain target platform.

When the C256 target is enabled it tells the compiler to generate code for the hardware math unit and enables hosted behavior.

For the F256 target is enables hosted behavior.

When the SNES target is enabled it tells the compiler to generate code for the hardware math unit as well as utilizing the PPU hardware multiplier. The use of the PPU multiplier can not be used when rendering in Mode 7. You can disable the use of the PPU multiplier using the `--no-ppu-mul` option..

### `--code-model`

This options tells the compiler which code model to use. See [Code model](#) for more information.

### `--data-model`

This options tells the compiler which data model to use. See [Data model](#) for more information.

This chapter describes the available data types in detail.

## 11.1 Basic types

All standard integer types are supported. The following table lists the built-in integer types and their ranges.

Table 11.1: Integer types

| Type name          | Size    | Range                     |
|--------------------|---------|---------------------------|
| bool               | 8 bits  | 0 to 1                    |
| char               | 8 bits  | 0 to 255                  |
| signed char        | 8 bits  | -128 to 127               |
| unsigned char      | 8 bits  | 0 to 255                  |
| short              | 16 bits | -32768 to 32767           |
| signed short       | 16 bits | -32768 to 32767           |
| unsigned short     | 16 bits | 0 to 65535                |
| int                | 16 bits | -32768 to 32767           |
| signed int         | 16 bits | -32768 to 32767           |
| unsigned int       | 16 bits | 0 to 65535                |
| long               | 32 bits | $-2^{31}$ to $2^{31} - 1$ |
| signed long        | 32 bits | $-2^{31}$ to $2^{31} - 1$ |
| unsigned long      | 32 bits | 0 to $2^{32} - 1$         |
| signed long long   | 64 bits | $-2^{63}$ to $2^{63} - 1$ |
| long long          | 64 bits | $-2^{63}$ to $2^{63} - 1$ |
| unsigned long long | 64 bits | 0 to $2^{64} - 1$         |

### 11.1.1 bool

To use the `bool` data type, include `stdbool.h`, which also defines `true` and `false`. The boolean data type is also available as `_Bool` without requiring `stdbool.h`.

### 11.1.2 char

The `char` type is unsigned by default. To enable signed `char`, compile with the `--char-is-signed` option. Note that the supplied C library uses unsigned `char`.

---

Note: The `char` type differs from `short`, `int` and `long` in that it defaults to being unsigned in this compiler. Standard C allows it to be either signed or unsigned. The rationale for making `char` unsigned is that it is meant to represent a character code and in encoding standards like ASCII, ISO-8859-1 and Unicode, character values are unsigned entities.

If you intend to use 8 bits data types in expressions, you should consider using `int8_t` or `uint8_t` instead. They are defined in `stdint.h`.

---

### 11.1.3 wchar\_t

The wide character type `wchar_t` is defined if you include `stddef.h`.

### 11.1.4 Bit fields

Bit fields are supported based on any integer type. A bit field value has the same type (and signedness) as the integer base type it is defined in and are subject to the usual conversion rules when used in expressions.

A bit field is allocated starting from the least significant available bitposition in its container. If there are not enough bits available in the container to represent the bit field, a new container is allocated.

Consider the following declaration:

```
struct bf {
    uint16_t a:7;
    int16_t b:5;
    uint32_t c:24;
    uint8_t d:4;
    uint8_t e:2;
    int8_t f:3;
};
```

The two bit fields `a` and `b` are allocated in the same 16 bits container. Bit field `c` gets a 32 bits container of its own. Finally, `d` and `e` will share the same 8 bits container while `f` is allocated in a separate 8 bits container as there is not room to store it together with `d` and `e`.

Bit fields should be used with care. They can be a convenient way to pack several small values into some structure when data space is limited. However, accessing bit fields is in general more costly in terms of produced code, compared to using normal integer types.

---

Note: Sometimes it can be tempting to try and map bit fields to hardware registers. This can work, but it makes the code more sensitive to using a particular compiler. It may also require some thinking to get it

right. The alternative way of accessing hardware registers in its intended access width and manually apply shift and mask operations is often more robust.

---

### 11.1.5 Floating-point types

Floating point values follows the IEEE 754 format and is supported in two different sizes.

Table 11.2: Floating-point types

| Type name   | Size    | Approximate range (normal values)                         |
|-------------|---------|-----------------------------------------------------------|
| float       | 32 bits | $\pm 1.18 \times 10^{-38}$ to $\pm 3.40 \times 10^{38}$   |
| double      | 32/64   | as float or long double                                   |
| long double | 64 bits | $\pm 2.23 \times 10^{-308}$ to $\pm 1.80 \times 10^{308}$ |

Floating point numbers are represented in binary floating point form. The size of double is 32 bits by default. This can be changed to 64 bits by using the command-line option `--64bit-doubles`. The runtime library comes both in variants compiled with double set to 32 bits as well as 64 bits.

Subnormal numbers, infinity and NaN (not a number) are supported. The ranges stated in the table are for normal floating point numbers. Subnormal floating-point numbers extends the range of the exponent further at the cost of gradual loss of precision in the mantissa.

Floating point exceptions and changing the rounding mode are not supported.

---

Note: When the C256 target is enabled (using `--target=C256` command-line option), the built-in hardware math module is used for 32-bit floating point operations. This math module does not support subnormal numbers; operations producing them will generate zero.

---

#### 32 bits format

In the 32 bits format the exponent is 8 bits and the mantissa is 23 bits. The precision is between 6 and 7 decimal digits.

#### 64 bits format

In the 64 bits format the exponent is 11 bits and the mantissa is 52 bits. The precision is between 15 and 16 decimal digits.

### 11.1.6 Function pointer types

The 65816 architecture supports 16 bits short function pointers (`nearfunc`) and 24 bits far function pointers (`farfunc`).

### 11.1.7 Data pointer types

The following data pointers are available:

Table 11.3: Data pointers

| Name or keyword                 | Size    | Index type  | Address range                    |
|---------------------------------|---------|-------------|----------------------------------|
| Direct page <code>__tiny</code> | 8 bits  | signed char | <code>0x00-0xff</code>           |
| Near bank <code>__near</code>   | 16 bits | signed int  | <code>0x0000-0xffff</code>       |
| Far memory <code>__far</code>   | 32 bits | signed int  | <code>0x000000-0xffffffff</code> |
| Far memory <code>__far24</code> | 24 bits | signed int  | <code>0x000000-0xffffffff</code> |
| Huge memory <code>__huge</code> | 32 bits | signed long | <code>0x000000-0xffffffff</code> |

Each data pointer has an associated index type, always a signed integer type. This type is used in address calculations, such as array access with an index or advancing a pointer by adding or subtracting an integer value.

### 11.1.8 Pointer conversions

Function and data pointers are treated as being unsigned values. In general, casting to a type that has fewer bits means a pointer value gets truncated. Casting to a wider type results in zero extension.

### 11.1.9 `size_t`

This unsigned integer type holds the maximum size of an object. On the 65816 the size of `size_t` is 16 bits.

### 11.1.10 `ptrdiff_t`

This signed integer type represents a distance within the largest possible object. On the 65816 the size of `ptrdiff_t` is 16 bits.

Subtracting two data pointers (within the same object) yields a `ptrdiff_t` value, representing the number of elements between the pointers, not the number of bytes.

---

Note: Standard C allows referring to an element one beyond the actual object. Subtracting the end address from the start address may result in a negative result if the value exceeds the range of `ptrdiff_t`.

---

## 11.2 Structure types

Structure types are fully supported and can be nested. Structure members are stored sequentially in the order they appear in the declaration.

## 11.3 Union types

Union types are fully supported. The compiler also supports a useful extension that allows anonymous unions within a structure. Consider the following code:

```
struct Scope {
    union {
        char alpha;
        int num;
    } u;
    int b;
};

int main() {
    struct Scope x;
    x.u.num = 65;
    x.u.alpha = 'A';
    return 0;
}
```

The use of a declarator on the union within the structure requires an extra step to access its members. This was relaxed in the C11 standard, allowing you to write:

```
struct Scope {
    // Anonymous union
    union {
        char alpha;
        int num;
    };
    int b;
};

int main() {
    struct Scope x;
    x.num = 65;
    x.alpha = 'A';
    return 0;
}
```

This C11 extension, enabled by default, allows anonymous unions to be compiled without diagnostic messages, even though the compiler actually supports C99.

You can enable warnings for such extensions using the `-Wc11-extensions` or `-Wpedantic` command-line options.

You can also disable such extensions with the command-line option `--pedantic-errors`. In that case you will get an error instead.

## 11.4 Enumeration types

Enumeration types are represented as `int`. To use a smaller storage representation, choose a smaller integral type. For a better name, use a `typedef`:

```
enum fruit { apple, orange, banana };

typedef char fruit_t;

fruit_t active;

void citrus(void)
{
    active = orange;
}
```

## 11.5 Type qualifiers

Standard C provides two type qualifiers: `volatile` and `const`.

### 11.5.1 Volatile objects

C has the concept of volatile objects, typically used for hardware access. Both writing and reading volatile objects are considered side effects that will occur.

Related is the concept of sequence points in a program. A volatile access between two sequence points occurs between those points and cannot be moved past a sequence point. To ensure ordered memory accesses, make them volatile and separate them with a sequence point. The semicolon after a statement is an example of a sequence point:

```
uint8_t volatile * mem1;
uint8_t volatile * mem2;

void foo () {
    *mem1 = 2;
    *mem2;
}
```

In this case the write to `mem1` is guaranteed to be performed before the read of `mem2`. The read of `mem2` will also occur even if the result of the read is not used.

If multiple volatile accesses are done between two sequence points the order they happen in is undefined:

```
uint8_t volatile * mem1;
uint8_t volatile * mem2;

int foo () {
    return *mem1 + *mem2;
}
```

In this case both `mem1` and `mem2` are read, but the order in which they are read is undefined.

## Access size

Accessing a volatile object wider than the natural register size on the target results in access performed in several steps, as dictated by the natural register size.

Reading and using only a portion of a scalar volatile object still results in the entire object being read:

```
volatile uint64_t wide;

uint16_t foo () {
    return wide;
}
```

Here the volatile object is 64 bits, but we are only interested in the lower 16 bits. In this case all 64 bits are read, the upper 48 bits are then discarded and the function returns the lower 16 bits.

If `wide` was not volatile, the compiler may instead choose to only read the lower 16 bits of the 64-bit `wide` variable.

## Assignment results

Assignments in C has an expression value. Consider:

```
volatile int var;

int foo () {
    return var = 4;
}
```

The assignment writes 4 to the volatile variable, but the function return value (either 4 or the value read from `var` after the assignment) is implementation-defined.

Avoid such constructs in your programs. Be more explicit about your intent. To force a read after the assignment:

```
volatile int var;

int foo () {
    var = 4;
    return var;
}
```

If you want to be sure the function returns 4:

```
volatile int var;

int foo () {
    var = 4;
    return 4;
}
```

This clarifies the intent and ensures consistent behavior across C compilers.

## Bit fields

Accessing volatile bit fields has undefined behavior. A bit field describes a subset of bits within its storage unit. Adjacent bits may be accessed depending on layout. Using volatile bit fields is discouraged.

---

Note: Rather than using bit fields, define normal scalar values so that they cover hardware registers, following the defined or intended size of hardware register access. Then use expressions to extract or manipulate the part of the register you want.

---

## 11.5.2 Const objects

The `const` type qualifier indicates a read-only object. It can be applied to data objects and pointers.

When a static object is defined as `const`, the compiler attempts to place it in a read-only memory section.

A pointer to a `const` object can point to both read-only and writable objects. Such a pointer signifies that its user should not attempt to alter memory. This can aid the optimizer and is considered good practice, as it prevents unintended data alteration by parts of the application.

## 11.6 Type definitions

Using type definitions (the `typedef` keyword) is highly recommended, offering several benefits:

1. Code becomes more readable with meaningful type names. For instance, `speed_t` is clearer than `long`. Should the type definition change, only one location requires modification, preventing the need to search and selectively update `long` instances that represent speed.
2. A `struct fish` can be concisely named `fish_t`, which also hides its structure, if desired.

`typedef` incurs no cost; the generated code remains identical.

## 11.7 Alignment

The 65816 imposes no data alignment. Data objects can start at any address, and no padding is introduced between structure elements.

An attribute is a property attachable to functions, data objects, or types, specified as a keyword. Standard C includes built-in keywords like `const` and `volatile`.

Extended attributes provide access to behaviors or properties beyond Standard C. They are either target-specific or useful for embedded systems.

### 12.1 Overview

Attributes can be applied with either keyword syntax (e.g., `__tiny`) or attribute syntax (e.g., `__attribute__((tiny))`). Both are functionally equivalent, but the C parser may not accept the keyword form in some situations.

Usage is largely a matter of preference. Preprocessor macros can rename attributes for portability, enabling them to be toggled off for other targets or renamed to match different compilers or targets.

---

**Note:** The C parser sometimes produces unexpected errors with the keyword form of attributes (e.g., `__tiny`). If this occurs, use the `__attribute__((tiny))` form instead.

---

### 12.2 Using attributes

Type attributes can be applied to type declarations, following the same syntax as type qualifiers like `const` and `volatile`.

### 12.2.1 Syntax for data objects

You can apply attributes to data objects as follows:

```
__attribute__((tiny)) int a, b;  
int __tiny c, d;
```

When applied to an object, the attribute's location is irrelevant. The example above applies the Direct page attribute to all defined objects (a, b, c, and d).

### 12.2.2 Syntax for pointer types

Attributes can also be applied to pointer types, where their location is significant. This determines whether the pointer itself is constant or what it points to is constant.

The easiest way to decipher attributes in function types is to read the type from right to left.

```
int __attribute__((tiny)) * p1;  
long * __attribute__((tiny)) p2;
```

Here, p1 is a pointer stored in default memory that points to an int in Direct page memory. p2 is a pointer stored in Direct page memory that points to a long in default memory.

## 12.3 Attribute reference

This section goes through all available extension keywords and attributes.

### 12.3.1 Summary of attributes

The following table summarizes available attributes. For the keyword form, prefix with two underscores (e.g., \_\_tiny). For attribute syntax, use the attribute name with \_\_attribute\_\_ (e.g., \_\_attribute\_\_((tiny))).

Table 12.1: Extended attributes summary

| Attribute name  | Description                                                                  |
|-----------------|------------------------------------------------------------------------------|
| aligned(nn)     | Specify alignment of data object or function                                 |
| section("name") | Specify section name to use for a data object or function                    |
| tiny            | Control storage of data object to direct page                                |
| near            | Control storage of data object to near bank                                  |
| far             | Control storage of data object to far memory                                 |
| huge            | Control storage of data object to huge memory                                |
| interrupt       | Used to define an interrupt function                                         |
| intrinsic       | Used to declare an intrinsic function                                        |
| saveds          | Entry point that needs to initialize the direct page and data bank registers |
| task            | Relaxes preserving registers                                                 |

### 12.3.2 Description of attributes

This section describes each attribute in detail.

#### aligned

This attribute can be applied to functions, global and static data objects to force a certain minimal alignment. The `aligned` attribute takes an argument which is the alignment to use:

```
__attribute__((aligned(16))) struct sprite ship;
```

Note: Certain data types may impose an alignment by themselves. The actual alignment is chosen so that all alignment constraints are satisfied.

#### section

This attribute can be applied to functions, global and static data objects to control the name of the section it is placed in.

The `section` attribute takes an argument which is the section name to use:

```
// Place in vram
__attribute__((section("vram")))
const char tiles[256] = { .. };

__attribute__((section("trueCode")))
long foo () {
    return 42;
}
```

See [Description of pragma directives](#) for how you can specify a section for multiple functions, global and static data objects.

#### tiny

This specifies a data object or a pointer to a data object that resides in the direct page. Direct page addressing results in shorter and faster code. It is especially efficient when working with pointers on the 65816.

Note: Auto variables are normally allocated to registers or direct page pseudo registers which means you will automatically get the benefits of direct page addressing when using auto variables. The `tiny` attribute is best used for smaller frequently used static data objects.

### near

This specifies a data object or a pointer to a data object that resides in a single 64K bank. Near bank addressing is a middle ground that is fairly efficient, but not as efficient as using the direct page. It results in shorter code than the `far` and `huge` attributes and being 64K it provides ample space.

### far

This specifies a data object or a pointer to a data object that can reside anywhere in memory. The only limitation is that a single data object can be at most 64K minus one byte large.

The code needed to access far memory tends to be somewhat larger than the near bank, but the 65816 is still capable of using various addressing modes for reasonable good performance.

### far24

This specifies a data object or a pointer to a data object that can reside anywhere in memory. The only limitation is that a single data object can be at most 64K minus one byte large.

A `far24` pointer is similar to `far` with the difference that it occupies 24 bits in memory.

---

Note: The `far24` attribute is intended to be used when writing to data structures where you actually want a write to be exactly 24 bits. This may be required by hardware or function API. The compiler will normally allocate 32 bits for wide pointers such as `far` and `huge` in order to make more efficient code on the 65816.

---

### huge

This specifies a data object or a pointer to a data object that can reside anywhere in memory and that has no size limitations other than the addressable 16MB range.

The generated code for using this attribute tends to be larger than any of the other alternatives. At best you can get something that is the same or similar as `far`, but in many cases the code will be larger and slower.

### interrupt

An interrupt function is meant to serve as an interrupt handler. The interrupt has the following effects:

1. An interrupt function cannot take any parameters
2. An interrupt function will preserve all registers used
3. Leaving the interrupt function uses a different instruction sequence compared to normal functions
4. The interrupt attribute may optionally be given a vector address as an argument

The interrupt vector specified is given as an argument to the interrupt attribute:

```
int counter;

__attribute__((interrupt(0xffff)))
void irq () {
    counter++;
}
```

---

Note: The vector argument is optional. Omitting it means that there will be no vector section entry generated for that interrupt function. You can also suppress all vector sections from being generated by using the `--no-vector-sections` command-line option.

---

---

Note: It is strongly recommended to make your interrupt functions small and simple. Also avoid making function calls from an interrupt function. A function call means the interrupt needs to preserve scratch registers, which is costly. If you must use function calls, consider forcing them to be inline expanded See [Function inlining](#) for more information.

---

### `intrinsic`

The `intrinsic` attribute is used to declare intrinsic built-in functions. This can only be done on intrinsic functions that is already known to the compiler. Normally you use this by including the `calypsi/intrinsics65816.h` file which contains all such valid declarations.

### `saveds`

A `saveds` function is useful for API functions called from another context. The direct page and data bank are initialized to suit the current runtime environment. The previous values of these registers are restored when the function returns. This makes it possible to call an API function from another context where different values may be used for these base registers.

---

Note: For `saveds` API calls, you should generally prevent arguments from being passed on the direct page. Use either a single register argument or the `simple_call` calling convention to force arguments onto the stack. Direct page differences when using `saveds` mean accessing caller-placed arguments there will not work.

---

---

Note: When using `saveds` for API functions you most likely want to use the `saveds` attribute on C written interrupt functions as well, as interrupts may be triggered at any time.

---

### `task`

A `task` attribute can be used on functions such as `main` which is the start of the application. You will not normally call such functions from any C code. In that case you can apply the `task` attribute which relaxes preserving registers that would otherwise be saved on the stack. This can save a little stack space and will make the application a tiny bit smaller.



The Calypsi C compiler supports several C language extensions that can be convenient. However, consider their portability implications.

## 13.1 Overloaded functions

The overloaded functions support provides overloading of functions in C similar to C++. Overloading in C is introduced using the `overloadable` attribute. For example, you might provide several overloaded versions of a `sine` function that invokes the appropriate standard function computing the sine of a value with `float`, `double`, or `long double` precision:

```
#include <math.h>

float __attribute__((overloadable)) sine(float x) { return sinf(x); }
double __attribute__((overloadable)) sine(double x) { return sin(x); }
long double __attribute__((overloadable)) sine(long double x) { return sinl(x); }
```

The compiler calls the most suitable function based on argument types. Overloaded functions are name mangled as C uses a single global namespace.

## 13.2 Statement expressions

This GCC extension allows a statement to appear where an expression is expected. Use it to allow loops, switches, and local variables within an expression.

A compound statement is a sequence of statements enclosed by braces. In the example below, parentheses around the braces create a statement expression:

```
({ int y = foo (); int z;
   if (y > 0) z = y;
```

(continues on next page)

(continued from previous page)

```
else z = - y;  
z; })
```

The last statement in a statement expression is an expression followed by a semicolon. In this case, the variable `z` is used as the result of the entire statement expression.

### 13.2.1 Safe macros

Statement expressions can also implement safe macros that evaluate their parameters only once.

Recall the well known `max` macro:

```
#define max(a,b) ((a) > (b) ? (a) : (b))
```

When used, this expands the largest value twice, which may be undesirable if it has side effects or represents an expensive computation.

With a statement expression this can be expressed as:

```
#define maxint(a,b) \  
    ({int _a = (a), _b = (b); _a > _b ? _a : _b; })
```

While this solves the double-evaluation issue, it introduces other subtle problems. First, you need to specify the type (see also `__auto_type` extension below), and it may cause variable shadowing.

In most cases, inline functions offer a better alternative to statement expressions. Inline functions avoid subtle variable shadowing problems, result in more readable code, and generate equally efficient code when inlined. The minor caveat is that inline functions may not be inlined, while statement expressions are always expanded in place.

## 13.3 Auto type

The `auto` type extension allows specifying a type that is automatically selected, using the `__auto_type` keyword. The declaration must declare only one variable, whose declarator must be an identifier. The declaration must be initialized, and the variable type is determined by the initializer type.

Using `__auto_type`, the “max” macro in the previous section can be written to select the type automatically:

```
#define max(a,b) \  
    ({__auto_type _a = (a); \  
      __auto_type _b = (b); \  
      _a > _b ? _a : _b; })
```

## 13.4 Typedef

Refer to the type of an expression using `__typeof`. The syntax resembles `sizeof`, but semantically, it acts like a typename defined with `typedef`.

Similar to `sizeof`, there are two ways to write the argument to `__typeof`: with an expression or with a type. For example, with an expression:

```
__typeof (x)
```

This creates a type identical to the identifier `x` in the current context.

You can also use a typename as an argument:

```
__typeof (int *)
```

In this case, the type is simply a pointer to `int`.

`Typedef` can also be used to implement the “max” macro:

```
#define max(a,b) \
    ({__typeof (a) _a = (a); \
      __typeof (b) _b = (b); \
      _a > _b ? _a : _b; })
```

## 13.5 Generics

C11 style `_Generic` is now supported. This allows expressions to expand to different outcomes based on a controlling expression. For example:

```
extern void stringFunc(char*);
extern void otherStringFunc(char*);
extern void string4Func(char[4]);

#define F(X) _Generic(&(X), \
    default: otherStringFunc, \
    char*: stringFunc, \
    char(*)[4]: string4Func, \
    char const*: stringFunc, \
    char const(*)[4]: string4Func \
)(X)

void foo(char *p) {
    F(p);
    F("foo");
    F("longer string");
}
```

The macro `F` calls different functions based on its input. For example, a string literal “foo” (3 characters plus null terminator) is treated as a `char` array of 4 bytes, leading to a call to `string4Func()`. The `p` parameter, an ordinary `char*`, results in a call to `stringFunc()`. Finally, “longer string” is treated as a `char` array with a length other than four, matching the `default:` alternative and resulting in a call to `otherStringFunc()`.



## Pragma directives

Pragma directives are a standard C mechanism allowing vendor-specific extensions to the C language while maintaining portability.

You can use pragma directives to control compiler behavior, such as suppressing warnings or placing objects in specific memory areas using custom sections.

Pragma directives are always enabled and can be used with either the `#pragma` preprocessor directive or the `_Pragma` preprocessor operator.

### 14.1 Pragma directives reference

This section details all available pragma directives.

#### 14.1.1 Summary of pragmas

The following table summarizes the recognized pragma directives:

| Pragma directive                      | Description                                      |
|---------------------------------------|--------------------------------------------------|
| <code>#pragma clang section</code>    | Control custom section for code or data          |
| <code>#pragma message</code>          | Generate custom warning                          |
| <code>#pragma GCC warning</code>      | Generate custom warning                          |
| <code>#pragma GCC error</code>        | Generate custom error                            |
| <code>#pragma clang diagnostic</code> | Control diagnostic messages inside a source file |
| <code>#pragma require</code>          | Force inclusion of other module                  |
| <code>#pragma rtattr</code>           | Define a runtime attribute                       |

Note: Since Clang is the C front end, some pragmas are inherited and function similarly. This simplifies compiler switching, despite pragmas being vendor-specific extensions.

### 14.1.2 Description of pragma directives

#### section

The `#pragma clang section` directive assigns section names to functions, global, and static objects.

The compiler typically places global and static objects into predefined sections, but you can override this to control placement to a specific memory area.

The section names can be specified as:

```
#pragma clang section bss="myBSS" data="myData" rodata="myRodata" text="myText"
```

The section names can be reverted back to default name by supplying an empty string to the section kind, for example:

```
#pragma clang section bss="" data="" text="" rodata=""
```

The new section name applies to all functions, global and static objects that follow from the pragma directive.

You are not required to define a name for every section category. If you omit some, the most recently specified one is used.

The `section` attribute takes precedence over the `clang section` pragma directive; a section name specified with `__attribute__((section("myname")))` has precedence.

---

Note: Section names are not interpreted. For example, naming a section `.bss.mySec` does not mean it will be a BSS section name.

---

#### message

You can generate a custom warning with `message` pragma. In addition there are a couple of GCC variants of this supported:

```
// The following will generate warning messages
#pragma message "my own diagnostic message"
#pragma GCC warning "my own diagnostic message"

// This will give and error
#pragma GCC error "not supported"
```

#### diagnostics

You can control diagnostic messages in source code using pragmas, which is useful for temporarily disabling specific warnings in a section of code.

The pragma controls any command-line-configurable warning. Warnings can be set to ignored, warning, error, or fatal.

You can also push and pop the current warning state. This is useful when writing a header file that others will compile, as their warning flags are unknown.

In the example below, `-Wextra-tokens` is ignored for a few lines; diagnostics then revert to their previous state.

```
#if foo
#endif foo // warning: extra tokens at end of #endif directive

#pragma clang diagnostic push
#pragma clang diagnostic ignored "-Wextra-tokens"

#if foo
#endif foo // no warning

#pragma clang diagnostic pop
```

The push and pop pragmas will save and restore the full diagnostic state, regardless of how it was set.

### require

The `#pragma require` directive ensures that a module is included at the link stage, even if no functions or data objects within it are called or referenced.

This is primarily for libraries that request code slices or objects. For example, the C library's heap system uses a `require` pragma in its `malloc()` module to enable heap initialization during system startup:

```
#pragma require __call_heap_initialize
```

### rtattr

The `#pragma rtattr` directive defines a runtime attribute within a C source file.

This is primarily for building libraries with alternative modules. While runtime attributes can be defined using the `--rtattr` command-line option, defining them in the source file may simplify build rules.

```
#pragma rtattr myAttribute "someValue"
```



This chapter covers predefined intrinsic and built-in functions.

Intrinsics functions appear as ordinary calls but are special compiler constructs, emitting specific instruction sequences. To enable them, include the `calypsi/intrinsics65816.h` header file.

Built-in functions are similar, but they are not declared in `calypsi/intrinsics65816.h`.

## 15.1 Intrinsic functions reference

The `calypsi/intrinsics65816.h` file appears as follows:

```
/******  
 *  
 * Copyright Håkan Thörnngren  
 *  
 * This file is part of the Calypsi C library.  
 * Permission to use with the Calypsi tool chain is hereby granted.  
 *  
 *****/  
  
#ifndef __INTRINSICS_65816_H  
#define __INTRINSICS_65816_H  
  
#ifdef __CALYPSI_TARGET_65816__  
  
typedef unsigned char __interrupt_state_t;  
typedef void (*__return_address_t)(void);  
typedef unsigned short __return_address16_t;  
  
__attribute__((intrinsic)) void __disable_interrupts(void);  
__attribute__((intrinsic)) void __enable_interrupts(void);  
__attribute__((intrinsic)) void __break_instruction(void);  
__attribute__((intrinsic)) void __break_with(unsigned char);
```

(continues on next page)

(continued from previous page)

```

__attribute__((intrinsic)) void __coprocessor_with(unsigned char);
__attribute__((intrinsic)) void __no_operation(void);
__attribute__((intrinsic)) __interrupt_state_t __get_interrupt_state(void);
__attribute__((intrinsic)) void __restore_interrupt_state(__interrupt_state_t);

__attribute__((intrinsic)) __return_address_t __get_return_address(void);
__attribute__((intrinsic)) void __set_return_address(__return_address_t);
__attribute__((intrinsic)) void __set_return_address_low16(__return_address16_t);

__attribute__((intrinsic)) void __wait_for_interrupt(void);
__attribute__((intrinsic)) void __stop_processor(void);

#endif // __CALYPSI_TARGET_65816__

#endif // __INTRINSICS_65816_H

```

### 15.1.1 Summary of intrinsic functions

Table 15.1: Intrinsic functions summary

| Function name                          | Description                                    |
|----------------------------------------|------------------------------------------------|
| <code>__disable_interrupts</code>      | Disable interrupts                             |
| <code>__enable_interrupts</code>       | Enable interrupts                              |
| <code>__get_interrupt_state</code>     | Get the current interrupt state                |
| <code>__restore_interrupt_state</code> | Restore to a previous interrupt state          |
| <code>__get_return_address</code>      | Get the return address of the current function |
| <code>__set_return_address</code>      | Alter the return address of current function   |
| <code>__break_instruction</code>       | Generate a BRK instruction                     |
| <code>__break_with</code>              | Generate a BRK instruction with an argument    |
| <code>__coprocessor_with</code>        | Generate a COP instruction with an argument    |
| <code>__no_operation</code>            | Generate a NOP instruction                     |
| <code>__wait_for_interrupt</code>      | Wait for an interrupt                          |
| <code>__stop_processor</code>          | Stop the processor                             |
| <code>__builtin_clz</code>             | Count leading zeroes in an int                 |
| <code>__builtin_clzl</code>            | Count leading zeroes in a long                 |
| <code>__builtin_clzll</code>           | Count leading zeroes in a long long            |
| <code>__builtin_signbit</code>         | Get the value of the sign bit                  |

### 15.1.2 Description of intrinsic functions

This section details each intrinsic function.

`__disable_interrupts`

Emits a machine instruction that disables interrupts.

`__enable_interrupts`

Emits a machine instruction that enables interrupts.

`__get_interrupt_state`

Returns the current interrupt state as a value of type `__interrupt_state_t`. This can be used in the following way:

```
int global;

void safe_increment () {
    __interrupt_state_t state = __get_interrupt_state();
    __disable_interrupts;
    global++;
    __restore_interrupt_state(state);
}
```

In the example the current state of interrupts is obtained and saved in `state`. The code then ensures interrupts are disabled and then do the real work. Finally, the interrupt state is restored to what it was when the function was entered.

`__restore_interrupt_state`

Restores to a previous interrupt state, see example above.

`__get_return_address`

Reads the return address of the current function from its the stack frame.

`__set_return_address`

Replaces the return address of the current function with the address found in the argument of `__set_return_address()`.

`__no_operation`

Emits a NOP machine instruction.

`__break_instruction`

Emits a machine instruction that causes a break exception.

`__break_with`

Emits a machine instruction that causes a break exception. This variant take an opcode argument for the second byte of the BRK instruction.

`__coprocessor_with`

Emits a COP machine instruction with its argument is the second byte of the COP instruction.

`__wait_for_interrupt`

Emits an instruction that causes the 65816 to wait until an interrupt is triggered. Execution resumes with the interrupt and then normal execution.

`__stop_processor`

Emits an instruction that causes the 65816 to halt execution.

`__builtin_clz`

Count the number of leading zeroes in a provided `int` value. Returns an undefined result if input is zero.

`__builtin_clzl`

Count the number of leading zeroes in a provided `long` value. Returns an undefined result if input is zero.

`__builtin_clzll`

Count the number of leading zeroes in a provided `long long` value. Returns an undefined result if input is zero.

`__builtin_signbit`

Returns the value (0 or 1) of the signbit of its input.



This chapter briefly describes the preprocessor and covers commonly used macros.

## 16.1 Overview

The preprocessor adheres to Standard C, providing:

1. Predefined macros for inspecting the compilation environment, allowing application tuning based on compiler properties.
2. User-defined macros, specified via the command line or within C source files.
3. The ability to dump the preprocessor environment for inspection.
4. The ability to dump the source file after preprocessing.

### 16.1.1 Predefined macros

The following describes the macros which are the ones you are most likely to use in conditional builds.

`__STDC__`

This macro is set to 1, indicating that this compiler adheres to Standard C.

### `__STDC_HOSTED__`

This macro is defined as 0, indicating a cross-compiler. You can test this as follows:

```
#if defined(__STDC_HOSTED__) && __STDC_HOSTED__ == 0
```

### `__STDC_VERSION__`

This macro is set to 199901L to indicate that this compiler adheres to the ISO/IEC 9899:1999 standard.

### `__CALYPSI__`

This macro is set to 1, indicating that either the assembler or compiler in use is from Calypsi.

### `__CALYPSI_CC__`

This is set to 1, indicating that this is a Calypsi C compiler.

### `__CALYPSI_ASSEMBLER__`

This macro is set to 1 when a Calypsi assembler is used.

### `__CALYPSI_DEBUG__`

This macro is defined and set to 1 when compiling with debug information enabled, i.e. `--debug` or `-g` command-line option used.

### `__CALYPSI_VERSION_MAJOR__`

This macro is set to the first number of the version.

### `__CALYPSI_VERSION_MINOR__`

This macro is set to the second number of the version.

### `__BIG_ENDIAN__`

This macro is defined if the target has big-endian byte order.

## `__LITTLE_ENDIAN__`

This macro is defined if the target has little-endian byte order.

## `__BYTE_ORDER__`

This macro defines the byte order, taking the value of `__ORDER_LITTLE_ENDIAN__`, `__ORDER_BIG_ENDIAN__`, or `__ORDER_PDP_ENDIAN__`.

It is usually simpler to use `__BIG_ENDIAN__` or `__LITTLE_ENDIAN__` macros for conditional inclusion of endian-dependent source code.

## `__ORDER_LITTLE_ENDIAN__`

This macro is defined to 1234 to describe the byte order in memory for a little-endian target.

## `__ORDER_BIG_ENDIAN__`

This macro is defined as 4321, describing the byte order in memory for a big-endian target.

## `__ORDER_PDP_ENDIAN__`

This macro is defined as 3412, describing the byte order in memory for a PDP-endian target.

## `__FILE__`

This macro expands to the name of the current source file.

## `__LINE__`

This macro expands to the current line number in the source file.

## `__DATE__`

This macro expands to the current date.

## `__TIME__`

This macro expands to the current time of the day.

`__COUNTER__`

This macro expands to a number that increments every time the macro is used.

`__CALYPSI_TARGET_65816__`

Set to 1 when the target is the 65816 family.

`__CALYPSI_CORE_65816__`

Set to 1 when the 65816 core is selected. You can test for this core using `#ifdef __CALYPSI_CORE_65816__`.

`__CALYPSI_TARGET_SYSTEM_EMBEDDED__`

Set to 1 when the compiler is configured for an embedded system, generating code suitable for flash or ROM memory. You can test if the compiler is configured for an embedded system using `#ifdef __CALYPSI_TARGET_SYSTEM_EMBEDDED__`.

`__CALYPSI_TARGET_SYSTEM_C256__`

Set to 1 when the compiler is configured for the Foenix C256.

`__CALYPSI_TARGET_SYSTEM_F256__`

Set to 1 when the compiler is configured for the Foenix F256.

`__CALYPSI_TARGET_SYSTEM_SNES__`

Set to 1 when the compiler is configured for the Super Nintendo Entertainment System.

`__CALYPSI_CODE_MODEL_SMALL__`

Set to 1 if the compiler uses the Small code model.

`__CALYPSI_CODE_MODEL_COMPACT__`

Set to 1 if the compiler uses the Compact code model.

`__CALYPSI_CODE_MODEL_LARGE__`

Set to 1 if the compiler uses the Large code model.

`__CALYPSI_DATA_MODEL_SMALL__`

Set to 1 if the compiler uses the Small data model.

`__CALYPSI_DATA_MODEL_MEDIUM__`

Set to 1 if the compiler uses the Medium data model.

`__CALYPSI_DATA_MODEL_LARGE__`

Set to 1 if the compiler uses the Large data model.

`__CALYPSI_DATA_MODEL_HUGE__`

Set to 1 if the compiler uses the Huge data model.

`__CALYPSI_HUGE_ATTRIBUTE_ENABLED__`

Set to 1 if the Huge attribute is enabled via `--enable-huge-attribute`, which also makes `size_t` 32 bits.



# CHAPTER 17

---

## Section reference

---

Compiler-generated code and data objects are organized into sections within the object file. Understanding these sections is crucial for writing linker placement files and interpreting compiler and linker list files.

### 17.1 Section overview

Code and data objects occupy space in target memory, either as program-bits (actual data written to memory) or no-bits (sections represented only by their size in the ELF object file).

The Calypsi C compiler tool chain has adapted names that trace back to the origins to the early days of computing. These may not always be the best possible names, but they are widely known and familiar to many.

Sections are broadly categorized into three groups: text, data, and bss. Text sections are read-only program-bits. Data sections are read-write program-bits. BSS sections are read-write no-bits.

---

**Note:** If you ever wondered about what bss section stands for, it originally is block started by symbol, a pseudo operation in an assembler for IBM 704 from the mid fifties. Some people suggest it is easier to remember as better save space, as it is a way to save space in the output file by just storing the size of the section without any explicit data bits.

---

### 17.1.1 Section names

Each section has a name, and multiple section fragments can share the same name. Section names serve two purposes:

- To provide a descriptive identifier for the section.
- To control memory placement through rules that reference these names.

### 17.1.2 Section types

Each section has a type: text, data, or bss. For embedded systems, rodata (read-only data) and no-init sections are also available. Read-only data sections are similar to text sections but identify data rather than executable code. No-init sections resemble bss but lack the memory-clearing mechanism.

### 17.1.3 Data sections

From the compiler's perspective, data sections are read-write program-bits sections, meaning they have initializer values that reside in memory. In a hosted environment, such sections are loaded into memory before program execution. This approach is unsuitable for programs stored in flash memory where execution is expected to begin immediately upon power-on.

To address this, the linker clones data sections. The initializers for the original data section are moved and placed in the clone. The clone has a section name with an "i" prepended (e.g., `idata` for `data`) and is intended for flash (read-only) memory.

To initialize data sections on power-on, a copy routine is inserted before the `main()` function. This routine copies initializers from flash to RAM, initializing data memory and clearing bss sections.

The linker generates a `data-init_table` describing data to be copied and cleared.

---

Note: For some targets, the Calypsi C compiler tool chain can cross-compile executables for OS loading. The `--hosted` option prevents data section cloning, avoiding duplication. However, the table-driven initializer is still needed for BSS sections in such environments.

---

## 17.2 Sections used by the compiler

The following sections are sections used by the Calypsi C compiler tool chain.

Table 17.1: Sections

| Section name    | Type   | Memory kind | Description                                                         |
|-----------------|--------|-------------|---------------------------------------------------------------------|
| code            | text   | ROM         | Executable code                                                     |
| compactcode     | text   | ROM         | Executable code                                                     |
| farcode         | text   | ROM         | Executable code                                                     |
| ztiny           | bss    | RAM         | Zero initialized (bss) direct page data                             |
| tiny            | data   | RAM         | Initialized direct page data                                        |
| zdata           | bss    | RAM         | Zero initialized (bss) data, default memory, Small data model only  |
| data            | data   | RAM         | Initialized near data, default memory, Small data model only        |
| cdata           | rodata | ROM         | Constant near data, default memory, Small data model only           |
| znear           | bss    | RAM         | Zero initialized (bss) near data, not available in Small data model |
| near            | data   | RAM         | Initialized near data, not available in Small data model            |
| cnear           | rodata | ROM         | Constant near data, not available in Small data model               |
| zfar            | bss    | RAM         | Zero initialized (bss) far data                                     |
| far             | data   | RAM         | Initialized far data                                                |
| cfar            | rodata | ROM         | Constant far data                                                   |
| zhuge           | bss    | RAM         | Zero initialized (bss) huge data                                    |
| huge            | data   | RAM         | Initialized huge data                                               |
| chuge           | rodata | ROM         | Constant huge data                                                  |
| switch          | rodata | ROM         | Switch tables                                                       |
| itiny           | rodata | ROM         | Tiny data initializers                                              |
| idata           | rodata | ROM         | Data initializers, default memory, Small data model only            |
| inear           | rodata | ROM         | Near data initializers, not available in Small data model           |
| ifar            | rodata | ROM         | Far data initializers                                               |
| ihuge           | rodata | ROM         | huge data initializers                                              |
| data_init_table | rodata | ROM         | Data initializer table                                              |
| registers       | noinit | RAM         | Pseudo registers in direct page                                     |
| reset           | text   | ROM         | Reset vector, when used                                             |
| heap            | noinit | RAM         | Heap memory, for <code>malloc()</code>                              |
| stack           | noinit | RAM         | CPU stack, anywhere in bank 0                                       |

The sections `itiny`, `inear`, `ifar`, `ihuge` and `data_init_table` in the table above are linker generated.

Note: It is assumed that there is no ROM in direct page. Constants placed in `tiny` address space are handled in the same way as non-constant data objects in the `tiny` address space. The only difference is that they have the `const` attribute in the C type system.

Note: The table assigns sections to ROM and RAM for ROM-based applications that start on power-up. For hosted systems loading from storage to RAM, this distinction is not applicable. Always consider ROM-marked sections as read-only.

### 17.2.1 The vector section

An interrupt function will have an associated vector. This is a specially named section for the purpose of holding a single vector. The name looks something like `$$interruptVector_0xfffe`. The intended address of the vector is encoded in the section name and the linker recognizes these and will place the vector at the address specified. This is handled without the help of any section placement rules.

### 17.2.2 Section reference

The following goes through the available sections in detail.

#### `code`

Holds program code, address range `0x0000-0xffff`. This section is intended to be placed in a flash or ROM memory.

#### `compactcode`

Holds program code, address range `0xNN0000-0xNNffff`. Where NN is some bank, typically not `00`. This section is intended to be placed in a flash or ROM memory.

#### `farcode`

Holds program code, address range `0x000000-0xffffffff`. This section is intended to be placed in a flash or ROM memory.

#### `zdata`

Holds zero initialized (bss) data in the main memory, address range `0x0000-0xffff` and is used in the Small data model only.

In the other data models the `znear` section is used instead.

#### `data`

Holds non-zero initialized data intended to go into RAM in the main memory, address range `0x0000-0xffff`.

In the other data models the `near` section is used instead.

#### `idata`

Holds initializers for the `data` section. This section is intended to be placed in a flash or ROM memory and is copied to the `data` section before the application starts.

When running in a hosted system this section is not used. In that case the initializer values are part of the `data` section and data is initialized by loading the application.

In the other data models the `inear` section is used instead.

### `cdata`

Holds constant data in the main memory, address range `0x0000-0xffff`. This section is intended to be placed in a flash or ROM memory.

In the other data models the `cnear` section is used instead.

### `znear`

Holds zero initialized (bss) data in the main memory, address range `0xNN0000-0xNNffff` in one selected 64K bank page where NN is the bank page.

This section is not used in the Small data model, see `zdata`.

### `near`

Holds non-zero initialized data in the main memory, address range `0xNN0000-0xNNffff` in one selected 64K bank page where NN is the bank page.

This section is not used in the Small data model, see `data`.

### `inear`

Holds initializers for the `near` section. This section is intended to be placed in a flash or ROM memory and is copied to the `near` section before the application starts.

When running in a hosted system this section is not used. In that case the initializer values are part of the `near` section and data is initialized by loading the application.

This section is not used in the Small data model, see `idata`.

### `cnear`

Holds constant data in the main memory, address range `0xNN0000-0xNNffff` in one selected 64K bank page where NN is the bank page.

This section is not used in the Small data model, see `cdata`.

### `ztiny`

Holds zero initialized (bss) data in the main memory, address range `0x00-0xff`. This 256 bytes is located somewhere in the first 64K of memory.

### `tiny`

Holds non-zero initialized data intended to go into RAM in the direct page memory, address range `0x00-0xff`. This 256 bytes is located somewhere in the first 64K of memory.

#### `itiny`

Holds initializers for the `tiny` section. This section is intended to be placed in a flash or ROM memory and is copied to the `tiny` section before the application starts.

When running in a hosted system this section is not used. In that case the initializer values are part of the `tiny` section and data is initialized by loading the application.

#### `zfar`

Holds zero initialized (bss) data in the main memory, address range `0x000000-0xffffffff`.

#### `far`

Holds non-zero initialized data intended to go into RAM in the main memory, address range `0x000000-0xffffffff`.

#### `ifar`

Holds initializers for the `far` section. This section is intended to be placed in a flash or ROM memory and is copied to the `far` section before the application starts.

When running in a hosted system this section is not used. In that case the initializer values are part of the `far` section and data is initialized by loading the application.

#### `cfar`

Holds initialized constant data in the main memory, address range `0x000000-0xffffffff`. This section is intended to be placed in a flash or ROM memory.

#### `zhuge`

Holds zero initialized (bss) data in the main memory, address range `0x000000-0xffffffff`.

#### `huge`

Holds non-zero initialized data intended to go into RAM in the main memory, address range `0x000000-0xffffffff`.

#### `ihuge`

Holds initializers for the `huge` section. This section is intended to be placed in a flash or ROM memory and is copied to the `huge` section before the application starts.

When running in a hosted system this section is not used. In that case the initializer values are part of the `huge` section and data is initialized by loading the application.

### `chuge`

Holds initialized constant data in the main memory, address range `0x000000-0xffffffff`. This section is intended to be placed in a flash or ROM memory.

### `switch`

Holds switch tables in the main memory, address range `0x0000-0xffff`. This section is intended to be placed in a flash or ROM memory. On a hosted system it is normally placed in the same memory as `code`.

### `farswitch`

Holds switch tables in the main memory, address range `0x000000-0xffffffff`. This section is intended to be placed in a flash or ROM memory. On a hosted system it is normally placed in the same memory as `code`.

### `idata`

Holds initializers for the `data` section. This section is created by the linker by cloning the `data` section provided by the compiler. This section is placed in the main memory, address range `0x0000-0xffff` and needs to be placed in a flash or ROM. This section is not used when linking for a hosted system.

### `data_init_table`

This section is created by the linker and filled in with information to the C startup code on how to copy and clear memory regions to properly initialize the data object before giving control to `main()`. This section can be placed anywhere in memory, address range `0x000000-0xffffffff`. As it is a constant table, it can usually go in the same memory as `farcode` (or `code` when using the small Code and Data models).



This chapter describes the supplied C runtime library, adapted from the Apache NuttX library.

### 18.1 Design considerations

The original Apache NuttX library is a collection of routines from the C standard library, POSIX standard, and other components common in real-time operating systems.

It is a scalable and highly configurable runtime library that works in both 8-bits and 32-bits environments.

The adapted version becomes a Standard C library, with POSIX and RTOS-style components either removed or disabled.

Apart from that, the following are the key changes:

- Replaced header files closely related to the compiler (e.g., `stddef.h`, `stdarg.h`, `setjmp.h`, `stdint.h`).
- Internal names are changed to start with either double underscores or a single underscore followed by a capital letter. This avoids namespace pollution, as such identifiers are reserved for the compiler vendor.
- The stub interface uses names prefixed by `_Stub` instead of `up_` for clarity and to avoid namespace pollution.

---

Note: The Apache NuttX library's configuration utility, designed for environmental tailoring, is not used here; the build system is different. This allows the library to be used without extensive configuration, as the compiler is known and target-board-specific matters are excluded. Some library functions have different versions to reduce memory footprint by making certain capabilities optional (e.g., formatters; see [Library on a diet](#)).

---

## 18.2 Using the library

The C library consists of header files that are immediately available to use by the compiler and can be included:

```
#include <stdio.h>

int main () {
    printf("Hello World!\n");
    return 0;
}
```

When linking, the C library does not need to be explicitly added on the command line, as the linker automatically finds the correct C runtime library by examining the settings used during C object file compilation:

```
$ ln65816 main.o linker-rules.scm
```

C library variants matching different compiler settings are provided in the installation directory; the appropriate one is automatically selected.

The linker detects mixing object files compiled with different settings or incompatible third-party libraries. This relies on matching runtime attributes in object files. A mismatch results in a descriptive linker error.

## 18.3 Provided library files

The following table lists the provided ready to use library files:

Table 18.1: Library variants

| Library name          | Code model | Data model | Size of double | Target system |
|-----------------------|------------|------------|----------------|---------------|
| clib-sc-sd.a          | Small      | Small      | 32 bits        | Generic       |
| clib-cc-sd.a          | Compact    | Small      | 32 bits        | Generic       |
| clib-lc-sd.a          | Large      | Small      | 32 bits        | Generic       |
| clib-sc-md.a          | Small      | Medium     | 32 bits        | Generic       |
| clib-cc-md.a          | Compact    | Medium     | 32 bits        | Generic       |
| clib-lc-md.a          | Large      | Medium     | 32 bits        | Generic       |
| clib-sc-ld.a          | Small      | Large      | 32 bits        | Generic       |
| clib-cc-ld.a          | Compact    | Large      | 32 bits        | Generic       |
| clib-lc-ld.a          | Large      | Large      | 32 bits        | Generic       |
| clib-sc-hd.a          | Small      | Huge       | 32 bits        | Generic       |
| clib-cc-hd.a          | Compact    | Huge       | 32 bits        | Generic       |
| clib-lc-hd.a          | Large      | Huge       | 32 bits        | Generic       |
| clib-sc-sd-double64.a | Small      | Small      | 64 bits        | Generic       |
| clib-cc-sd-double64.a | Compact    | Small      | 64 bits        | Generic       |
| clib-lc-sd-double64.a | Large      | Small      | 64 bits        | Generic       |
| clib-sc-md-double64.a | Small      | Medium     | 64 bits        | Generic       |
| clib-cc-md-double64.a | Compact    | Medium     | 64 bits        | Generic       |
| clib-lc-md-double64.a | Large      | Medium     | 64 bits        | Generic       |
| clib-sc-ld-double64.a | Small      | Large      | 64 bits        | Generic       |
| clib-cc-ld-double64.a | Compact    | Large      | 64 bits        | Generic       |
| clib-lc-ld-double64.a | Large      | Large      | 64 bits        | Generic       |
| clib-sc-hd-double64.a | Small      | Huge       | 64 bits        | Generic       |

continues on next page

Table 18.1 – continued from previous page

| Library name               | Code model | Data model | Size of double | Target system |
|----------------------------|------------|------------|----------------|---------------|
| clib-cc-hd-double64.a      | Compact    | Huge       | 64 bits        | Generic       |
| clib-lc-hd-double64.a      | Large      | Huge       | 64 bits        | Generic       |
| clib-sc-sd-c256.a          | Small      | Small      | 32 bits        | Foenix C256   |
| clib-cc-sd-c256.a          | Compact    | Small      | 32 bits        | Foenix C256   |
| clib-lc-sd-c256.a          | Large      | Small      | 32 bits        | Foenix C256   |
| clib-sc-md-c256.a          | Small      | Medium     | 32 bits        | Foenix C256   |
| clib-cc-md-c256.a          | Compact    | Medium     | 32 bits        | Foenix C256   |
| clib-lc-md-c256.a          | Large      | Medium     | 32 bits        | Foenix C256   |
| clib-sc-ld-c256.a          | Small      | Large      | 32 bits        | Foenix C256   |
| clib-cc-ld-c256.a          | Compact    | Large      | 32 bits        | Foenix C256   |
| clib-lc-ld-c256.a          | Large      | Large      | 32 bits        | Foenix C256   |
| clib-sc-hd-c256.a          | Small      | Huge       | 32 bits        | Foenix C256   |
| clib-cc-hd-c256.a          | Compact    | Huge       | 32 bits        | Foenix C256   |
| clib-lc-hd-c256.a          | Large      | Huge       | 32 bits        | Foenix C256   |
| clib-sc-sd-double64-c256.a | Small      | Small      | 64 bits        | Foenix C256   |
| clib-cc-sd-double64-c256.a | Compact    | Small      | 64 bits        | Foenix C256   |
| clib-lc-sd-double64-c256.a | Large      | Small      | 64 bits        | Foenix C256   |
| clib-sc-md-double64-c256.a | Small      | Medium     | 64 bits        | Foenix C256   |
| clib-cc-md-double64-c256.a | Compact    | Medium     | 64 bits        | Foenix C256   |
| clib-lc-md-double64-c256.a | Large      | Medium     | 64 bits        | Foenix C256   |
| clib-sc-ld-double64-c256.a | Small      | Large      | 64 bits        | Foenix C256   |
| clib-cc-ld-double64-c256.a | Compact    | Large      | 64 bits        | Foenix C256   |
| clib-lc-ld-double64-c256.a | Large      | Large      | 64 bits        | Foenix C256   |
| clib-sc-hd-double64-c256.a | Small      | Huge       | 64 bits        | Foenix C256   |
| clib-cc-hd-double64-c256.a | Compact    | Huge       | 64 bits        | Foenix C256   |
| clib-lc-hd-double64-c256.a | Large      | Huge       | 64 bits        | Foenix C256   |

Note: In addition to the Foenix C256 there are libraries provided for Foenix F256 with the `-f256` suffix, SNES with the suffix `--snes` and SNES not using the PPU multiplier with the `-snes-noppu` suffix.

Note: There is normally no need to specify the C library to use when linking. The linker is able to automatically pick the correct C library by inspecting the C object files.

## 18.4 Stubs interface

How do you provide functions such as `fopen()`, `fprintf()`, or `assert()` when using a cross-compiler that may not know the final execution environment of an application, or even have a file system?

Functions like `fprintf()` pass through several library layers, handling formatting, buffered I/O, and stream interfaces. Ultimately, they output characters to a display or file. This is managed by a simple stubs API defined in the `calypsi/stubs.h` header file.

### 18.4.1 Semi-hosting

Semi-hosting is the concept where the debugger implements operations, such as I/O, on behalf of the target. This involves using a single breakpoint where the target temporarily stops to exchange data with the debugger, which performs the actual operation on the host. After the operation, any return value is passed back, and execution resumes on the target.

The Calypsi C compiler tool chain includes a semi-hosted debug implementation of the stubs API within the standard C library. To enable semi-hosting, add the `--semi-hosted` command-line option to the linker.

File operations are performed in the host filesystem by the debugger in a directory specified by the `--semi-hosted-root` command-line option, which defaults to the current directory.

The standard streams `stdin`, `stdout` and `stderr` are fully supported. Depending on your IDE you may have a console window for these streams.

Semi-hosting can be used during development for tracing or prototyping before a real I/O system is implemented in the application. In a final application, you will most likely use target-specific stub functions, not semi-hosted variants.

---

Note: Even if you implement some stub actions, you can still link with `--semi-hosted` for partial semi-hosted support. This allows you to implement a file system while retaining `assert()` through the semi-hosted mechanism. The C library's semi-hosted stubs are compiled with `--weak-symbols`, so your implementations take precedence. Unimplemented stubs will use their weak semi-hosted variants.

---

### 18.4.2 Stubs header

The `calypsi/stubs.h` file looks as follows:

```
/******
 *
 * Copyright Håkan Thörngren
 *
 * This file is part of the Calypsi C library.
 * Permission to use with the Calypsi tool chain is hereby granted.
 *
 *****/

#ifndef __INCLUDE_CALYPSI_STUBS_H
#define __INCLUDE_CALYPSI_STUBS_H

#include <calypsi/config.h>
#include <stddef.h>

/* Currently not implemented. */
#define __noreturn_function

/******
 * Public Data
 *****/

#if defined(__cplusplus)
extern "C"
{
#endif
```

(continues on next page)

(continued from previous page)

```

/*****
 * Debug interfaces exported by the architecture-specific logic
 *****/

/*****
 * Name: _Stub_open
 *
 * Description:
 *   Open a file.
 *   The oflag argument are POSIX style mode flags, e.g O_RDONLY which
 *   are defined in fcntl.h.
 *   This function is variadic as it optionally can take a mode_t that
 *   are permissions, e.g 0666. If the file system does not handle
 *   permissions you can ignore that this function is variadic.
 *   The return file descriptor shall be a positive number, larger
 *   than 2 (as 0-2 are used for stdin, stdout and stderr).
 *   The actual number does not matter and they need not to be
 *   consecutive, multiple numeric series with gaps between can be used.
 *
 * Return the obtained file descriptor or the desired errno value negated
 * if there is an error.
 *****/

int _Stub_open(const char *path, int oflag, ...);

/*****
 * Name: _Stub_close
 *
 * Description:
 *   Close a file
 *
 * Return 0 if operation was OK or the desired errno value negated
 * if there is an error.
 * Note: This will only be invoked for streams opened by _Stub_open(),
 *       there is no need to check for the standard descriptor 0-2.
 *****/

int _Stub_close(int fd);

/*****
 * Name: _Stub_lseek
 *
 * Description:
 *   Change position in a file
 *
 * Returns the new position in the file in bytes from the beginning of the
 * file, or the desired errno value negated if there is an error.
 *****/

long _Stub_lseek(int fd, long offset, int whence);

/*****

```

(continues on next page)

(continued from previous page)

```

* Name: _Stub_fgetpos
*
* Description:
*   Get current position in a file
*
* Returns 0 on success, or the desired errno value negated if there is an error.
*
*****/

int _Stub_fgetpos(int fd, fpos_t *pos);

/*****
* Name: _Stub_fsetpos
*
* Description:
*   Change position in a file
*
* Returns 0 on success, or the desired errno value negated if there is an error.
*
*****/

int _Stub_fsetpos(int fd, const fpos_t *pos);

/*****
* Name: _Stub_read
*
* Description:
*   Read from a file
*
* Returns the number of characters read or the desired errno value negated
* if there is an error.
*
*****/

size_t _Stub_read(int fd, void *buf, size_t count);

/*****
* Name: _Stub_write
*
* Description:
*   Write to a file
*
* Returns the number of characters actually written or the desired errno
* value negated if there is an error.
*
*****/

size_t _Stub_write(int fd, const void *buf, size_t count);

/*****
* Name: _Stub_rename
*
* Description:
*   Rename a file or directory
*
* Return 0 on success or the desired errno value negated if there is an error.

```

(continues on next page)

(continued from previous page)

```

*
*****/

int _Stub_rename(const char *oldpath, const char *newpath);

/*****
* Name: _Stub_remove
*
* Description:
*   Remove a file or directory
*
* Return 0 on success or the desired errno value negated if there is an error.
*
*****/

int _Stub_remove(const char *path);

/*****
* Name: _Stub_exit
*
* Description:
*   Terminate the program with an exit code, exit clean-ups are done
*   before this function is (finally) called.
*
*****/

void _Stub_exit(int exitCode) __noreturn_function;

/*****
* Name: _Stub_environ
*
* Description:
*   Get the environment. On UNIX this is typically a global variable
*   'environ', but in order to make it more flexible and avoid having
*   such global variable (which is not part of the C standard) it is
*   obtained using the stub interface.
*
* Note:
*   This stub function is not implemented by the semi-hosted debug stub
*   interface.
*
*****/

char** _Stub_environ(void);

/*****
* Name: _Stub_assert
*
* Description:
*   Handle an assertion
*
*****/

void _Stub_assert(const char *filename, int linenum) __noreturn_function;

#if defined(__cplusplus)

```

(continues on next page)

(continued from previous page)

```
}  
#endif  
  
#endif /* __INCLUDE_CALYPSI_STUBS_H */
```

### 18.4.3 Custom I/O

You can provide your own low-level I/O stubs. File descriptors are integers indexing an internal lookup table (typically an array). If the I/O system is used, the library statically allocates the first three streams: `stdin`, `stdout`, and `stderr`. Additional file streams are dynamically allocated from the heap.

### 18.4.4 Error handling

If a stub action encounters an error, it should return the negated `errno` value. The calling C library function then performs appropriate error handling, setting `errno` to the corresponding positive error code.

## 18.5 Examine use of library

When working with a memory-constrained system, the risk of running out of memory is always present. To help you understand memory usage, the linker can generate a list file with cross-reference information. This valuable tool details what is placed in memory, its location, and why it occupies that space.

You can tell the linker to provide a list file in the following way:

```
$ ln65816 main.o clib-lc-sd.a linker-rules.scm --list-file=project.lst --cross-reference
```

This will result in a list file names `project.lst` which contains several parts showing things such as:

- A summary of the memories
- An overview of sections and where they take space in the memories
- The object files used and from which library they have been extracted from
- Complete cross reference, showing where every section fragment is located with its size, together with:
  - What symbols are defined
  - What symbols are referenced
  - Who is referencing me
- An overall summary of total memory size

## 18.6 Library on a diet

If you use the C library in a memory-constrained system, you may find it rapidly consumes code space. This is because the C library provides substantial functionality, with many functions acting as simple facades to complex underlying implementations.

There are ways to tune things to reduce the memory footprint. However, first understand the current memory usage before making changes.

### 18.6.1 Formatters

The C library provides various `printf()` and `scanf()` variants, each with different capabilities and memory requirements.

While you can specify a format function, the compiler inspects used format strings and outputs this information to the object file. The linker then automatically selects the smallest variant compatible with those format strings.

---

Note: Automatic format function selection requires using string literals in calls. This is good practice, as it also allows the compiler to validate argument list consistency with the format string.

---

If you for some reason have to override the formatter used and dictate a specific one, you can do so using the `--rtattr` attribute which for `printf()` would be:

```
$ ln65816 --rtattr printf=nofloat files...
```

The default print formatter as mentioned above, is selected automatically to match your needs. The following table describes the available print formatters:

Table 18.2: Print formatters

| Capability                  | printf=reduced | printf=medium | printf=nofloat | printf=float |
|-----------------------------|----------------|---------------|----------------|--------------|
| basics, c d i o p s u X x % | yes            | yes           | yes            | yes          |
| format flag, 0 + - #        | no             | yes           | yes            | yes          |
| field width                 | no             | yes           | yes            | yes          |
| long long                   | no             | no            | yes            | yes          |
| float, e f g E F G          | no             | no            | no             | yes          |

The `scanf()` function exists in three variants, mainly depending on if you want support for floating point numbers or not:

Table 18.3: Read formatters

| Capability                  | scanf=medium | scanf=nofloat | scanf=float |
|-----------------------------|--------------|---------------|-------------|
| basics, c d i o p s u X x % | yes          | yes           | yes         |
| long long                   | no           | yes           | yes         |
| float, e f g E F G          | no           | no            | yes         |

---

Note: An appropriate default formatter for `scanf()` is automatically picked by the linker based on the needs of your application.

---

### 18.6.2 Reduced exit

Terminating an application involves closing open files and executing `atexit()` functions. For embedded applications, which may never exit, termination code can be redundant; closing files also adds related code. An implicit call to `exit()` occurs when `main()` returns. If proper exit code is undesirable, you can exclude it.

The simplest way to reduce exit code overhead is to add `--rtattr exit=simplified` to the linker command line. This prevents closing files or calling `atexit()` handlers. Instead, `exit()` will immediately jump to `_Stub_exit()`, the lowest-level termination routine.

### 18.6.3 Consider alternatives

If the library still occupies too much space, consider replacing certain functions with smaller, less flexible alternatives.

## 18.7 Override library functions

You may sometimes need to override a library function. Consider using an alternative function with a different name rather than replacing a standard one. Replacing a library function is appropriate if you have a more suitable application-specific implementation, especially when other library functions call it.

To replace a library function, add a new source file to your project using the same function name and prototype, then build your application normally.

---

Note: Function replacement works because library functions use weak symbols. Your replacement functions are non-weak by default, taking precedence over library functions during linking.

---

## 18.8 C startup

The initial system configuration is performed by the C startup object which is included in the standard C library.

The C startup is responsible for setting up the execution environment before giving control to the `main()` function. This typically includes setting up the stack pointer and providing initial values to static variables.

The C startup object also contains a couple of small optional sections that are only included if needed. They are responsible for initializing optional parts of the runtime system, such as the heap and the file streams.

If you do not use functions such as `malloc()`, the heap is not needed and the code to initialize the heap memory system is omitted. The same happens for file streams. If there are no file streams used, the code related to initializing and terminating them are omitted. This is done in order to reduce the memory footprint of the final application.

In case you want to study the source code of the C startup it can be found at `src/lib/lowlevel/cstartup.s` under the installation directory.

### 18.8.1 Customizing the startup

In many cases the default C startup will work fine without any changes. However, there are a couple of typical situations when some custom configuration is needed. One example is that the memory system needs to be configured immediately at power on, or if your application is going to run under an operating system that may impose special rules on initialization and termination.

If you only need to run some code to perform early initialization of the hardware you can provide your own `__low_level_init()` function which is called early in the C startup object, before any static C variables are given their start values.

If you need to make actual changes to the C startup object the easiest way is to copy the existing one to your project, make changes as required and include it in your build system.

To be able to properly suppress the C startup object in the C library you need to provide a different value for the `cstartup` runtime attribute in your own `cstartup.s`. The `cstartup.s` assembly source file starts with:

```
;;; C startup variant, change attribute value if you make your own
        .rtmodel cstartup,"normal"
```

You need to change the value `normal` to something else. What value you use does not matter so much, but using a descriptive value is probably a good idea:

```
;;; C startup variant, change attribute value if you make your own
        .rtmodel cstartup,"mycustom"
```

The modified C startup source file needs to be included in your build system. The linker also needs to be informed that it should use a specific C startup object, which is done using the `--cstartup` command line option:

```
% ln65816 <object-files> <your-startup.o> clib-lc-sd.a --cstartup=mycustom
```

Note: The C startup object has been carefully crafted to have a small footprint and to properly initialize the C runtime system. Even though it makes calls to handle the file system, initialize data areas and the heap, it is done in such way that the actual code is only included in the final application if these subsystems are used.

A good understanding of the tools and the provided C runtime is needed in order to make non-trivial changes to the C startup object. Incorrectly made changes may cause linker errors, unused subsystem being pulled in, or result in a runtime that is not properly initialized.

## 18.9 Time and date

The provided C library implements the functions and definitions in the `time.h` system header file. The time related types are defined as 64-bit types to avoid the year 2038 problem and reduce issues with overflow.

### 18.9.1 Providing a time

C defines two functions `clock()` and `time()` that provides a concept of a time. The library provides a default implementation for them that returns a value with all bits set, which means that the time is not available.

To use actual times you will need to provide your own implementation of `clock()` or `time()` and include it in your project.

The `CLOCKS_PER_SEC` macro is set to `1000` by the `stdint.h` header file.

## 18.10 Rebuilding the C library

The C library is designed to allow for a great deal of configuration without resorting to rebuild it. Situations where you may need to rebuild the C library are when you need a specific variant build, or to enable debug information to debug the library itself.

For the most part, building the library is straightforward, but there are certain files that needs to be built in certain ways, e.g. variants of `printf()`. To make it easier to build a variant C library there are build scripts provided in the `library-build` folder in the installation directory.

There is one build file for each library variant that are included in the installation. Select one that has a similar configuration as a starting point. It is a good idea to make copy of the build script and modify it to suit your needs.

The compiler applies various optimizations to reduce code size and increase performance. This chapter describes how to control optimizations and highlights useful aspects, rather than detailing every optimization applied.

### 19.1 Overview

By default, the compiler performs basic rewrites, optimizes code selection for efficient instruction set utilization, and makes good use of the CPU's internal registers.

The resulting application is generally well-suited for debugging, as many optimization passes are disabled by default.

#### 19.1.1 General settings

Optimizations can be broadly enabled using the following command-line options.

`-O1`

Enable some optimization passes.

--02

Enable all provided optimization passes.

--speed

Allow application size to grow in order to make the application run faster.

--space

Tune optimizations more towards making the application small. This is the default.

## 19.2 Function inlining

Function inlining expands a function's body at its call site, typically increasing performance but potentially increasing code size.

Function inlining may be beneficial on code size for small functions as the inserted function body becomes tailored to the particular code surrounding it. Normally when making a function call there is a calling convention that has to be followed which puts restrictions on how registers are used.

Inlining is enabled by specifying at least -01.

### 19.2.1 The inline keyword

The `inline` keyword in C is sometimes misunderstood. Its purpose is to expose a function body to many compilation units and hint that the function may be inline expanded instead of making an ordinary call to it. There is no guarantee that a function marked as `inline` is actually inlined, this decision is taken by the compiler. Furthermore, the compiler may also choose to inline functions that are not marked as `inline`.

### 19.2.2 Using the inline keyword

A function marked as `inline` is normally placed in a header file to allow it to be used by multiple compilation units:

```
inline max (int a, int b) {  
    if (a > b) {  
        return a;  
    } else {  
        return b;  
    }  
}
```

An inline function can also be marked as `static inline` in a similar way:

```
static inline max (int a, int b) {  
    if (a > b) {  
        return a;  
    } else {  
        return b;  
    }  
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

You can use functions marked as `inline` or `static inline` in the same way as any other function. They can be called and in that case they may be subject to inline expansion and you can store it as a function pointer and make calls to it using that function pointer.

There are two cases when a function marked as `inline` or `static inline` requires a separate copy of the function. First, the compiler may choose to make an ordinary function call to it. Second, a function marked as `inline` may be stored in a function pointer.

### 19.2.3 Variants of inline functions

As hinted in the previous section, there are two kinds of inline functions, plain `inline` and `static inline`. As long as the compiler chose to do inline expansion these two variants work identically. They differ when the compiler choose to not inline expand it and make an ordinary function call instead.

#### `inline`

A function is marked as `inline` does not cause the compiler to generate a separate definition of the function. If such separate function is needed, it is assumed that one compilation unit provides the definition using the `extern` keyword:

```
extern inline max (int a, int b);
```

This should be placed in an ordinary C source file, not a header file.

In most cases you will need to tell the compiler to actually generate such separate function using an `extern` declaration somewhere. If you do not do this and the function is required by the application, a linker error will result that tells you that the function is not defined.

The main benefit of using `inline` rather than `static inline` is that you are ensured that there will be at most one separate definition of the function.

---

Note: One subtle benefit which `inline` has over `static inline` is that if you really must expand the function inline, then you can omit the `extern` declaration of the function. The linker will then let you know that the required inline expansion did not happen. In this case you most likely want to specify the command-line option `--always-inline` to ensure that the inliner is always used.

---

#### `static inline`

The compiler will generate a `static` standalone version of a `static inline` function whenever there is a need for it. Being `static` it is local to the compilation unit with no external linkage.

This means that if you make use of the same `static inline` function in several compilation units, there may be multiple copies of the same function.

---

Note: If you take the address of the same `static inline` function in different translation units then the result will not compare equal.

---

---

Note: While it is slightly easier to use a `static inline` function as you do not need to provide a single `extern` declaration, it comes with the cost of potentially having duplicate code in the final application.

---

### 19.2.4 Controlling inline expansion

The compiler will apply its own decisions on what to inline and when to use normal function calls. Small simple functions and static functions with only a single use are prime candidates for inline expansion.

There are a couple of command-line options that allow for tuning which functions are considered for inlining:

#### `--always-inline`

Always enable the inliner, no matter which optimization level is used.

This is useful if you have functions that you always want to be inlined.

#### `--no-inline`

Do not perform any inlining at all.

#### `--only-marked-as-inline`

Only consider functions that have the `inline` keyword. This prevents functions not marked as `inline` from being considered for inlining.

#### `--strong-inline`

Try to obey the `inline` keyword. This relaxes some of the requirements, such as that the function must be small enough. This is a strong hint that we really want functions marked for inlining to be inlined. The compiler may still refuse to inline a function.

---

Note: There are a number of reasons why the compiler may refuse to inline a function, e.g. used recursively, there is no prototype, it uses a variable argument list or it uses variable length arrays.

---

#### `--inline-on-matching-custom-text-section`

Functions being placed in a custom section (using `#pragma clang section`) are only allowed to be inlined into a function that belongs to the same custom section.

This is useful in situations when the memory system is being manipulated and functions are placed in a custom section to control placement to obey the memory setup. Especially if you have several such configurations you may want to prevent functions from being inlined across custom sections as they may make assumptions on which section (memory placement) they are being executed from.

## 19.3 Cross call

This optimization pass can sometimes result in quite significant savings on the application size. It examines the almost final program looking for instruction sequences that are identical. Such sequences are lifted out to separate subroutines which are called instead. This takes in account the size of the code sequences as well as how many times they appear.

### 19.3.1 Controlling cross call

Cross call is enabled by specifying at least `-O2`. Even though cross call is focusing on the application size, it is enabled also when the `--speed` command-line option is specified, as the savings on code space can sometimes be quite significant.

`--no-cross-call`

This option disables the cross call optimization. This can be used with `-O2` when you want to avoid the subroutine call overhead that cross call will introduce.

`--no-interprocedural-cross-jump`

This option disables the cross jump between functions optimization. This can be used if you are manually mapping in functions by some bank mechanism in a way so that not all functions in the same compilation are visible at the same time.



---

## Assembly language interface

---

Assembly language provides symbolic access to target machine instructions. You might need this control for specific instructions unrepresentable in C, precise hardware interaction, exception stack frame manipulation, exact timing sequences, or performance-critical routines.

### 20.1 Intrinsic functions

The compiler provides intrinsic functions (declared in the `calypsi/intrinsics65816.h` file) that resemble ordinary functions. Instead of a function call, an intrinsic generates a specific instruction sequence. For example, `__disable_interrupts()` emits machine instructions to disable normal interrupts. See [Intrinsic functions](#) for details.

### 20.2 Assembly functions

You can implement functions in assembly language and call them from C like any other function. An assembly function must adhere to the C calling convention, which dictates how values are passed to the called function and where the return value is placed.

---

**Note:** To simplify your assembly routine interface, consider using the `simple_call` calling convention.

---

You can choose between a separate assembly source file or inline assembly; each has pros and cons. Assembly functions offer better separation between C and assembly, aiding portability. However, function call overhead and adherence to calling conventions may be undesirable.

Minimal boilerplate assembly code is required to place the routine in a suitable section and declare public symbols. This assembly code resides in a separate file, which must be added to the build system.

Assembly language files typically use the `.s` or `.asm` extension; however, this varies due to the lack of standardization in assembly language itself.

The assembler provided by the Calypsi C compiler tool chain, similar to UNIX assemblers, uses directives starting with a dot. This avoids name clashes with instructions, whose naming conventions vary widely across targets, ensuring consistent directive names.

As a minimum, you must declare the section and export your function name using the `.public` directive:

```
        .section code
        .public myFunction      ; export myFunction
        .extern _Dp             ; pseudo direct page registers
myFunction:
        clc
        adc     dp:_Zp
        rts
```

To call this function correctly, you must provide a prototype in C:

```
#include <stdint.h>

extern int8_t myFunction(int8_t a, int8_t b);

int caller(int8_t x) {
    return myFunction(5, x);
}
```

### 20.2.1 Generate skeleton code

The easiest way to generate an assembly source file is to have the compiler create it using the `--assembly-source` command-line option. A simplified C source file containing the desired functions and declarations can be used for this purpose.

You can provide desired function definitions with simple parameter uses to study how they are passed:

```
extern int intvar;
extern char charvar;

extern void externalFunction(int*);

int myFunction(int i, char c) {
    int local = i;
    intvar = i;
    charvar = c;
    externalFunction(&local);
    return local;
}

int main() {
    myFunction(intvar, charvar);
    return 0;
}
```

```
$ cc65816 skeleton.c --assembly-source=skeleton.s
```

```
; Generated by Calypsi ISO C compiler for 65816
```

```
        .rtmodel version,"1"
        .rtmodel codeModel,"large"
        .rtmodel dataModel,"small"
```

(continues on next page)

(continued from previous page)

```

        .rtmodel core,"65816"
        .rtmodel huge,"0"
        .rtmodel target,"none-specified"
        .extern _Dp
        .extern _Vfp
        .extern charvar
        .extern externalFunction
        .extern intvar
; extern int intvar;
; extern char charvar;
;
; extern void externalFunction(int*);
;
; int myFunction(int i, char c) {
        .section farcode,text
        .public myFunction
myFunction: phy
;   int local = i;
        sta     1,s
;   intvar = i;
        sta     intvar
;   charvar = c;
        sep     #32
        lda     6,s
        sta     charvar
        rep     #32
;   externalFunction(&local);
        tsc
        inc     a
        jsr     long:externalFunction
;   return local;
        lda     1,s
; }
        ply
        rtl
;
; int main() {
        .section farcode,text
        .public main
main:
;   myFunction(intvar, charvar);
        lda     charvar
        pha
        lda     intvar
        jsr     long:myFunction
;   return 0;
        pla
        lda     ##0
; }
        rtl

```

## 20.3 Calling convention

The default calling convention is complex in detail, but straightforward in most common scenarios. An alternative calling convention `simple_call` is also provided.

If parameters are passed on the stack, the caller is responsible for cleanup. The called function may use any register resource but must preserve certain registers, saving and restoring them before returning.

### 20.3.1 Simple calling convention

Use `__attribute__((simple_call))` or `__simple_call` on a function declaration to enable the simple calling convention.

In this calling convention the first parameter is passed in register if it fits, the remaining parameters are passed on the stack.

Registers A, X, Y and pseudo registers `_Dp[0-7]` are destroyed by a function call. Pseudo registers `_Dp[8-15]` must be preserved.

---

Note: The pseudo registers are clobbered for consistency with the normal calling convention. If you implement an assembly or foreign language interface that do not use the `_Dp` style pseudo registers you can simply ignore this detail as it does not matter.

---

Table 20.1: first parameter register location

| Register | Size | Types                              |
|----------|------|------------------------------------|
| A        | 8    | char                               |
| C        | 16   | int, short, tiny and near pointers |
| X and C  | 32   | long, float, far and huge pointers |

Table 20.2: Return values

| Register | Size | Types                        |
|----------|------|------------------------------|
| A        | 8    | char, tiny pointer           |
| C        | 16   | int, short, tiny pointer     |
| X and C  | 24   | far pointer, farfunc pointer |
| X and C  | 32   | long, float                  |

### 20.3.2 Normal calling convention

Parameters are passed in the A accumulator, index register X and pseudo registers `_Dp[0-7]`.

Registers A, X, Y and pseudo registers `_Dp[0-7]` are destroyed by a function call. Pseudo registers `_Dp[8-15]` must be preserved.

---

Note: The A register is the 8 bits accumulator. The same register in 16 bits version is called C. As the 65816 shares mnemonics with the 6502 and can switch modes, the term register name A is often used for the 16 bits version as well. Technically it should be named C when being the 16 bits version.

---

Table 20.3: Parameter registers

| Register | Size | Types                                     |
|----------|------|-------------------------------------------|
| A        | 8    | char                                      |
| C        | 16   | int, short, tiny and near pointers        |
| _Dp[0-1] | 16   | int, short, tiny and near pointers        |
| _Dp[2-3] | 16   | int, short, tiny and near pointers        |
| _Dp[4-5] | 16   | int, short, tiny and near pointers        |
| _Dp[0-3] | 32   | long, float, far, far24 and huge pointers |
| _Dp[4-7] | 32   | long, float, far, far24 and huge pointers |
| X and C  | 32   | long, float                               |

Parameters are bound to register left to right on a first fit basis. If a parameter register has to be skipped over, it will be considered again for later parameters. Parameters that cannot be fit into registers are passed on the stack.

Note: The 16 bit register \_Dp[2-3] is only used in the small data model. In the other data models this parameter is passed in \_Dp[4-5] instead. The reason for this is that in the small data model it is possible to use plain 16 bits pointers. In the other data models the 16 bit register may need to be converted to a 24 bit register value.

Note: 24 bit values are passed in three bytes, but the actual pseudo register covered is four bytes. The last byte is ignored and can be regarded as unused garbage. This is done in order to avoid data size shifts in the code.

Table 20.4: Return values

| Register | Size | Types                        |
|----------|------|------------------------------|
| A        | 8    | char, tiny pointer           |
| C        | 16   | int, short, tiny pointer     |
| X and C  | 24   | far pointer, farfunc pointer |
| X and C  | 32   | long, float                  |

### 20.3.3 Index and data size

The 65816 is used in native mode with 16 bit registers. In some situations the runtime needs to switch to 8 bit register mode to implement the correct behavior. This is done automatically and the compiler will then switch back to 16 bits mode.

There are two flags in the CPU that control the width of registers, one for index registers and the other for the accumulator and data memory access. The index width is always kept at 16 bits by the runtime, only the accumulator/memory width is changed as needed.

A function call expects the width to be set to 16 bits and must return with the same setting. This means that if you write a function in assembly language, it will get called from a C function with 16 bit register width and you need to return that way. Inside your own routine you are free to switch register width as you may see fit.

Note: The compiler may generate code that reads 8 and 24 bit objects using 16 bit access, ignoring the extra byte. Writes are always done obeying the object width. If you are reading from hardware registers

that requires reads to follow the width, you must use the `volatile` keyword.

It can be a good idea to group multiple 8 bit accesses next to each others as the compiler is able to reduce mode shifts back and forth if it detects it. However, you need to enable the optimizer for this to happen.

---

### 20.3.4 64 bit values

The above tables do not mention 64 bit values. A 64 bit value such as `long long` and `long double` is passed by reference, that is, as a pointer to the value. The size of this pointer is the same as the default pointer size and is either 16 bits or 24 bits depending on the data model used.

If a function returns a 64 bit value, the caller is responsible for allocating space. An extra ‘invisible’ parameter, a pointer to this space, is added to the function. The called function is expected to return this pointer.

### 20.3.5 Structure passing

Structure parameters are passed on the stack. Structure return values are handled like 64 bit values. The caller allocates space and adds an extra ‘invisible’ parameter (a pointer to that space) to the function call. The called function is expected to return this pointer.

## 20.4 Inline assembler

The inline assembler allows you to insert and interface assembly code slices within a C function. This avoids call overhead and can improve parameter adaptation. However, the optimizer must be more cautious, which may affect the performance of the surrounding C code.

### 20.4.1 Basic inline assembly

You can insert a slice of assembly code using an `__asm` block:

```
int counter;

void foo(char xx) {
    __asm(" sei\n"
          " inc counter\n"
          " bne skip\n"
          " inc counter+2\n"
          "skip: cli\n"
          );
}
```

Each line without a label requires at least one leading space and must be terminated by a newline character (`\n`).

The inline assembler supports the full assembly instruction set, including literal bytes, volatile operations, local labels, and register allocation for parameters and return values.

Goto labels and most assembler directives are currently not supported by the inline assembler. If you need better control with placement, use a separate assembly source file instead.

When inline assembly is inserted, the compiler adapts it to fit the generated C code. Variables can be passed as parameters, and a single result variable is supported. The compiler reasonably understands the inserted assembly, mixing it with C-generated code. Inline assembly is subject to low-level optimizations when the optimizer is enabled.

### 20.4.2 Volatile

An assembly block can be marked as `volatile`:

```
__asm volatile { ... }
```

This has the effect that all memory accesses in the assembly slice are treated as side effects. Otherwise the optimizer may remove reads from memory when the value read is not used.

### 20.4.3 Local labels

Local labels can be used and their names will not clash with C identifiers. When an inline assembly slice is inserted, local labels are converted to internal C labels, preventing name clashes with C identifiers.

### 20.4.4 External symbols

Inline assembly can refer to symbols defined outside its code slice, provided such symbols are visible at the C level within the same compilation unit.

### 20.4.5 Constraints

An inline assembly code slice can refer to C variables and return a value. The inline assembly construct optionally accepts three lists:

1. Output variable: Specifies a C variable to represent the returned value and a register class where the assembly code block places it. The result is prefixed by `=` in its single-value list.
2. Input expressions: Typically variables. The compiler evaluates the expression and places it in the specified register class.
3. Clobbered registers: Any register resource clobbered by the inline assembly must be specified here.

Multiple entries in a list are comma-separated.

#### Register classes

A register class is a register resource that can represent either a single register or a set of equivalent registers. They are used for allocating parameters and determining the return value location.

The following register classes are defined:

Table 20.5: Register classes

| Register class | Description                              |
|----------------|------------------------------------------|
| a              | Accumulator A, 8 bit                     |
| c              | Accumulator C (B and A combined), 16 bit |
| x              | X index register                         |
| y              | Y index register                         |
| dp8            | An 8-bit direct page pseudo register     |
| dp16           | A 16-bit direct page pseudo register     |
| dp32           | A 32-bit direct page pseudo register     |

Note: Register classes are used internally during code generation. Internal code generator rules ensure safe register allocation by adhering to specific rules and invariants. Rather than attempting to diagnose inline assembly constraints or impose conservative limitations, the compiler trusts you. Overusing resources or violating internal invariants may lead to a register allocation error. In such cases, ease the register resources to find a working allocation.

## Registers and constraints

The following code shows how constraints for an inline assembly code slice are defined:

```
char foo(int xx) {
    int out;
    __asm("  inx\n"
          "  inx\n"
          "  inx\n"
          "  inx\n"
          "  txa\n"
          : "=Kc" (out)
          : "Kx"  (xx)
          : "c", "x"
    );
    return out;
}
```

Currently, the only supported constraint is 'K', followed by a register class.

An empty list can be entered by using a colon character followed by nothing.

The first list is the optional output parameter, describing the C variable where the output is visible after the inline assembly slice executes. The inline assembly slice must leave the result in the specified register class; the compiler will automatically insert code to store this value in the C variable.

The second constraint specifies input variables and their register classes. The compiler ensures these variables are available in the specified register classes before passing control to the inline assembly code slice.

The third list specifies the actual registers clobbered by the inline assembly code slice. These must be register classes describing a single register.

## Substitutions

A register class may specify a register resource with multiple alternatives. The register allocator selects the actual register used. You can refer to the register resource using substitutions, given by the %N syntax, where N is 0 for the result (if present), 1 for the first input, 2 for the second, and so on:

```
int foo(int xx, int *yy) {
    int out;
    __asm(" ldy %1\n"
          " iny\n"
          " iny\n"
          " lda (%2),y\n"
          " clc\n"
          " adc %1\n"
          " sta %0\n"
          : "=Kdp16"(out)
          : "Kdp16"(xx), "Kdp16"(yy)
          : "c", "y"
    );
    return out;
}
```

Note: If the output is empty then the input list starts with %0.

If you find substitutions using numbers to be unreadable, you can specify a symbol for each substitution:

```
int foo(int xx, int *yy) {
    int out;
    __asm(" ldy %[xx]\n"
          " iny\n"
          " iny\n"
          " lda %[second],y\n"
          " clc\n"
          " adc %[xx]\n"
          " sta %[result]\n"
          : [result] "=Kdp16"(out)
          : [xx] "Kdp16"(xx), [second] "Kdp16"(yy)
          : "c", "y"
    );
    return out;
}
```



The assembler enables writing assembly source files for C projects or standalone assembly language projects.

## 21.1 Overview

Unlike Standard C, there is no standard assembly language. The syntax for assembly instructions used by the 65816 assembler mostly follows what is outlined in guides made by the vendor of the instruction set.

Assembly source expressions are heavily inspired by C-style operators and precedence rules.

The assembler processes assembly source code using the C preprocessor, allowing C comments, include files, and conditional inclusion.

Directives are largely specific to the Calypsi C compiler tool chain, appearing consistent across products from Calypsi. Directive names begin with a leading dot, a common practice in many, but not all, assemblers.

---

**Note:** The dot prefix is used to prevent potential future conflicts where a mnemonic name might clash with a commonly used directive in the product line. This also makes directives stand out in the source code.

---

## 21.2 Syntax

### 21.2.1 Source file format

Assembly source lines follow the traditional format: a label field in the first column, followed by an instruction that may take one or more operands. Comments are preceded by a semicolon:

|                                                              |
|--------------------------------------------------------------|
| <code>[label[:]] [instruction [operands]] [; comment]</code> |
|--------------------------------------------------------------|

A label typically starts in the first column. If leading spaces are present, a label must be followed by a colon or enclosed in back quotes (see Symbol syntax below).

An instruction can be a mnemonic (target instruction name) or a directive. All directives begin with a leading dot as part of their name.

A label, a symbol describing a location, is defined by placing it in the first column of a source line, optionally followed by a colon.

Labels can also be declared by importing them with the `.extern` directive or by value using the `.equib` directive.

Predefined words in instructions (mnemonics and operands) are case insensitive, while symbols are case sensitive. Some examples follow:

```
; Assembler comments start with a semicolon
;
Loop:      dex      ; assembler comment
           nop
           BNE     Loop
```

Predefined words are case-insensitive because both uppercase and lowercase assembly styles are commonly used.

### 21.2.2 Symbol syntax

Symbols are case sensitive and can be of arbitrary length. A symbol starts with a letter or underscore, followed by letters, underscores, and digits. Examples: `_4`, `a_symbol`, `abc`, `Test1`.

Any character is permitted in a symbol, provided the symbol is quoted with a back-tick. Examples: ``another symbol`` and ``Table: 5``.

Intermixing symbols with the same name but different cases is possible, though considered poor style. Symbols in `as65816` are case-sensitive primarily to support the C compiler, which also has case-sensitive symbols. This also encourages consistent mixed-case symbols in assembly source code, which is considered good practice.

### 21.2.3 Preprocessor

The assembler uses a full-featured C preprocessor for input source files. It provides standard C preprocessor features, including header file inclusion, macro expansions, conditional compilation, and C-style comments.

For an introduction with examples on using the C preprocessor<sup>3</sup>, refer to Wikipedia.

See [Predefined macros](#) for available macros, most of which are also relevant to the assembler.

## 21.3 65816 assembly

The use of direct page is quite central on the 65816 target. If you have an operand that cannot be solved at assembly time, the assembler will use the longer absolute addressing. You can prefix an operand with `dp:` to force the shorter direct page addressing, the following example taken from a list file shows how it works:

```
0001      .extern foo
0002 000000 ad...    lda foo
0003 000003 a50a     lda 10
0004 000005 a5..     lda dp:foo
```

---

<sup>3</sup> [http://en.wikipedia.org/wiki/C\\_preprocessor](http://en.wikipedia.org/wiki/C_preprocessor)

Loading individual parts of an address can be done using relocation operators, such as `.word0` and `.word2`, see [Relocation operators](#) for more information. In many 65816 assemblers angle brackets are used for these (< and >). Here they are named somewhat more uniform with directives and a notation that should be fairly easy to read and understand.

```
0001                                .extern foo, bar
0002 000000 a9..                   lda    ##.word0 foo
0003 000002 85..                   sta    dp:bar
0004 000004 a9..                   lda    ##.word2 foo
0005 000006 85..                   sta    dp:bar+2
```

### 21.3.1 Immediate size

The 65816 can operate in 8-bit or 16-bit mode. There are two such modes, accumulator mode and index register mode, which can be set to either 8-bit or 16-bit independently.

Contrary to some 65816 assemblers, the Calypsi assembler does not attempt to keep track of the current runtime mode settings when it processes instructions. Thus, each instruction is treated independently of each other and from the actual mode settings. Properly matching instructions with runtime mode settings is the sole responsibility of the developer.

This mainly affects the immediate addressing mode where the size of immediate operand depends on the active mode settings at runtime. The assembler can emit either 8-bit or 16-bit immediate instruction operands and you specify which by using either a single or double hash mark:

```
0001 000000 a23412                 ldX    ##0x1234      ; 16-bit immediate
0002 000003 a9ab                   lda    #0xAB         ; 8-bit immediate
```

### 21.3.2 Addressing mode range

As mentioned there is a `dp:` facility that can be put before an address expression to specify that the direct page is intended. There are also `abs:` and `long:` that can be used in to describe 16 and 24 bits addressing to be used.

For expressions that can be solved in the assembler, the shortest possible addressing is assumed if no such prefix modifier is present.

```
0001 000000 b50a                   lda    dp:10,x
0002 000002 bd0a00                 lda    abs:10,x
0003 000005 bf0a0000               lda    long:10,x
0004 000009 b50a                   lda    10,x
0005 00000b bd0010                 lda    0x1000,x
0006 00000e bf001033               lda    0x331000,x
```

## 21.4 Sections

The assembly source file is divided into sections using the `.section` directive. Each section, an indivisible unit of code or data, is laid out in memory by the linker according to rules from a placement rules file.

If no `.section` directive is specified, the assembler defaults to a `.section code` in the source file, meaning you start in a section named `code` unless otherwise directed.

Multiple sections enable more flexible placement of code and data by the linker than a single section. A section name can be used multiple times, with each instance creating an individual section fragment.

All sections are relocatable due to the ELF file format and must be defined in the linker rules file.

### 21.4.1 Section kinds

The following table describes the supported section kinds. If unspecified, the section defaults to `text`.

Section kinds and modifiers are case-insensitive.

| Section kind        | Description                                                                                                                                                                               |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>text</code>   | Executable code.                                                                                                                                                                          |
| <code>data</code>   | An initialized read/write data section in memory (RAM).                                                                                                                                   |
| <code>rodata</code> | An initialized read-only data section in memory (ROM).                                                                                                                                    |
| <code>bss</code>    | “Block Started by Symbol”; holds zero initialized variables and variables not given an initializer value. The C runtime normally zero fill such area before calling <code>main()</code> . |

### 21.4.2 Section modifiers

Section modifiers describe further behavior and can be specified as positive or negative.

| Section modifier       | Description                                                                                                                                |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <code>noreorder</code> | Maintains the relative order of section fragments with the same name within a translation unit.                                            |
| <code>reorder</code>   | Allows section fragments of the same name to be placed arbitrarily, regardless of other fragments with the same name. This is the default. |
| <code>root</code>      | Always include this section fragment in the program. This is the default for object files.                                                 |
| <code>noroot</code>    | Only include this section fragment in the program if someone refers to a label inside it. This is the default for library files.           |

---

Note: Section fragments with the same name and the `noreorder` modifier are combined by the linker into a single placement group, ensuring contiguous placement. The absence of `noreorder` (or explicit `reorder`) allows the linker to place each section fragment arbitrarily, regardless of other fragments with the same name.

---

### 21.4.3 Section alignment

The `.align` directive specifies an alignment in address units. It advances the location counter and inserts fillers if needed, ensuring the next location has the specified alignment.

```
; Ensure that "table" label is placed at an address that can be
; evenly divided by 4.
    .section data
    .align 4
table: ...
```

## 21.5 Expressions

Numeric expressions operate in signed (2-complement) mode. A range check occurs based on value usage; exceeding the possible range results in an error. The [Operators](#) table lists standard operators, supplemented by specialized relocation and section operators (see [Relocation operators](#) and [Section operators](#)).

You can optionally use spaces between values and operators in an expression.

Table 21.1: Operators

| Operator | Precedence | Purpose               |
|----------|------------|-----------------------|
| ~        | 9          | Unary bit-wise not    |
| !        | 9          | Unary logical not     |
| -        | 9          | Unary negate          |
| +        | 9          | Unary plus            |
| *        | 8          | Multiply              |
| /        | 8          | Divide                |
| %        | 8          | Modulo                |
| +        | 7          | Add                   |
| -        | 7          | Subtract              |
| <<       | 6          | Bit shift left        |
| >>       | 6          | Bit shift right       |
| >        | 5          | Greater than          |
| <        | 5          | Less than             |
| >=       | 5          | Greater than or equal |
| <=       | 5          | Less than or equal    |
| ==       | 4          | Equal                 |
| !=       | 4          | Not equal             |
| &        | 3          | Bit-wise and          |
| ^        | 2          | Bit-wise exclusive or |
|          | 1          | Bit-wise or           |

## 21.6 Numeric constants

Integer constants can be entered in decimal, binary, octal, or hexadecimal. Decimal integers are sequences of digits not starting with zero. Prefixes `0x` indicate hexadecimal, `0b` binary, and `0` (followed by digits) octal.

Character constants are also supported, replaced by their ASCII value.

Some examples:

```
table:    .byte    0x1ff        ; hexadecimal number
          .byte    077         ; octal number (corresponds to decimal 63)
          .byte    65          ; decimal 65
          .byte    'A'         ; ASCII 65 (decimal)
          .byte    0b1011      ; binary (corresponds to decimal 11)
          .byte    0           ; zero (actually octal, but it is written
                                ; the same way in decimal)
```

## 21.7 Location counter

The assembler converts source programs into machine code for the 65816 processor. Instructions occupy consecutive memory locations until the next `.section` directive or end of file. The current instruction address is accessed by a single period (`.`), often called “dot”.

The dot label, representing the current instruction location, is the location counter. It increments after each instruction, always holding the start address of the current instruction.

The location counter address is resolved by the linker but can be used in expressions. Short branches can use the location counter:

```
test:     ldx      #0
          lda      (table),y
          bpl      .+3          ; skip next instruction of positive
          dex
```

However, using labels or local labels is often preferred (see below).

## 21.8 Local labels

Local labels are useful for local branch destinations in assembly source files and come in two variants, dollar-postfix alphanumeric and plus/minus sign character labels.

### 21.8.1 Dollar postfix style

Local labels end with a single `$`; the name can be an identifier or numeric (e.g., `loop$` or `3$`). A local label is active only between two non-local labels and cannot be exported to the linker. They serve as temporary locations for short-distance branching, typically for skips or local loops.

```
loop$:    lda      (ptr),y
          beq      15$
          iny
          dex
          bne      loop$
```

(continues on next page)

(continued from previous page)

```

15$:      bra    50$
          lda    #7

```

After a non-local label, you can no longer refer to any local label before it. Consider if the following code is added below the one above, the presence of the non-local label `foo` resets the local labels:

```

foo:      ldx    #5
          bne    loop$      ; error, loop$ above no longer visible

```

As an alternative, you can use ordinary labels related to the context and add a number to make it unique. What you do is mostly a matter of taste.

## 21.8.2 Sign style

Local labels can also be created using sequences of `+` or `-` characters. The entire label name must use the same character, and its length is used for matching. References to labels with minus characters go backwards to the closest match, while plus characters go forward to the closest match.

This allows easy determination of whether the destination label is before or after an instruction. Sign-style local labels can pass over non-local labels.

```

+          bne    +          ; this one goes to first + below
--
          beq    --          ; backward
+:          beq    ++++       ; goes over foobar
foobar:
          nop
++++:

```

Note: Sign-style local labels are only usable with branch-style instructions. They are not allowed in more elaborate operands where an ordinary label is allowed.

## 21.9 Directives

All directives start with a single dot character.

The following table summarizes the directives known to the assembler:

| Directive                                         | Purpose                                                     |
|---------------------------------------------------|-------------------------------------------------------------|
| <code>.section section-name, argument-list</code> | Generate code for given section                             |
| <code>.equ expr</code>                            | Define a symbol value                                       |
| <code>.equib expr</code>                          | Define a symbol value that corresponds to a memory location |
| <code>.public symbol-list</code>                  | Export symbols to the linker                                |
| <code>.global symbol-list</code>                  | Synonym to <code>.public</code>                             |
| <code>.globl symbol-list</code>                   | Synonym to <code>.public</code>                             |
| <code>.pubweak symbol-list</code>                 | Export weak symbols to the linker                           |
| <code>.extern symbol-list</code>                  | Import symbol from other module                             |
| <code>.require symbol-list</code>                 | Require symbol from other module                            |
| <code>.incbin filepath</code>                     | Include a binary file as data                               |
| <code>.rtmodel symbol, string</code>              | Define a runtime model attribute                            |
| <code>.byte expr-list</code>                      | Emit 8 bit values                                           |
| <code>.word expr-list</code>                      | Emit 16 bit values                                          |
| <code>.address expr-list</code>                   | Emit 24 bit values                                          |
| <code>.long expr-list</code>                      | Emit 32 bit values                                          |
| <code>.quad expr-list</code>                      | Emit 64 bit values                                          |
| <code>.ascii text</code>                          | Emit the given string                                       |
| <code>.asciz text</code>                          | Emit the given string with a terminating 0 (C string)       |
| <code>.fillto address [, value]</code>            | Fill memory with value                                      |
| <code>.space count [, value]</code>               | Fill memory with value                                      |
| <code>.align value</code>                         | Specify alignment                                           |
| <code>.macro parameter-list</code>                | Define a macro                                              |
| <code>.endm</code>                                | End a macro definition                                      |
| <code>.end</code>                                 | Stop processing the source file                             |

A symbol-list is a list of symbols separated by commas. An expr-list is a list of expressions separated by commas.

### 21.9.1 Directives in detail

#### `.section`

The `.section` directive takes the name of the section as the first argument. It can optionally be followed by a kind and modifiers to describe the section further.

```

; A code section named "code" (text)
    .section code

; A code section named "code" (text) that are not stored
; in relative order to other "code" section fragments in
; the same compilation unit.
    .section code, reorder

; A data section named "storage"
    .section storage, data

; A constant area in ROM that are always included in output,
; even when put in a library (provided that anything
; in the compilation unit is referenced).
    .section table, rodata, root

```

`.equ`

Introduces a new symbol and gives it a value. The symbol appears in the label column and may have an optional colon after it:

```
BufNo      .equ 7
BufSize:   .equ 2 + ContentSize
```

Any expression can be used as a value, however using external symbols is subject to certain limitations imposed by relocations in the ELF object file format. In the case of valid expressions with external symbols, the value will be resolved by the linker.

`.equlab`

Constants can also be defined with the `.equlab` directive. This is similar to the `.equ` directive, with the difference that it defines a label, which describes a location, rather than a plain number:

```
MyLocation: .equlab 0xD7
```

Where this matters is in the interpretation of the debugging information. Labels defined using `.equlab` are treated the same as other location labels. The debugger will understand that a symbol defined using `.equlab` can be used as a location when generating disassembly listings, while symbols defined using `.equ` will not be used for locations in the disassembly listing.

`.public`

Exports a symbol in the current file and makes it visible to other modules. Symbols are local to the file being assembled by default.

For shared definitions, an alternative is to put them in an include file and define them using the `.equ` directive.

`.global`

Synonym to `.public` for compatibility with other assemblers.

`.globl`

Synonym to `.public` for compatibility with other assemblers.

`.pubweak`

A weak symbol is created using the `.pubweak` directive in a similar way to `.public`. The difference is that a weak symbol may exist in multiple copies. Of these potentially multiple copies, one is selected by the linker. If there is a non-weak symbol among the weak ones, it will be picked by the linker.

Weak symbols serve a couple of purposes. They can be used for library replaceable objects where you can override a default library object using a non-weak public symbol. They are also useful for tools that generate assembly code where an identical construct may be generated multiple times, though only one is needed in the end.

`.extern`

Imports a symbol that is defined in some other module and make it visible in the current source file.

`.require`

A required symbol can be specified with the `.require` directive. It works similar to the `.extern` directive, with the difference that you do not need to actually use the symbol in any expression, it is being pulled in by the linker regardless.

This is mostly of interest when building modular software using libraries.

The typical use of this is that you have initialization code somewhere else built up using section fragments with the no-reorder property. The no-reorder property ensures that the code fragments appear next to each other. A section fragment is only active if someone actually refers to it. In this case the `.require` directive can be used to refer to it without actually using it. The result is that code which relies on that initialization code fragment exists can request that such code fragment becomes active.

`.incbin`

Include a binary file as embedded data directly into the program. This is useful if you have binary data assets as you can include them in the program without having to convert them to suitable source code.

```
my_data:
    .incbin "rawdata.bin"
```

`.rtmodel`

It is possible to define runtime model attributes using the `.rtmodel` directive. Such attributes are checked at link time to ensure object file consistency. The attribute is an identifier and its value is a string:

```
; activities suitable for winter
    .rtmodel season, "winter"
```

Another source file could specify `season` to have another value:

```
; activities suitable for summer
    .rtmodel season, "summer"
```

If you try to link these two modules together will result in an error message describing that runtime model attribute `season` has a mismatch.

Source files that do not define the `season` attribute can be linked with either. It is also possible to use the special `*` value which also means it works with either. It is an explicit way of saying that I aware of this attribute and it works with whatever value it has:

```
; I can work with either season
    .rtmodel season, "*"
```

Functionally it is equivalent to not having the attribute defined. The difference is more in the eye of the reader, you have actively considered the attribute and concluded it works with whatever interpretation there may be.

---

Note: A defined runtime attribute affects the entire compilation unit it appears in. If you need to have a more narrow scope for some runtime model attribute, you need to break up the source file into smaller pieces.

---

#### `.byte`

Define one or more 8 bit values expressed as a list of comma separated expressions.

#### `.word`

Define one or more 16 bit values expressed as a list of comma separated expressions.

#### `.address`

Defines one or more 24 bit values expressed as a list of comma separated expressions.

#### `.long`

Define one or more 32 bit values expressed as a list of comma separated expressions.

#### `.quad`

Define one or more 64 bit values expressed as a list of comma separated expressions.

#### `.ascii`

Take a string arguments and emit the bytes.

#### `.asciz`

Take a string arguments and emit the bytes followed by a terminating zero byte.

#### `.fillto`

Take an address offset and an optional filler value. The address value is relative to the start of the current section fragment. Emits the filler value until the location counter is equal address offset.

If the current section is bss, the filler value cannot be non-zero.

`.space`

Take an `count` value and an optional filler value. Emits `count` number of filler values in the output. If the current section is `bss`, the filler value cannot be non-zero and it only advances the location counter by `count`.

`.align`

Emit zero fillers to bring the location counter in alignment with the specified alignment argument. To long word align a table, use:

|                    |                            |
|--------------------|----------------------------|
| <code>table</code> | <code>.align 4</code>      |
|                    | <code>.byte 1,2,3,4</code> |

---

Note: The `.align` directive also have the effect of applying alignment on the section itself in the object file which forces the linker to place the section according to the alignment. If there are multiple `.align` directives with different alignment values in a section, the assembler will determine the overall alignment needed and emit that as the alignment of the section.

---

`.end`

End the source file and stop processing any further input. This directive is optional and processing will otherwise stop when the input file has been fully consumed.

`.macro`

Define a macro, see [Macro language](#).

`.endm`

End a macro definition.

`.argdelim`

This directive defines delimiter characters that can be used include commas in an argument, see [Macro language](#).

## 21.10 Relocations

The assembler and the C compiler produces relocatable output, which means that the program output does not have fixed addresses. The linker is responsible for placing sections at suitable memory addresses.

The compiler may need to refer to addresses that are unknown at compiler time. Such references are represented by relocations which are part of the object file. The linker processes the relocations after section placement to finalize the program.

Relocations are automatically generated by the assembler and C compiler when it encounters expressions or values that depend on final section placement.

### 21.10.1 Relocation operators

In certain contexts some additional prefix operators are available. They can be seen as relocation operators as they can be used to optionally introduce a relocation.

As they are relocations, they allow the operator to be executed at link time when the final addresses are known. However, they have the limitation that they must appear at the top level of the expression.

On the 65816 you often need to load word parts of an address. This can be done with the `.word0` and `.word2` operators. There are also `.byte0`, `.byte1`, `.byte2` operators that return individual bytes of an expression at link time.

```
; Set up a pointer
    lda    #.byte0 table
    sta    zp:ptr
    lda    #.byte1 table
    sta    zp:ptr+1
```

---

Note: Relocation operators have high precedence like other unary prefix operators. If you want to access an address with an offset, you need to surround the expression by parentheses.

---

`.byte0`

The `.byte0` operator gives the first byte (count starts from 0) and is allowed to be resolved at link time.

`.byte1`

The `.byte1` operator gives the second byte (count starts from 0) and is allowed to be resolved at link time.

`.byte2`

The `.byte2` operator gives the third byte (count starts from 0) and is allowed to be resolved at link time.

`.word0`

The `.word0` operator gives the lower 16 bits word and is allowed to be resolved at link time.

`.word2`

The `.word2` operator gives the upper 16 bits word and is allowed to be resolved at link time.

### .tiny

This relocation gives the address on a symbol in the Tiny area, which corresponds to the direct page. Typically it is given an expression to be resolved at link time. The expression is an address that is converted to a relative offset to the linker defined direct page base symbol `DirectPageStart`:

```
0001                      .extern foo
0002  000000 a9..         lda    dp: .tiny (foo+2)
```

### .near

This relocation gives the relative address of a symbol in the Near area, which corresponds to the 16 bits absolute area. Typically it is given an expression to be resolved at link time. The expression is an address that is converted to a relative offset to the linker defined Near base symbol `NearBaseAddress`:

```
0001                      .extern foo
0002  000000 a9....         lda    abs: .near (foo+2)
```

### .kbank

This relocation is used for `jmp` and `jsr` inside the current 64K program bank. It takes an address and strips out the lower 16 bits, in the same way as `.word0` does. In addition it checks that the destination is within the same program bank (K register) as the instruction itself.

```
0000  020034 4c....         jmp    .kbank localLabel
```

## 21.10.2 Section operators

Section operators makes it possible to get hold of where a section is placed in memory.

Section names must be known to the assembler. If you want to refer to a section that is not used otherwise in the current assembly source file, simply declare it without inserting any code below it.

```
;;; Forward declaration
    .section elsewhere

    .section code
    ...
```

All these operators are resolved by the linker as placement is not known before link time. An error is given if a section spans multiple memories.

Table 21.2: Section operators

| Operator                               | Precedence | Purpose                                 |
|----------------------------------------|------------|-----------------------------------------|
| <code>.sectionStart sectionName</code> | 9          | The first address of the given section. |
| <code>.sectionEnd sectionName</code>   | 9          | The last address of the given section.  |
| <code>.sectionSize sectionName</code>  | 9          | The size of the given section.          |

---

Note: Section size corresponds to `1 + end - start`.

---

Warning: If you allow the linker to intermix different sections in the same allocation range, these operators will base their values on the first and last section fragment of the given name. Different sections that are interleaved inside are silently included in the address range given by these operators.

## 21.11 Macro language

The `.macro` directive allows you to generate new commands that can create assembler output. A simple example follows:

```
foo      .macro  a, b
        .byte   \a
        .word   \b - 1
        .long   0
        .endm
```

This creates a new macro named `foo` which takes two arguments `a` and `b`. To use an argument inside the macro, prefix the parameter name with a backslash `\`.

### 21.11.1 Rules for argument substitutions

When looking for argument substitutions, the longest match is favored. This means if you have parameters called `a` and `aa`, substituting the longer name is always tried before shorter names. As there is no way to explicitly specify the end of a parameter name inside the body, a parameter may accidentally try to match characters that comes after the parameter. A good rule of thumb is to make use of space to separate entities whenever possible. This also tends to improve readability.

### 21.11.2 Use of local labels

Each macro expansion will create a new unique context for local labels inside the macro body. Any previous local label context is restored after the macro is expanded. Thus, local labels inside a macro will not clash or interfere with any local labels surrounding the use of the macro.

This also works when using nested macro expansions. If a macro uses another macro inside its body, that inner macro expansion will have its own private local label context, and the previous context of the outer macro expansion will be restored when the inner macro has been expanded.

Thus, you are able to use the same label name inside a macro and in the code that uses it without any clash:

```
waitreg  .macro  reg
1$:      de\reg
        bne    1$
        .endm

        ldx    #100
        waitreg x
1$:      lda    #0
```

Which would create the following list file:

```
#####
#
# Calypsi assembler for 65816                                version 5.16 #
#                                                           14/Apr/2026 16:42:19 #
# Command line: example/macro-local.s -l                      #
#                                                           #
#####

0001          waitreg      .macro  reg
0002          1$:          de\reg
0003                      bne      1$
0004                      .endm
0005
0006 000000 a264          ldx      #100
0007                      waitreg x
    \ 000002 ca      `1$`:  dex
    \ 000003 d0fd          bne      `1$`
0008 000005 a900      1$:      lda      #0

#####
#                               #
# Memory sizes (decimal) #
#                               #
#####

Executable (Text): 7 bytes
```

### 21.11.3 Arguments with comma

Arguments to a macro are comma separated. This poses a problem in a situation where you want an argument to contain a comma character. The `.argdelim` directive defines a start and stop character that can be used to create an argument that contains a comma character.

```
.argdelim <>
access    .macro  arg1, arg2
...
.endm

access    0, <2,a>
```

Here `arg1` is bound to `0` and `arg2` is bound to the value `2,a`.

The delimiter can be either one or two characters and you can pick any suitable character combination. By default there are no delimiter characters defined.

If a single character combination is not suitable, you can use two characters, e.g. `<-` and `->` which would be defined as follows:

```
.argdelim <-->
access    .macro  arg1, arg2
...
.endm

access    0, <-2,a->
```

All delimiter characters are stripped and `arg2` is bound to `2,a` here as well.

---

## Running the assembler

---

The assembler is run from the command line, or from an IDE that interfaces with the assembler using the command line.

### 22.1 Basic invocation

The assembler is invoked in the following way:

```
$ as65816 [options] sourcefile [options]
```

As an example, to assemble a source file with debugging information and a list file, you can use:

```
$ as65816 --debug source.c -l
```

Command line options are optional arguments that tune the behavior of the assembler. They always start with a dash character. There are two variants, single letter options (`-l` to instruct the assembler to create a list file) starts with a single dash. The other variant is long descriptive options that starts with two dashes.

Some options require arguments. These follow the option, separated by a space or an equals sign (`=`). For single-argument options, the separator is optional:

```
$ as65816 -Iinclude source.c -D VERBOSE=2 --list-file=tiny-source.lst
```

In this case the `-I` option adds `include` as a directory to scan for header files. The symbol `VERBOSE` is defined in the preprocessor with value 2. A list file with a specific name is also produced.

The order in which the options appear normally do matter, except for the `-I` option that adds directories to search for header files. Such directories are searched in the order in which they appear on the command line.

To display the version of the assembler, use `--version`:

```
$ as65816 --version
Calypsi assembler for 65816 version 5.16
```

Command line options are described in [Command line options](#).

## 22.2 Include search path

Header files are included by surrounding the filename with either double quotes or angled brackets:

```
#include "myheader.h"
#include <system.h>
```

Angled brackets are typically for system header files, while double quotes are for application-specific header files.

---

**Note:** The installation provides no system header files for the as65816 assembler. However, the preprocessor retains its standard behavior of supporting angled include files and searching the installation directory.

---

The search order for include files is as follows:

1. Relative to compilation directory (double quote include only)
2. In directories specified in the `-I` option in the order they appear on the command line
3. The system directory, which is the Calypsi installation directory

Header files specified in angle brackets are not searched relative to the compilation directory. Otherwise, their behavior is identical.

---

**Note:** More precisely, the system header file directory is relative to the as65816 executable. If you move an installation, it will still find the correct system header file directory.

---

---

**Note:** No header files are currently provided with the assembler. The preprocessor search path exists for consistency with the C compiler.

---

## 22.3 Assembler output

The assembler outputs one or two files, an object file that can be linked with other object files and libraries by `ln65816` to produce an executable file, and optionally a list file.

### 22.3.1 Object file

The object file is in ELF format, optionally including DWARF debugging information if `--debug` (or `-g`) is specified.

It contains:

- A symbol table
- Relocatable sections for code and data
- Relocations, allowing linker to modify address values after section placement
- DWARF-formatted source-level debugging information (if `--debug` was specified)

- Vendor-specific data, including runtime attributes and additional the Calypsi C compiler tool chain-specific information not covered by ELF/DWARF, used by the linker and debugger, with section types 0x8000000a and 0x8000000b.

Note: DWARF version 5 is used. The implementation covers the needs of the Calypsi C compiler tool chain and includes vendor extensions.

Note: DWARF debug information from the assembler includes source line information. The object file also contains a symbol table.

### 22.3.2 List file

The list file is a text file meant to be shown with a fixed width font. It contains a header that shows assembler, version, time when it was created and the command line used. The assembly source is shown with corresponding hex machine code. It also shows macro expansions and includes a summary of code and data sizes.

The following example is the builtin source file from the C runtime library:

```

/*****
 *
 * Copyright Håkan Thörngren
 *
 * This file is part of the Calypsi C library.
 * Permission to use with the Calypsi tool chain is hereby granted.
 *
 *****/

;;; -----
;;;
;;; Builtin support.
;;;
;;; -----

        .rtmodel version, "1"
        .rtmodel cpu, "*"

#include "macros.h"

        .extern _Dp

;;; ****
;;;
;;; _CLZ - count leading zeroes
;;;
;;; In: A - integer value
;;;
;;; Out: A - number of leading zeroes
;;;
;;; Uses: Y
;;;
;;; Result is undefined if A is 0.
;;;

```

(continues on next page)

(continued from previous page)

```

;;; *****
                .public _CLZ
                .section libcode
_CLZ:           tay                ; check for zero
                beq     30$
                ldy     ##-1
10$:           iny
                asl     a
                bcc     10$
20$:           tya
30$:           return

;;; *****
;;;
;;; _CLZL - count leading zeroes
;;;
;;; In: X:A - 32 bit integer value
;;;
;;; Out: A - number of leading zeroes
;;;
;;; Uses: X, Y
;;;
;;; Result is undefined if X:C is 0.
;;;
;;; *****

                .public _CLZL
                .section libcode
_CLZL:         txy                ; upper part non-zero
                bne     10$
                tax                ; lower part zero?
                beq     60$         ; yes, done
                ldy     ##15
                bra     30$
10$:           txa
                ldy     ##-1
30$:           iny
                asl     a
                bcc     30$
50$:           tya
60$:           return

```

Compiling with the -l option produces a list file:

```

#####
#
# Calypsi assembler for 65816                      version 5.16 #
#                                           14/Apr/2026 16:42:19 #
# Command line: builtin.s -l                      #
#                                           #
#####

0001          /*****
0002          *
0003          * Copyright Håkan Thörnengren

```

(continues on next page)

(continued from previous page)

```

0004      *
0005      * This file is part of the Calypsi C library.
0006      * Permission to use with the Calypsi tool chain is hereby granted.
0007      *
0008      *****/
0009
0010      ;; -----
0011      ;;
0012      ;; Builtin support.
0013      ;;
0014      ;; -----
0015
0016      .rtmodel version, "1"
0017      .rtmodel cpu, "*"
0018
0019      #include "macros.h"
0020
0021      .extern _Dp
0022
0023      ;; *****
0024      ;;
0025      ;; _CLZ - count leading zeroes
0026      ;;
0027      ;; In: A - integer value
0028      ;;
0029      ;; Out: A - number of leading zeroes
0030      ;;
0031      ;; Uses: Y
0032      ;;
0033      ;; Result is undefined if A is 0.
0034      ;;
0035      ;; *****
0036
0037      .public _CLZ
0038      .section libcode
0039 000000 a8      _CLZ:      tay          ; check for zero
0040 000001 f008      beq      30$
0041 000003 a0ffff      ldy      #-1
0042 000006 c8      10$:      iny
0043 000007 0a      asl      a
0044 000008 90fc      bcc      10$
0045 00000a 98      20$:      tya
0046 00000b 6b      30$:      return
0047      \ 00000b 6b      rtl
0048      ;; *****
0049      ;;
0050      ;; _CLZL - count leading zeroes
0051      ;;
0052      ;; In: X:A - 32 bit integer value
0053      ;;
0054      ;; Out: A - number of leading zeroes
0055      ;;
0056      ;; Uses: X, Y
0057      ;;
0058      ;; Result is undefined if X:C is 0.

```

(continues on next page)

(continued from previous page)

```

0059                ;;
0060                ;; *****
0061
0062                .public _CLZL
0063                .section libcode
0064 000000 9b      _CLZL:    txy            ; upper part non-zero
0065 000001 d008    bne      10$
0066 000003 aa      tax            ; lower part zero?
0067 000004 f00e    beq      60$      ; yes, done
0068 000006 a00f00    ldy      ##15
0069 000009 8004    bra      30$
0070 00000b 8a      10$:    txa
0071 00000c a0ffff    ldy      #-1
0072 00000f c8      30$:    iny
0073 000010 0a      asl      a
0074 000011 90fc    bcc      30$
0075 000013 98      50$:    tya
0076 000014 6b      60$:    return
    \ 000014 6b      rtl

#####
#                                     #
# Memory sizes (decimal) #
#                                     #
#####

Executable (Text): 33 bytes

```

## 22.4 Command line options

This section covers the as65816 command-line options in detail.

### 22.4.1 Options overview

Running the assembler from the command line without arguments results in a missing input file error, followed by a short help message:

```

$ as65816
Missing: FILE

Usage: as65816 [--version] [-o|--output-file OUTPUT-FILE] [-l]
             [--list-file LIST-FILE] [-I DIRECTORY] [-D IDENTIFIER]
             [-U IDENTIFIER] [-g|--debug] [--rtattr NAME=VALUE]
             [--weak-symbols] [--core CORE] [--target TARGET]
             [--code-model NAME] [--data-model NAME] FILE
    use 'as65816 --help' for detailed help

```

For more detailed help, use the `--help` option:

```

$ as65816 --help
Calypsi assembler for 65816 version 5.16

```

(continues on next page)

(continued from previous page)

```

Usage: as65816 [--version] [-o|--output-file OUTPUT-FILE] [-l]
           [--list-file LIST-FILE] [-I DIRECTORY] [-D IDENTIFIER]
           [-U IDENTIFIER] [-g|--debug] [--rtattr NAME=VALUE]
           [--weak-symbols] [--core CORE] [--target TARGET]
           [--code-model NAME] [--data-model NAME] FILE
use 'as65816 --help' for detailed help

Available options:
--version           Display version number
-o,--output-file OUTPUT-FILE
                    Name of output file
-l                 Generate a list file, named by appending '.lst' to
                    input file
--list-file LIST-FILE
                    Generate list file, using given name
-I DIRECTORY       Include directory
-D IDENTIFIER      Predefine a macro
-U IDENTIFIER      #undef a predefined macro
-g,--debug         Produce debugging information
--rtattr NAME=VALUE
                    Define a runtime attribute (identifier or quoted
                    string value accepted)
--weak-symbols     Make all public symbols entries weak
--core CORE        Core, one of '65816' (defaults to '65816')
--target TARGET    Target system, one of 'C256', 'F256' or 'SNES'
                    (defaults to embedded/ROM use, if omitted)
--code-model NAME  Code model, one of 'small', 'compact' or 'large'
                    (defaults to 'large')
--data-model NAME  Data model, one of 'small', 'medium', 'large' or
                    'huge' (defaults to 'small')
-h,--help         Show this help text

```

## 22.4.2 Options in detail

### --version

Displays the name and version of the assembler.

### --output-file, -o

Specify the output object file name. If omitted, the output file is derived from the input file name (excluding path) with a .o extension, written to the current directory by default.

This option can also alter the output file name and provide a directory path. The specified directory must already exist.

`-l`

Generate a list file. The name used is the name of the input file (ignoring any directory path) with a `.lst` file extension. See also `--list-file`.

`--list-file`

Generate a list file. The name of the list file is given as argument to this option. See also `-l` to generate a list file based on the source filename.

`-I`

Add a directory to the current include search path. This option can be used multiple times on a command line. The order in which they appear specifies the search order between the directories.

The system include directory is always added last to the search order list.

`--include-system`

Add a directory to the current system include search path before the provided system include directories. This option can be used multiple times on a command line. The order in which they appear specifies the search order between the directories.

`--include-system-after`

Add a directory to the current system include search path after the provided system include directories. This option can be used multiple times on a command line. The order in which they appear specifies the search order between the directories.

`-D`

Define a macro. This takes an argument with the symbol name and optionally an assignment value `-Dsymbol[=value]`. If no value is given, the macro is given the value 1.

`-U`

Undefine a macro. This takes an argument with the symbol name to be undefined.

`--debug`

Generate DWARF symbolic debugging information. In order to get debugging information all the way to the debugger, the linker must also be given this option.

When debugging information is enabled the `__CALYPSI_DEBUG__` macro is also defined and set to 1.

`-g`

Synonym for `--debug`.

`--rtattr NAME=VALUE`

This defines a runtime attribute, written to the object file, that the linker can use for object file consistency checks.

`--weak-symbols`

Make all public symbols in the object file weak. This is normally not needed, but can provide a default library function implementation that is used if no replacement is provided.

In the assembler you can also specify symbols to be exported weak individually using the `.pubweak` directive.

`--core`

This option is mostly available for symmetry with other tool chains from Calypsi and future expansion. At the moment only the 65816 core is recognized.

`--target`

This option is available for symmetry with the compiler. It has the effect of setting the corresponding target preprocessor macro.

`--code-model`

This option is available for symmetry with the compiler. It has the effect of setting the corresponding code model preprocessor macro.

`--data-model`

This option is available for symmetry with the compiler. It has the effect of setting the corresponding data model preprocessor macro.



The linker combines object files from the C compiler and assembler into an executable. Object files can also originate from libraries, such as the provided C runtime library.

Cross compilers need information about the target system, and these systems can differ significantly. A host linker, in contrast, is already tailored for its system.

## 23.1 Running the linker

Like other tools, the linker is command-line based and can also be used from an IDE. The IDE runs the linker via its command-line interface.

### 23.1.1 Basic invocation

The linker is invoked in the following way:

```
$ ln65816 [options] [object-files] [library-files] rules-file
```

A rules file is required to describe the target. This `.scm` file contains information about the target system's memory and section placement rules.

Command-line options and input files can appear in any order.

Command-line options are optional arguments that tune linker behavior. They always start with a dash. Two variants exist: single-letter options (e.g., `-l` for a list file) begin with a single dash, while long descriptive options start with two dashes.

Some options require an argument, which appears after the option. It can be separated by a space or an equal (=) sign. For single-argument options, the argument may appear without a separator:

```
$ ln65816 --debug file.o clib-lc-sd.a placement.scm -o rocket.elf
```

In this example debugging information is retained in the output executable which is also named `rocket.elf`.

To display the version of the linker, use `--version`:

```
$ ln65816 --version
Calypsi linker for 65816 version 5.16
```

## 23.2 Linking process

The goal of the linker is to combine object files produced by the compiler and the assembler. Object files have the file extensions `.o`. From now on we will refer to object files as “objects” as that is how they are seen by the linker once have been read from the file system.

The linker selects and combines objects in several stages:

1. Objects are selected for inclusion in the application.
2. Objects are placed at memory addresses by the linker following placement rules.
3. After placement the final address of each object is known. Using this knowledge the linker resolves relocations. This fixes all address references in the application.
4. The final program is emitted as an executable file.

## 23.3 Simplified placement rules

Understanding the linker process can be daunting due to its flexibility and numerous operations. Fortunately, it can often be used without in-depth knowledge, as it automatically derives placement rules from internal tool knowledge.

The linker cannot know available memory ranges or their types. This section describes placement in a simplified way; for more details, refer to [Memory](#).

### 23.3.1 Minimal rules

A minimal (or simplified) linker rules file primarily specifies memory locations:

```
(define memories
  '((memory flash (address (#x8000 . #xffff)) (type ROM))
    (memory RAM   (address (#x0100 . #x7fff)) (type RAM))
  ))
```

Each memory has a name, address range, and type.

The available types are shown in the following table.

| Memory type         | Description                                                                                                                                                                 |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RAM                 | This roughly corresponds to <code>bss</code> , and also includes the <code>data</code> area in an embedded system.                                                          |
| ROM                 | This is a combination of executable code, read-only variables, and everything else that should go into read-only memory, e.g. switch tables, string literals and constants. |
| ANY                 | This can be used in a hosted environment when the application is loaded into a single RAM memory. It combines allocation of RAM and ROM sections in the same memory range.  |
| <code>rodata</code> | This type describes constants that are stored in read-only memory.                                                                                                          |
| <code>data</code>   | Initialized data area. Only exists when building for hosted use.                                                                                                            |
| <code>text</code>   | This type describes executable code.                                                                                                                                        |
| <code>bss</code>    | This is the data area for an embedded system or bare metal system. In a hosted system this holds the zero initialized data.                                                 |

An error message will be given if the linker fails to automatically bind sections to memory, indicating the sections which are causing problems. In such case you can bind those sections to a memory using a [section rule](#). See [Section placement](#).

In some cases you may have areas that need specific placement. This may be due to API jump tables or entry points. In such case you can explicitly place such sections using a section rule with an explicit address.

---

Note: If you have a special memory, such as a video memory that requires special access methods and cannot handle normal code or data, you may want to prevent the linker from doing automatic section binding there. This can be done by not specifying any `type` for it. In such case only sections with explicit binding will be placed in that memory.

---

## 23.4 ROM based systems

A ROM based system is typically an embedded or bare-metal system that runs without an operating system, taking control at power-on. The Calypsi linker assumes a ROM based system unless configured otherwise.

On a ROM based system, the application provides the reset vector. The C runtime configures itself by initializing static variables, setting up the stack, the heap (if used), and the C runtime itself. The I/O system is also initialized, but requires you to provide stubs for low-level stream-style I/O operations; see [Stubs interface](#) for more information.

If the hardware for your application has specific initialization needs, these must be provided by the application.

Initializing variables are done by clearing memory and filling in non-zero initializers using information stored in the ROM. The linker will split the initialized variables into two parts: the value part is stored in ROM and is copied to RAM during initialization before `main()` is called.

## 23.5 Host based systems

The alternative to a ROM based system is a hosted system that uses an operating system to load the application. Such operating systems vary from simple kernels to more elaborate systems.

Support for some hosted systems is provided, but given the vast number of alternatives, you may need to provide your own support.

The `--target` and `--hosted` command-line options configure the linker for hosted builds. Selecting a suitable application output file format is also part of this process.

When building for a hosted system, variables are initialized by loading their values in place. This avoids the duplication found in ROM based systems where initializers are stored in ROM and copied to RAM. The downside is that initialization occurs only once, during loading. If the application restarts, variables are not reinitialized. Linker options can override this behavior for detailed control, if needed.

See [Target specifics](#) for details about target specific support for supported hosts.

### 23.5.1 Simple hosted placement

The combination of an often simple executable file format and an application being loaded into RAM with both code and data gives certain problems. An application consists of both “bits” areas (code, tables and constant data) and “nobits” which is typically zero initialized data areas. In addition there are “data” areas where initializers are loaded in place, so they are a form of “bits” memory as well. Due to limitations in executable file formats, it is desirable to group “bits” and “nobits” separated, as the latter is often not represented in the executable file.

This can be solved by using a combined memory of type “any”. The linker creates placement groups in that memory for different areas, one for each memory type. Placement is tailored for the executable file format.

#### Specific address placement

Using placement groups clashes with restricting a section to a specific address or limited memory area. If bits sections are placed first, then nobits sections, and there are also sections with address restrictions in the same memory, conflicts can arise.

To avoid conflicts, the linker forbids section rules with arbitrary restrictions when placement groups are used. However, some specific placement may be desirable, such as stub code placed first in memory. Such fixed address placement is allowed, but it must be specified in the memory that contains the placement groups.

## 23.6 Objects

This section looks at the objects being linked, how they are selected and how to control it.

### 23.6.1 C startup

In a C program the `main()` function starts the application. Before `main()` is called, code initializes the C runtime. This prebuilt system startup can be customized in rare cases; see [C startup](#) for details.

Although the C startup is a small module, it is intricate. Knowledge of how Calypsi handles its C runtime is necessary to modify it.

In general the C startup handles hardware initialization, stack setup, heap initialization and configuring the I/O system. It may also do things related to the `exit()` handling.

The startup code is written so that subsystems which are not used by the application are omitted. This selection process is a combination of the C startup, the linker and the C library.

### 23.6.2 Libraries

Multiple object files can be combined into a library file with a `.a` extension. This extension stands for archive, the common name for link libraries on UNIX.

The standard C library is provided as a library file.

---

Note: There are several variants of such library files provided to handle different compiler settings. You will normally not need to specify it as the linker in almost every situation is able to pick a suitable library automatically based on the provided object files.

---

### 23.6.3 Selected objects

Objects are selected from object files specified on the linker command line. All mentioned object files are included in the linker process.

Libraries also contribute objects as needed. The linker includes only objects from a library that are actually used by the application, which helps reduce the application size.

### 23.6.4 Object selection process

The linker selects object based on a tree shake algorithm. This is a worklist algorithm that starts with a root object. The root object references one or more external symbols. These symbols are placed in the work list. The linker will then process the first symbol in the work list and try to fulfill it using the available objects. Doing so may produce further external symbol which are placed in the work list unless they have already been processed.

This process continues until the work list is empty. If there are unresolved symbols they will be emitted as undefined symbol error messages.

When all symbols have been processed, the linker knows which objects are needed for the application.

### 23.6.5 The root

The actual root is a symbol named `__program_root_section` which identifies the section that must be included first. For a ROM based bare metal system, this is typically the section that holds the reset vector. For an operating system based system startup, it may be some stub or simply the section that contains the first executable code.

---

Note: If you study the system startup code you will find the `__program_start` symbol which can be seen as a close relative to the `__program_root_section` symbol. The purpose of the `__program_start` symbol is to tell the debugger about the first executable instruction of the application. It has no special meaning to the linker.

---

### 23.6.6 Placement

The actual placement of objects is described in [Section placement](#).

### 23.6.7 Relocations

Consider a function in your application. This is described by a symbol with the same name as the function. In order to make a function call, the compiler needs to emit an instruction to transfer control to the function. How can the compiler know the address of the function when it does not know where the function is located? The answer is that it emits a relocation, which instructs the linker to adjust the application code when the final address is known.

Somewhat simplified, the relocation describes a location inside the object code that has to be altered. The called function is mentioned by its symbol. With this relocation information, the linker can figure out the location in the application to alter, and what to put there, to allow the actual call to be correctly made.

Relocating the program can be done by the linker once all objects have been given their final locations in the memory space of the application.

---

Note: In some supported operating systems relocations are only done partially by the linker and a further relocation is done when the operating system loads the application to memory.

---

## 23.7 Memory

When describing the target system, the linker uses memories. A memory is a named entity specifying a continuous range of storage locations. Memory descriptions have the following attributes:

**name** The memory name identifies a particular memory area.

**address range** The start and end addresses (inclusive) specify the memory's address range.

**section** Names of sections to bind to the current memory.

**placement group** A memory can only hold sections of one type; mixing code with zero- initialized data within a single memory is disallowed. In some output formats, it may be desirable to place different types of areas immediately adjacent within a single memory range. Placement groups function as sub-memories within a parent memory, each capable of holding sections of a distinct type. For example, one placement group can hold code sections, another data sections, and a third BSS sections. The

parent memory defines the overall range, and placement groups allocate memory within this range flexibly, allowing one group to begin immediately after the previous one concludes.

**type** Specifies which section types can be placed in this memory. See [Minimal rules](#) for more information.

**qualifier** Defines the address spaces that can be placed in this memory.

**fill word** Value used to fill unused locations in the memory, defaults to zero.

## 23.8 Linker rules file

The rules file (extension `.scm`) describes the available memory areas and defines rules for section placement. Sections can be explicitly bound to memory areas, or implicitly placed by omitting a binding rule. Implicit placement requires specifying a type for the memory; see [Minimal rules](#).

You can often use a simplified rules file as described in [Simplified placement rules](#). This section details linker rules. These rules can be elaborate and often unnecessary. You can intermix simplified and detailed rules to gain control without over-complicating the file. This approach helps you balance memory and placement requirements.

A full rules file can look as follows:

```
(define memories
  '((memory flash (address (#x8000 . #xffff))
    (section code switch idata izpage
      cdata data_init_table (reset #xfffc)))
    (memory directPage
      (address (#x0 . #xff))
      (section registers zzpage zpage))
    (memory RAM (address (#x0100 . #x7fff))
      (section stack heap data zdata))
    (block stack (size #x800)) ; machine stack size
    (block heap (size #x800)) ; heap size
  ))
```

The use of spacing for indentation here is unimportant to the linker tool, it is used it to make the file easier to read.

### 23.8.1 Section placement

Sections are bound to a specific memory by being mentioned after the `section` keyword. Placement of a section may be further restricted as described below.

#### Free placement

The `code` and `switch` sections appear alone (no surrounding parentheses). Sections with these names can be placed anywhere in the memory they belong to.

### Restricted placement

The `stack` section appears with a specified range (`(#x100 . #x1ff)`). If the range is larger than the size of the section, the section is placed somewhere inside the allowed range.

### Fixed placement

The `reset` section is placed at the fixed address `0xfffc` by having only a given address, not a range.

## 23.8.2 Block rule

Most sections are created by the compiler or the assembler and passed to the linker in the object file. It is also possible to create a section to represent a memory block in the linker rules file. Such blocks are used for the stack and the heap.

A block rule is defined among the memory rules. It has a name which is its section name and a size:

```
(define memories
  '(memory flash (address (#x8000 . #xffff))
    ...
    (block stack (size #x800))
    (block heap (size #x800))
```

You can omit mentioning blocks as the linker will attempt to create them automatically if left out. They are typically used for dynamic memory areas, such as heap and stack. There are some of these built in and the desired size can be specified on the command line. There is also a default size that the linker will fall back if it is not specified.

## 23.8.3 Stack size

The stack keeps track of the dynamic execution state, which includes how to return from function calls and local variables that are not held in registers. It also provides temporary storage during execution. The size needed depends on the application and the stack size can be defined in the linker rules files using a `block` rule:

```
(define memories
  '(memory LoRAM (address (#x4000 . #xefff))
    (section stack data zdata))
  (block stack (size #x1000))
))
```

Here the stack size is set to `1000` hexadecimal (4096 bytes decimal) and is bound to the memory area named `LoRAM`.

---

Note: You can override a stack size defined in the linker rules file using the `--stack-size` command-line option. This is often the preferred method.

---

### 23.8.4 Heap size

The heap is the area from which memory is obtained when using library functions such as `malloc()`. You only need to define a heap if your application actually uses functions such as `malloc()` or `calloc()`. It is perfectly fine for an application to run without a heap, and it may in some cases be desirable to do so.

If you decide to use a heap in your application you need to define its size in the linker rules file using a `block` rule:

```
(define memories
  '( (memory LoRAM (address (#x4000 . #xefff))
      (section stack heap data zdata))
    (block stack (size #x1000))
    (block heap (size #x2000))
  ))
```

Here the heap size is set to `2000` hexadecimal (8192 bytes decimal) and is bound to a memory area named `LoRAM`.

Note: You can override or define the heap size in the linker rules file using the `--heap-size` command-line option. This is often the preferred method.

### 23.8.5 Base address

A base address is a memory region permanently pointed to by a 65816 CPU register. Using base addressing for object access is typically more efficient due to optimized addressing modes on the 65816.

Defined in the linker rules file via a `base-address` rule, it works by associating a special symbol with the actual memory area it covers.

The 65816 defines two base addresses, one for the tiny area which corresponds to the direct page and one for the bank register which describes a single 64K bank.

They are defined in the linker rules file in the following way:

```
(define memories
  '( ...
    (memory NearRAM (address (#x20000 . #x2ffff))
      (section znear))
    (memory DirectPage (address (#xf000 . #xf0ff))
      (section (registers ztiny)))
    ...
    (base-address _DirectPageStart DirectPage 0)
    (base-address _NearBaseAddress NearRAM 0)
  ))
```

The `_DirectPageStart` symbol is tied to the `DirectPage` memory using an offset of zero (meaning start of the memory area). here it is given the address value `00f000`.

The `_NearBaseAddress` symbol is tied to the `NearRAM` memory using an offset of zero (meaning start of the memory area). Here it is given the address `20000` hexadecimal, which is bank 2.

In both cases the C startup module takes care of setting up the corresponding CPU registers before the `main()` function is called.

### 23.8.6 Placement groups

ELF output format limitations prevent mixing section types within a single memory. However, placement groups allow flexible organization.

When an application changes, varying section sizes within a memory range can be cumbersome, requiring frequent adjustments to memory ranges. Placement groups resolve this by dynamically allocating within the overall parent memory address range.

A placement group acts as a sub-memory within a parent, each holding sections of a single type. Multiple placement groups of different types can exist within one parent memory. Sections are filled one at a time. Once a group is placed, the linker continues allocation from the next free address, maximizing memory utilization without manual range specification for each group.

You can define placement groups in the following way:

```
(define memories
  '(...
    (memory near-bank (address (#x10000 . #x1ffff))
      (placement-group near-bits (section near cnear))
      (placement-group near-nobits (section znear)))
```

In this example the `near-bank` memory contains two placement groups `near-bits` and `near-nobits`. The `near-bits` placement group contains two sections `near` and `cnear`. The `near-nobits` placement group has a single section `znear` which is a zero initialized BSS section.

---

Note: The use of `near` and `cnear` makes sense when compiling for a hosted environment where a writable `near` section and a read-only `cnear` section can be placed next to each other as both are loaded into system RAM memory. On a ROM based embedded system they would go into separate memories.

---

---

Note: The names `near-bits` and `near-nobits` have no special meaning to the linker.

---

## 23.9 Scatter

In some situations you have a memory bound to a given address range at runtime, but that is stored at a different address in the application image.

Examples of this are an overlay system, a banking system, or a ROM that is located at some address, but parts of it is copied to RAM at a different location.

The linker implements this using `scatter-to` which is a property that can be given to a memory:

```
(memory bankSlotRAM
  (address (#xa000 . #xbfff))
  (scatter-to RAM-banks)
  :generate-instances
  (section bankedcode))
```

In this case the `bankSlotRAM` memory has an address range `A000-BFFF`. Sections of type `bankedcode` are bound to it.

Using the `scatter-to` property, the entire memory is converted to a new section named `RAM-banks`. This section can be placed in another memory:

```
(memory bankedRAM (address (#x2000 . #x1fffff))
 (section RAM-banks))
```

The optional `:generate-instances` is useful when implementing an overlay system. It indicates that once the `bankSlotRAM` memory is full, a new empty instance can be generated to accommodate further `bankedcode` sections. Thus, memory is created on the fly and each instance is scattered to the `bankedRAM` memory.

If `:generate-instances` is not specified then only one instance of the memory is allowed.

Note: Relocations in the memory are done towards the runtime address, which is `bankSlotRAM` in the example above. This is the address area where the memory is intended to appear at runtime.

## 23.10 Linker list files

A list file can be generated from the linker using the `--list-file` (or `-l`) command-line option.

The list file is a valuable tool if you want to see how much memory your application actually needs. It can also answer questions about what objects are included from libraries, why they are included and where they are placed.

This is especially useful if want to understand how your application uses memory and can give insights to have it may be tuned.

As an example, consider the following example program:

```
#include <calypsi/stubs.h>

__task int main () {
    return 0;
}
```

This is intended to be a minimal program, but with debug support to allow the debugger to catch the call to `exit()`.

If built and linked with the standard library and you specify the command-line options `"-l --cross-reference --rtattr exit=simplified"` you will get a list file which contains something like:

```
#####
#                                     #
# Memories summary #
#                                     #
#####
```

| Name           | Range         | Size   | Used   | Checksum | Largest unallocated |
|----------------|---------------|--------|--------|----------|---------------------|
| LoRAM-tiny     | 004000-0040ff | 000100 | 7.8%   | none     | 0000ec              |
| LoRAM          | 004100-00efff | 00af00 | 9.1%   | none     | 009f00              |
| > LoRAM-nobits | 004100-0050ff | 001000 | 100.0% | none     | 009f00              |
| LoCode         | 00f000-00ffe3 | 000fe4 | 0.9%   | none     | 000fdc              |
| > LoCode-code  | 00f000-00f023 | 000024 | 100.0% | none     | 000fdc              |
| LoCode-vector  | 00ffe4-00ffff | 00001c | 7.1%   | none     | 000ffc              |
| flash          | 010000-01ffff | 010000 | 0.0%   | none     | 00ffe5              |
| NearRAM        | 020000-02ffff | 010000 | 0.0%   | none     | 010000              |
| CompactCode    | 040000-04ffff | 010000 | 0.0%   | none     | 010000              |

(continues on next page)

(continued from previous page)

```
#####
#
# Sections summary #
#
#####
```

| Name      | Range         | Size   | Memory        | Fragments |
|-----------|---------------|--------|---------------|-----------|
| registers | 004000-004013 | 000014 | LoRAM-tiny    | 1         |
| stack     | 004100-0050ff | 001000 | LoRAM-nobits  | 1         |
| code      | 00f000-00f023 | 000024 | LoCode-code   | 2         |
| reset     | 00fffc-00fffd | 000002 | LoCode-vector | 1         |
| farcode   | 010000-01001a | 00001b | flash         | 5         |

```
#####
#
# Placement rules #
#
#####
```

| Name             | Address range | Key                    |
|------------------|---------------|------------------------|
| CompactCode      | 040000-04ffff |                        |
| > compactcode    |               |                        |
| LoCode-vector    | 00ffe4-00ffff | Plain                  |
| > (reset 00fffc) |               |                        |
| LoCode           | 00f000-00ffe3 | Plain                  |
| LoCode-cbits     | 00f000-00ffe3 | Plain and RODATA       |
| LoCode-code      | 00f000-00ffe3 | PlainFunction and TEXT |
| > code           |               |                        |
| LoRAM            | 004100-00efff | Plain                  |
| LoRAM-nobits     | 004100-00efff | Plain and BSS          |
| > stack          |               |                        |
| LoRAM-bits       | 004100-00efff | Plain and DATA         |
| LoRAM-tiny       | 004000-0040ff | Tiny and BSS           |
| > registers      |               |                        |
| NearRAM          | 020000-02ffff | DATA                   |
| NearRAM-NoBits   | 020000-02ffff | BSS                    |
| flash            | 010000-01ffff | TEXT and RODATA        |
| > farcode        |               |                        |

```
Name  Size  Align
-----
stack 001000 no
```

| Name             | Memory     | Offset |
|------------------|------------|--------|
| _DirectPageStart | LoRAM-tiny | 000000 |

```
#####
#
# Object files #
#
#####
```

(continues on next page)

(continued from previous page)

```
#####
Unit Filename      Archive
-----
 0  main.o          -
    > farcode 000004
 2  cstartup.o      clib-lc-sd.a
    # picked based on cstartup=normal (built-in default)
    > code 000024
    > farcode 000001
    > reset 000002
 4  simplified_exit.o clib-lc-sd.a
    # picked based on exit=simplified (specified on the command line)
    > farcode 000005
 6  debug_exit.o    clib-lc-sd.a
    # picked based on stubs=semi_hosted (due to configured to use semi-hosting)
    > farcode 000010
 7  debug_break.o   clib-lc-sd.a
    # picked based on stubs=semi_hosted (due to configured to use semi-hosting)
    > farcode 000001
 8  pseudoRegisters.o clib-lc-sd.a
    > registers 000014

#####
#                      #
# Cross reference #
#                      #
#####

__program_start in section 'code'
placed at address 00f000-00f018 of size 000019
(cstartup.o (from clib-lc-sd.a) unit 2 section index 2)
Defines:
  __program_start = 00f000
References:
  _DirectPageStart
  .sectionEnd(stack)
  _Vfp in (pseudoRegisters.o (from clib-lc-sd.a) unit 8 section index 2)
  __low_level_init in (cstartup.o (from clib-lc-sd.a) unit 2 section index 8)
Referenced from:
  __program_root_section (cstartup.o (from clib-lc-sd.a) unit 2 section index 7)

Section 'code' placed at address 00f019-00f023 of size 00000b
(cstartup.o (from clib-lc-sd.a) unit 2 section index 6)
References:
  exit in (simplified_exit.o (from clib-lc-sd.a) unit 4 section index 2)
  main in (main.o unit 0 section index 2)

__program_root_section in section 'reset'
placed at address 00fffc-00fffd of size 000002
(cstartup.o (from clib-lc-sd.a) unit 2 section index 7)
Defines:
  __program_root_section = 00fffc
References:
  __program_start in (cstartup.o (from clib-lc-sd.a) unit 2 section index 2)
```

(continues on next page)

(continued from previous page)

```

Section 'stack' placed at address 004100-0050ff of size 001000
(linker generated)

_Dp in section 'registers' placed at address 004000-004013 of size 000014
(pseudoRegisters.o (from clib-lc-sd.a) unit 8 section index 2)
  Defines:
    _Vfp = 004010
    _Dp = 004000
  Referenced from:
    __program_start (cstartup.o (from clib-lc-sd.a) unit 2 section index 2)
    _Stub_exit (debug_exit.o (from clib-lc-sd.a) unit 6 section index 2)

_Stub_exit in section 'farcode' placed at address 010000-01000f of size 000010
(debug_exit.o (from clib-lc-sd.a) unit 6 section index 2)
  Defines:
    _Stub_exit = 010000
  References:
    _DebugBreak in (debug_break.o (from clib-lc-sd.a) unit 7 section index 2)
    _Dp in (pseudoRegisters.o (from clib-lc-sd.a) unit 8 section index 2)
  Referenced from:
    exit (simplified_exit.o (from clib-lc-sd.a) unit 4 section index 2)

exit in section 'farcode' placed at address 010010-010014 of size 000005
(simplified_exit.o (from clib-lc-sd.a) unit 4 section index 2)
  Defines:
    exit = 010010
  References:
    _Stub_exit in (debug_exit.o (from clib-lc-sd.a) unit 6 section index 2)
  Referenced from:
    (cstartup.o (from clib-lc-sd.a) unit 2 section index 6)

main in section 'farcode' placed at address 010015-010018 of size 000004
(main.o unit 0 section index 2)
  Defines:
    main = 010015
  Referenced from:
    (cstartup.o (from clib-lc-sd.a) unit 2 section index 6)

_DebugBreak in section 'farcode'
placed at address 010019-010019 of size 000001
(debug_break.o (from clib-lc-sd.a) unit 7 section index 2)
  Defines:
    _DebugBreak = 010019
  Referenced from:
    _Stub_exit (debug_exit.o (from clib-lc-sd.a) unit 6 section index 2)

__low_level_init in section 'farcode'
placed at address 01001a-01001a of size 000001
(cstartup.o (from clib-lc-sd.a) unit 2 section index 8)
  Defines:
    __low_level_init = 01001a
  Referenced from:
    __program_start (cstartup.o (from clib-lc-sd.a) unit 2 section index 2)

#####
#                                     #

```

(continues on next page)

(continued from previous page)

```
# Memory sizes (decimal) #
#                               #
#####

Executable      (Text):   65 bytes
Non-initialized      : 4116 bytes
```

The linker list file displays how section fragments are distributed across memories and the total allocated memory size.

The **Memories** summary area summarizes the application's memories, including the size, usage, largest unallocated block, and an optional checksum of each memory.

The **Sections** summary area summarizes section placement and fragment counts.

The **Object files** area lists object files, their originating libraries (if any), and their section and overall size contributions.

The **Cross reference** area provides detailed information for each code segment, including its entry point, name, address range, size, defined symbols, referenced symbols, and referencing code segments. This offers full bi-directional reference information for all application code segments.

Finally, the **Memory sizes** area summarizes the total memory used. Note that memory amounts are given in decimal here; elsewhere in the list file hexadecimal numbers are used.

## 23.11 Output formats

Additional executable output format can be selected with the `--output-format` command-line option.

---

Note: The linker will always produce an ELF output which is the format used by the db65816 debugger.

---

### 23.11.1 ELF/DWARF

This is the base output format of the executable file. It contains an ELF program image of the executable application. If you specify `--debug` all DWARF debugging information sections found in the object files are passed on to the final executable image. The output file has file extension `.elf`.

### 23.11.2 Intel hex

The Intel hex file format contains the output binary as ASCII text. Output files have the `.hex` extension.

### 23.11.3 Motorola S-record

This is the Motorola S-record, sometimes called SREC. The output binary is expressed in ASCII text form. The output file has file extension `.srec`.

The S-record format has variants: S19, S28, and S37, corresponding to 16, 24, and 32-bit addresses. Using these formats forces the linker to output in a specific address size. The output file extension used corresponds to the format variant chosen, `.s19`, `.s28` or `.s37`.

### 23.11.4 Raw

The `raw` format represents program memory as a raw binary image file. Output files use the `.raw` file extension.

---

Note: In the default setting only one memory area can be emitted in Raw format. Multiple areas can be emitted by specifying the `--raw-multiple-memories` command-line option, this also has the effect of changing the name of the output file(s). When this option is active the output file names are based on the name of the memory areas in the `.scm` file.

---

## 23.12 More on linker rules

The linker rules file is a Scheme programming language source file, hence the `.scm` extension.

You do not need to be familiar with Scheme to specify linker rules, simply follow the examples to set things up.

The examples present a Scheme “program” with a single `memories` variable bound to a data structure. The linker reads the file with a Scheme interpreter, then examines the `memories` variable’s contents. This has two implications: a) the file’s syntax is Scheme-dictated, based on [s-expressions](https://www.wikipedia.org/wiki/S-expression)<sup>4</sup>; and b) more elaborate programs using Scheme macros are possible. generate the memory rules.

## 23.13 Command line options

This section details `ln65816` command-line options.

### 23.13.1 Options overview

If the linker is run without command-line arguments, it will indicate that object files are required:

```
$ ln65816
the linker requires some object files to work with (use --help for help)
Terminating due to errors
```

A more detailed help message can be requested with the `--help` option:

---

<sup>4</sup> <https://www.wikipedia.org/wiki/S-expression>

```
$ ln65816 --help
```

```
Calypsi linker for 65816 version 5.16
```

```
Usage: ln65816 [--version] [-o|--output-file OUTPUT-FILE] [-g|--debug]
        [--semi-hosted] [--no-data-init-table-section]
        [--override IDENTIFIER] ([--cross-reference] |
        [--no-cross-reference]) [--rtattr NAME=VALUE] [--verbose] [-l]
        [--list-file LIST-FILE] [--memories-expression EXPRESSION]
        [--program-root SYMBOL] [--root-symbol SYMBOL]
        [--program-start SYMBOL] [--copy-initialize SECTION]
        [--no-copy-initialize SECTION] [--no-automatic-placement-rules]
        [--raw-multiple-memories] [--no-merge-raw-memories]
        [--force-output] [--no-auto-libraries] [--cstartup VALUE]
        [--no-tree-shaking] [--output-format FORMAT] [--stack-size SIZE]
        [--heap-size SIZE] ([--hosted] | [--rom-code])
        [--initialize-large-data] [--core CORE] [--target TARGET]
        [FILE...]
use 'ln65816 --help' for detailed help
```

```
Available options:
```

```
--version           Display version number
-o,--output-file OUTPUT-FILE
                    Name of output file
-g,--debug          Produce debugging information
--semi-hosted       Enable debug stubs for semi-hosting
--no-data-init-table-section
                    Do not generate any data_init_table section (mainly
                    useful for assembly projects)
--override IDENTIFIER
--cross-reference    Override symbol in archive library (treat it as weak)
                    Include cross reference and map information in the
                    list file (this is the default)
--no-cross-reference
                    Do not include cross reference and map information in
                    the list file
--rtattr NAME=VALUE  Specify runtime attribute to select specific
                    alternative from library (e.g. printf=float,
                    scanf=nofloat)
--verbose           Generate more detailed output
-l                 Generate a list file, defaults to 'ln65816.lst'
--list-file LIST-FILE
                    Generate list file, using given name
--memories-expression EXPRESSION
                    Expression that extracts the list of memory
                    descriptions from the .scm file, defaults to
                    'memories'
--program-root SYMBOL
                    Program root point, defaults to
                    '__program_root_section'
--root-symbol SYMBOL
                    Add a root symbol
--program-start SYMBOL
                    Program start symbol, defaults to '__program_start'
--copy-initialize SECTION
                    Override default and initialize this section by
                    copying
--no-copy-initialize SECTION
                    Override default and do not initialize this section
                    by copying
--no-automatic-placement-rules
                    Do not add rules to a .scm linker rules file
--raw-multiple-memories
                    Allow multiple memories for raw style output format
--no-merge-raw-memories
                    Never attempt to merge raw memories
```

(continues on next page)

(continued from previous page)

|                                      |                                                                                                                          |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <code>--force-output</code>          | Ignore (some) errors and attempt to generate output                                                                      |
| <code>--no-auto-libraries</code>     | Do not automatically add runtime libraries                                                                               |
| <code>--cstartup VALUE</code>        | Define the C startup to be used, synonym to <code>'--rtattr cstartup=VALUE'</code>                                       |
| <code>--no-tree-shaking</code>       | Do not perform tree shaking and keep unused section fragment in output                                                   |
| <code>--output-format FORMAT</code>  | Format, one of 'intel-hex', 'S-record', 'S19', 'S28', 'S37', 'raw', 'pgz' or 'prg' (in addition to the ELF/DWARF output) |
| <code>--stack-size SIZE</code>       | Stack size override (overrides size defined in the .scm file)                                                            |
| <code>--heap-size SIZE</code>        | Heap size override (size normally defined in the .scm file)                                                              |
| <code>--hosted</code>                | Initialize data sections in place by loading executable (hosted environment)                                             |
| <code>--rom-code</code>              | Run from ROM/Flash, initialize data sections by copying from ROM                                                         |
| <code>--initialize-large-data</code> | Enable initialization of data areas that can span over 64K boundaries                                                    |
| <code>--core CORE</code>             | Core, one of '65816' (defaults to '65816')                                                                               |
| <code>--target TARGET</code>         | Target system, one of 'C256', 'F256' or 'SNES' (defaults to embedded/ROM use, if omitted)                                |
| <code>-h,--help</code>               | Show this help text                                                                                                      |

### 23.13.2 Options in detail

#### `--output-file, -o`

Use this option to specify the name of the output executable file. If not given, the output file is `aout` with a file extension based on the format of the file, e.g. `aout.elf` or `aout.hex`.

#### `--debug`

Merge DWARF symbolic debugging information from the object file and output that in ELF executable file.

#### `-g`

Synonym for `--debug`.

#### `--semi-hosted`

Enable debug stubs for semi-hosting support. The program is linked with stubs for semi-hosting in the debugger. This allows the debugger to support standard I/O on behalf of the target and various other low level operations, such as capture `exit()` and implement `assert()`.

**--no-data-init-table-section**

Do not create any data initialization sections. This is mainly intended for assembly projects, but it can also be used for C if you do not want any initialization of static data objects to be done.

**--override**

Treat specified symbol as weak when taken from a library. This can be used if you want to replace a specific routine in a library with your own variant.

**--cross-reference**

Also include cross reference information in the list file. This is enabled by default.

**--no-cross-reference**

Do not include cross reference information in the list file.

**--rtattr NAME=VALUE**

Used to control selection of a specific variant of a module that exists in multiple weak versions. A typical situation is to select a custom C startup module or other variant of `exit()`.

It can also be used to override the automatic selection of `printf()` and `scanf()` formatter capability.

**--cstartup=VALUE**

Select a specific custom C startup module to be used. This is a slightly easier way than using the **--rtattr** option.

**--verbose**

Ask the linker to be more talkative. Mainly useful in cases where you run into errors. This option will generate a lot more information about the situation.

**-l**

Generate a list file. The name used is the name of the input file (ignoring any directory path) with a `.lst` file extension. See also **--list-file**.

### `--list-file`

Generate a list file. The name of the list file is given as argument to this option. See also `-l` to generate a list file based on the source filename.

### `--memories-expression`

This is the expression used to extract the memory description after reading a `.scm` rules file. This is set to `memories` by default and that should suffice in almost every situation.

### `--program-root`

This is the root of the application and everything referenced from this section, directly or indirectly is part of the application.

If you are building a normal C project you should not change this.

### `--root-symbol`

Add a root symbol that is regarded as something to always include in the build. This can be used to override a specific symbol, or if you have a root inside a library as they are otherwise ignored.

### `--program-start`

This symbol defines the first actual executable instruction in the program. This is used as the entry point which is defined in some output formats.

If you are building a normal C project you should not change this.

### `--output-format`

With this option you specify an additional file format for the produced executable application. An ELF executable is always produced.

### `--raw-multiple-memories`

This option enables output to multiple files in the Raw file format. By default only one area is allowed and an error will result if the application uses more than one program area. Specifying this option will also result in different output filenames that are based on the memory areas in the `.scm` file.

### `--no-merge-raw-memories`

The linker will attempt to merge memories when emitting in the Commodore PRG format for memories in the first bank (address `0x0000-0xffff`). This is to generate as few memories as possible due to limitations in the PRG format. You can avoid this automatic merging by specifying the `--no-merge-raw-memories` command-line option. This option also controls whether adjacent memories are merged in the RAW output format.

### `--heap-size`

The heap size can be defined in the linker control file (`.scm`). This option makes it possible to override the heap size and is useful if you use a common linker control file and want to avoid modifying the `.scm` file.

---

Note: The heap is completely removed if you set the size to 0. If the application tries to use the heap in that situation, the linker will complain that the heap section is undefined.

---

### `--stack-size`

The stack size can be defined in the linker control file (`.scm`). This option makes it possible to override the stack size and is useful if you use a common linker control file and want to avoid modifying the file.

### `--hosted`

Specifies a hosted environment. Data object initialization by copying from ROM is omitted; instead, the application is loaded into memory, and initialized data objects receive values during the load action.

### `--rom-code`

This is the opposite of `--hosted`. Data objects gets initialized by copying from ROM. This is the default and allows the application to be written to ROM or flash and started by turning the device on.

### `--copy-initialize`

This option takes a data section name as argument. You can use this option to override the default of whether the section is initialized by copying or by loading.

This option can be used together with the `--hosted` option to make an exception for a given data section and have that initialized by copying. It can be used when the load mechanism is unable to load directly into a memory area. You can specify this option multiple times.

### `--no-copy-initialize`

This option takes a data section name as argument. It is used to override the default of whether the section is initialized by copying or by loading. This option mostly exists for symmetry reasons. You can specify this option multiple times.

### `--core`

Specifies the core used. This is provided mainly for completeness and symmetry with the other tools. It has no practical meaning to the linker at the moment.

### `--target`

Specifies a certain target system. Using this option may affect the setting of `--hosted` versus `--rom-code` and is preferred if the target system used is supported by this option.

### `--no-automatic-placement-rules`

Do not attempt to add any rules, like section placements or missing blocks to a linker rules file. This is useful if you want to state everything explicitly, as it will result in errors if the linker rules (`.scm`) file is incomplete. In most cases you want to allow the linker to complete a partial linker rules file and this option should not be specified in such case.

### `--force-output`

Ignore certain errors and allow output to be generated anyway. This is mainly intended for getting a list file to help understanding the cause of errors for troubleshooting.

### `--no-auto-libraries`

The linker will attempt to pick a suitable C library based on the input. This command-line option disables this mechanism. If given you will need to specify the C library explicitly on the command line.

### `--initialize-large-data`

Data initialization is limited to Far objects which means that they cannot be allocated across 64K address banks. If you have custom data objects, e.g. graphics data with such needs, you can use this option.

When enabled two things happens. The data initializer table the is generated by the linker will use 32-bit size instead of 16-bit. The runtime attribute `initialize` is given the value `"large"` (rather than `"normal"` which is default). The result is that the linker will use a different data initialization module that can handle the larger initialization.

The `nlib` tool builds libraries (also called archives), which are single files containing multiple object files and a symbol index.

### 24.1 Creating a library

To create a library, list all object files (`.o` extension) and a single output file (`.a` extension).

```
$ nlib lib.a obj1.o obj2.o ashfx.o xtoal.o
```

The created library is typically used as input to the `ln65816` by specifying it on the command line. When linking with a library, only used objects are included in the final output.

### 24.2 File format

The file format is identical to `ar` on UNIX System V and GNU. The archive includes an index for faster object selection.

### 24.3 Command line options

Only `--version` and `--help` command-line options are supported.

To display the version of the linker, use `--version`:

```
$ nlib --version
Calypsi librarian for none version 5.16
```



The Calypsi tool chain debugger is a source code debugger that enables you to observe program execution. It offers both interactive and script-based control. Modeled after command-line debuggers like GDB and LLDB, it shares many similarities, making it familiar to users of those tools.

### 25.1 Command line debugger

A command-line debugger provides powerful commands that are easy to remember and understand. You can use descriptive commands like **step**, **break** and **continue** rather than a set of bindings to function keys and other more or less cryptic key bindings. Thanks to short forms and repeating the last command with a return key press, you can execute common and repetitive actions with single or few keystrokes.

The debugger also supports various front-ends like Visual Studio Code, Eclipse, and Emacs, leveraging the GNU project's MI (Machine Interface) protocol for this integration.

The debugger also provides powerful scripting abilities via the Lua language. You can use Lua to configure debugging sessions, create your own commands and even make your own debugger plug-ins. The debugger offers powerful scripting via Lua. You can use Lua to configure debugging sessions, create custom commands, and develop debugger plug-ins. This is possible because the complete debugger functional API is exposed to Lua. The Lua API further extends capabilities by allowing definition of new console commands and subscription to internal debugger event notifications.

Thus, with the Calypsi debugger, you gain the advantages of both the command-line environment and integrated graphical development environments.

## 25.2 Running the debugger

The debugger's interface is command-line based, allowing use from a terminal or connection to an IDE.

### 25.2.1 Basic invocation

The debugger is started in the following way:

```
$ db65816 [options] [debug-file] [options]
```

As everything on the command line is optional you can start the debugger without arguments. This will put you in an interactive debugger session:

```
$ db65816
```

Typically, specify the application program to debug on the command line:

```
$ db65816 application.elf
```

Command line options and the application to debug can appear in any order on the command line.

Command line options are optional arguments that tune debugger behavior, always starting with a dash. Two variants exist: single- letter options with one dash, and long descriptive options with two.

To display the version of the debugger, use `--version`:

```
$ db65816 --version
Calypsi debugger for 65816 version 5.16
```

## 25.3 Command line options

This section covers the `db65816` command-line options in detail.

### 25.3.1 Options overview

Use the `--help` option to display the available command-line options.

```
$ db65816 --help
Calypsi debugger for 65816 version 5.16

Usage: db65816 [--version] [-i|--interpreter INTERPRETER] [--nh] [--nx]
             [-e|--eval-command COMMAND] [-t|--tty TTY] [--batch]
             [--logging-enable] [--logging-to-file PATH] [--logging-overwrite]
             [--silent] [--extra-memory HEX-RANGE] [--stop-on-program-run]
             [--terminate-on-program-exit] [--exit-breakpoint] [--core CORE]
             [--target TARGET] [--program-start SYMBOL]
             [--semi-hosted-root PATH] [--target-remote DEVICE]
             [--serial-speed SPEED] [FILE]
    use 'db65816 --help' for detailed help

Available options:
  --version                Display version number
```

(continues on next page)

(continued from previous page)

```

-i,--interpreter INTERPRETER
                                Select command interpreter, one of 'Console', 'MI' or
                                'MI2' (defaults to 'Console')
--nh                            Do not read '~/.db65816'
--nx                            Do not read any '.db65816' file in any directory
-e,--eval-command COMMAND
                                Execute a single command (can be specified many
                                times)
-t,--tty TTY                    Use TTY for input/output by the program being
                                debugged
--batch                          Exit after processing options
--logging-enable                 Enable logging from start
--logging-to-file PATH          Enable logging from start to specified file
--logging-overwrite              Enable logging from start and overwrite the file
--silent                         Generate less messages
--extra-memory HEX-RANGE       Additional memory ranges to create (simulator use)
--stop-on-program-run           Stop and take control immediately when program is
                                started
--terminate-on-program-exit
                                Terminate execution of debugger if debuggee exits
                                (implies --exit-breakpoint)
--exit-breakpoint               Insert a breakpoint at 'exit()' to detect program
                                termination
--core CORE                     Core, one of '65816' (defaults to '65816')
--target TARGET                 Target system, one of 'C256', 'F256' or 'SNES'
                                (defaults to embedded/ROM use, if omitted)
--program-start SYMBOL          Program start symbol, defaults to '__program_start'
--semi-hosted-root PATH         Root directory of the semi-hosted file system
--target-remote DEVICE          Use remote device to communicate with target
--serial-speed SPEED            Set serial speed, one of 9600, 19200, 38400, 57600 or
                                115200 (defaults to 115200)
-h,--help                       Show this help text

```

### 25.3.2 Options in detail

#### --interpreter, -i

Use this option to specify the interpreter to use. The debugger can execute either with a console interpreter or a machine interface (MI) interpreter. The console interpreter is suitable for interactive debugger sessions from the command line, which is the default. The machine interface interpreter is used when the debugger is controlled by an IDE.

#### --eval-command, -e

This option specifies a command to execute before the interactive session gains control. It can be used multiple times to specify multiple commands.

### `--batch`

Terminate the debugger session after all `--eval-command` commands have been processed. This is useful for automated debugger execution from scripts.

### `--tty, -t`

Specifies the application's console `tty`. If the application writes to or reads from the standard stream, this provides the `tty` device for that operation.

If not specified, console output from the application is written to the console interaction. In MI mode such console output is marked so that the IDE can see that it is stream output from the application.

### `--semi-hosted-root`

Specifies the host-side directory serving as the target's semi-hosted file system. Defaults to the current directory. See [Semi-hosting](#) for details.

### `--logging-enable`

Enables logging at debugger startup.

### `--logging-to-file`

Specifies the logging output file.

### `--logging-overwrite`

Specifies whether logging to the file should append or overwrite previously existing log file content.

### `--extra-memory`

Specifies additional memory areas not part in the application that exists on a simulated target.

Memories in the simulator are normally created based on the memories found in the debug application image. This option allows additional memory regions to be created.

### `--stop-on-program-run`

When the program starts, immediately take control and stop execution. This enables debugging from the very beginning (before `main()` is reached) without needing to set a breakpoint.

**--terminate-on-program-exit**

If the application exits, the debugger session also terminates. The default behavior is for the debugger to notify of the exit and remain active.

**--exit-breakpoint**

Insert a breakpoint at the `exit()` function. This option is not needed if you link with debug stubs, as exit handling is provided.

**--target-remote**

Specifies the communications device to connect to a target for debugging. This disables the built-in simulator and configures the debugger for remote debugging via the gdbserver protocol, typically used for debugging on real hardware.

**--serial-speed**

Set the speed of the target remote communications serial port defined with `--target-remote`. The speed defaults to 115200.

**--program-start**

This symbol defines the program's first executable instruction. For a standard C project, avoid changing this setting.

**--core**

Specifies the core used.

**--target**

Specifies a certain target system.



---

## Console command line

---

The console interpreter provides an interactive command line with powerful features, including command completion, history, and editing.

Your command-line history is also saved between sessions in `~/.db65816/db65816.history`.

### 26.1 Command structure

Commands consist of one or more space-separated words. If a word is a unique match, it does not need to be fully spelled out.

An argument specific to the command (e.g., a string, file path, or boolean value) may follow the command.

#### 26.1.1 Command completion

Command completion is available via the TAB key. Pressing TAB directly at the prompt lists all words that can start a command:

```
(db65816) <TAB>
b          enable          n          s
break      exec-file       next       set
c          file           nexti      show
cd         frame          ni         si
clear      ignore          p          source
complete   info            platform  step
condition  interpreter-exec plug-in   stepi
continue   interrupt       print     thread
delete     kill            pwd       tty
directory  lua              quit
disable    maintenance     r
disassemble module          run
```

Pressing TAB completes the command to the left of the cursor as far as possible. If the word can be partially completed, the first TAB fills in as much as possible:

```
(db65816) int<TAB>
```

Results in:

```
(db65816) inter
```

Pressing TAB a second time results in:

```
(db65816) inter<TAB>
interpreter-exec interrupt
(db65816) inter
```

## 26.2 A simple debugging session

To get a feeling for how to run a very simple debugging session here is how it can look.

```
$ db65816 program.elf
Calypsi debugger for 65816
(db65816) b main
breakpoint 1
(db65816) info breakpoints


| Num | Type       | Disp | Enb | Lua/plugin | Address | What |
|-----|------------|------|-----|------------|---------|------|
| 1   | breakpoint | keep | y   | no         | 0x80f1  | main |


(db65816)
```

---

Note: A command requires enough characters to be unique. For `info breakpoints`, `i b` suffices, as the interpreter considers all words for multi-word commands. The completer operates left of the cursor: `i<TAB>` shows all commands starting with ‘i’. To complete to `info`, type `inf<TAB>` followed by `b<TAB>`. However, `i b<TAB>` expands directly to `i breakpoints` because no other command matches this two-word prefix.

---

To start debugging, use the run command (abbreviated to `r`), which begins execution from the start. The program will stop almost immediately upon hitting the breakpoint at the beginning of the `main()` function.

```
(db65816) r
running
    12 int const cglob_int = 12;
    13
    14 size_t fossasas;
    15
-> 16 int main () {
    17     basic_expr(cglob_int + 21, 2, 3);
    18     int xx = fact(4);
    19     int b = vla1(4);
    20     int c = struct1(2);
0x80f1 in main() at main.c:16
(db65816)
```

The breakpoint is hit at address `0x80f1` in `main.c:16`. Source lines are shown as context, with the current line indicated.

You can now perform a couple of single steps:

```

(db65816) s
13
14 size_t fossasas;
15
16 int main () {
-> 17  basic_expr(glob_int + 21, 2, 3);
18  int xx = fact(4);
19  int b = vla1(4);
20  int c = struct1(2);
21  int d = struct11(2);
0x80f8 in main() at main.c:17
(db65816)
10 int bar (long* p) {
11  return *p + 2;
12 }
13
-> 14 long basic_expr (long a, long a2, int a3) {
15  char cglob_int = 0; // block a variable in global scope
16  if (a) {
17      int b = 10;
18      a = foo(b + cglob_int);
0x8000 in basic_expr() at basic_expr.c:14
(db65816)

```

The **step** command (abbreviated as **s**) steps one source line. Subsequent steps can be performed by simply pressing RETURN, which repeats the most recent command.

Each step advances us one line, updating the source context shown.

The step command steps into functions, as demonstrated by entering **basic\_expr()**.

You can inspect variables with **info locals**.

```

(db65816) info locals
(int) a3 = 3
(long) a2 = 2
(long) a = 21
(db65816)

```

To trace your current position, use the **backtrace** (or **bt**) command.

```

(db65816) backtrace
#0 0x8000 in basic_expr() at basic_expr.c:14
#1 0x8137 in main() at main.c:17
(db65816)

```

To stop at a specific line further down the source code, use:

```

(db65816) b basic_expr.c:18
breakpoint 2
(db65816) c
running
program hit breakpoint 2
14 long basic_expr (long a, long a2, int a3) {
15  char cglob_int = 0; // block a variable in global scope
16  if (a) {
17      int b = 10;
-> 18  a = foo(b + cglob_int);

```

(continues on next page)

(continued from previous page)

```
19     a += bar(&a2); // take address of a2, forcing it to stack.
20 }
21 if (a > 2) {
22     // Make the variable in global scope visible
0x8030 in basic_expr() at basic_expr.c:18
(db65816)
```

The `continue` command (or `c`) resumes execution (unlike `run`, which starts from the beginning).

Lets take a look at the local variables again.

```
(db65816) info locals
(int) b = 10
(char) cglob_int = 0
(int) a3 = 3
(long) a2 = 2
(long) a - value not available, reason: "a" has no valid value here
(db65816)
```

The local context now includes two additional variables: `b` and `cglob_int`. At this point, `a` has no value because its last use was on line 16, and it will receive a new value on the current line. Thus, the debugger knows about `a` but it currently holds no examinable value.

Step over two function calls and inspect the value of `a` using the `print` command, which can be entered as `p`:

```
(db65816) n
15 char cglob_int = 0; // block a variable in global scope
16 if (a) {
17     int b = 10;
18     a = foo(b + cglob_int);
-> 19     a += bar(&a2); // take address of a2, forcing it to stack.
20 }
21 if (a > 2) {
22     // Make the variable in global scope visible
23     extern int const cglob_int;
0x8050 in basic_expr() at basic_expr.c:19
(db65816)
17     int b = 10;
18     a = foo(b + cglob_int);
19     a += bar(&a2); // take address of a2, forcing it to stack.
20 }
-> 21 if (a > 2) {
22     // Make the variable in global scope visible
23     extern int const cglob_int;
24     a = foo(cglob_int);
25 }
0x8082 in basic_expr() at basic_expr.c:21
(db65816) p a
(long) $8 = 8
(db65816)
```

The variable `a` now holds a value from the function calls.

The local `char cglob_int` shadows a global variable of the same name. Stepping to line 24 makes the global `cglob_int` visible:

```

(db65816) n
20  }
21  if (a > 2) {
22      // Make the variable in global scope visible
23      extern int const cglob_int;
-> 24      a = foo(cglob_int);
25  }
26  return a + a3;  // a3 alive and on the stack
27  }
28
0x8098 in basic_expr() at basic_expr.c:24
(db65816) p cglob_int
(int const) $9 = 12
(db65816)

```

The global variable is now visible again, with a different value and type.

If you no longer need the first breakpoint, you can disable it using the `disable` command, rather than deleting it:

```

(db65816) disable 1
(db65816) info breakpoints

```

| Num | Type                          | Disp | Enb | Lua/plugin | Address | What            |
|-----|-------------------------------|------|-----|------------|---------|-----------------|
| 1   | breakpoint                    | keep | n   | no         | 0x80f1  | main            |
|     | breakpoint already hit 1 time |      |     |            |         |                 |
| 2   | breakpoint                    | keep | y   | no         | 0x8030  | basic_expr.c:18 |
|     | breakpoint already hit 1 time |      |     |            |         |                 |

```

(db65816)

```

The first breakpoint still exists but is not enabled (n in the `Enb` column). The debugger also tracks breakpoint statistics.



---

## Remote debugging

---

This chapter describes how to set up remote debugging to a target. In this scenario, the debugger runs on the host and connects to, controlling, an application on actual hardware.

The gdbserver protocol is used for communication. It can be implemented by a small agent on the target or a bridge process on the host that translates between gdbserver and hardware-specific protocols.

### 27.1 Serial port

Target debug support can be implemented by placing a small debug agent software on the target. The RS-232 serial port can be used for this. Since most host computers lack these ports, a USB-to-serial converter (e.g., FTDI USB-RS232, often with bare wires requiring a soldered connector) can be used.

Most major operating systems include drivers for the FTDI USB-RS232 cable. If you choose other cables, you may need to install a driver, which you can obtain from the USB-RS232 cable vendor.

#### 27.1.1 Finding the device

On Linux and macOS, the serial port is typically found in `/dev/ttyUSBn`, `/dev/ttySn`, or similar. Use the `dmesg` command to query for USB devices:

```
$ sudo dmesg | grep USB
...
[109191.143623] ftdi_sio 1-1:1.0: FTDI USB Serial Device converter detected
[109191.151619] usb 1-1: FTDI USB Serial Device converter now attached to ttyUSB1
```

Here you can see that a FTDI serial device is attached to `/dev/ttyUSB1`.

---

Note: If you solder a serial port connector, be careful about which pin numbering. Also note that a DB-25 connector has the receive and transmit pins (numbers 2 and 3) are wired oppositely compared to a DE-9 connector.

---

### 27.1.2 Setting the speed

The speed can be set with the `--serial-speed` command-line option, which defaults to 115200 if not specified.

### 27.1.3 Testing the port

If you have the `screen` utility program installed it can be started with parameters to test that the serial port works:

```
$ screen /dev/ttyUSB1 115200
```

Once started, you can verify debugger agent responsiveness by pressing any character; a `$S13#b7` packet reply confirms it.

You can exit by pressing Control-A followed by a backslash (`c-a \`). Some screen versions require Control-A followed by Control-backslash.

### 27.1.4 Installing the agent

The Calypsi remote debugger agent for the serial port must be downloaded, built, and executed on the target hardware. Download it from <https://github.com/hth313/Calypsi-remote-debug>.

### 27.1.5 Starting the debugger

Once the debugger agent runs on the hardware and the communications channel is configured, you can start the Calypsi `db65816` using:

```
$ db65816 --target-remote /dev/ttyUSB1 application.elf
```

---

### Visual Studio Code

---

Visual Studio Code<sup>5</sup> is a lightweight, yet powerful source code editor that runs on multiple platforms.

With the C/C++ for Visual Studio Code extension<sup>6</sup> it is possible to use the Calypsi C compiler tool chain in Visual Studio Code.

Visual Studio Code can be tailored to use various build and debugger tools. Its debugger integration, based on MI (Machine Interface), is supported by the Calypsi db65816 debugger.

Visual Studio Code is an excellent GUI choice for Calypsi, being actively developed, lightweight, responsive, and highly configurable.

This guide does not fully document Visual Studio Code. For detailed information, refer to the official documentation<sup>7</sup>.


Some basics on getting started are covered here, as its operation may differ from traditional IDEs.

---

<sup>5</sup> <https://code.visualstudio.com/>

<sup>6</sup> <https://marketplace.visualstudio.com/items?itemName=ms-vscode.cpptools/>

<sup>7</sup> <https://code.visualstudio.com/docs>



## 28.1 Installation

Installation is easy to do by downloading it from the official site. After that you can update from inside Visual Studio Code itself.

Note: If you use Arch Linux or a derivative like Manjaro, install VS Code via its package system. The community code package currently cannot install the vscode-cpptools extension due to licensing. To make it work, install the Microsoft-branded visual-studio-code-bin from the AUR.

## 28.2 Project setup

In contrast to many IDEs you will not find any project setup in the menus.

To create a new project from scratch, use **File>Open** and create a new folder for your project and then press **Open**. You now have an empty project directory.

For an existing project, select **File>Open**, browse to its top folder, and press **Open** to load.

Project-specific files are stored in a `.vscode` directory within the project. Initially, this directory is empty.

## 28.3 Building

Once the project is populated and a Makefile created, add a build task. Select **Tasks>Configure Tasks...**, then **Create tasks.json file from template**, and **Others**. Edit the resulting `tasks.json` file, which may resemble:

```
{
  "version": "0.2.0",
  "tasks": [
    {
      "label": "build",
      "type": "shell",
      "command": "make -k",
      "options": {
        "cwd": "${workspaceRoot}"
      },
      "problemMatcher": [
        "$gcc"
      ]
    },
    {
      "label": "clean",
      "type": "shell",
      "command": "make -k clean",
      "problemMatcher": [],
    }
  ]
}
```

## 28.4 Running the debugger

The `db65816` debugger can be used inside Visual Studio Code. For basic use it uses a 65816 simulator and a memory system that is configured based on the executable image.

## 28.4.1 Configuration

Before starting the debugger, add a launch configuration.

```
{
  "version": "0.2.0",
  "configurations": [
    {
      "name": "C++ Launch",
      "type": "cppdbg",
      "request": "launch",
      "program": "${workspaceRoot}/src/myproject",
      "args": [],
      "stopAtEntry": false,
      "cwd": "${workspaceRoot}/src",
      "environment": [],
      "externalConsole": false,
      "linux": {
        "MIMode": "lldb",
        "miDebuggerPath": "/usr/local/bin/db65816",
        "launchCompleteCommand": "exec-run",
        "setupCommands": []
      },
      "osx": {
        "MIMode": "lldb",
        "miDebuggerPath": "/usr/local/bin/db65816",
        "launchCompleteCommand": "exec-run",
        "setupCommands": []
      },
      "windows": {
        "MIMode": "lldb"
      }
    }
  ]
}
```

There are a couple of things to specify here.

1. You need to specify that you are using `/usr/local/bin/db65816` as the debugger using `miDebuggerPath`
2. You also need to tell that you want to use `lldb` as `MIMode`<sup>8</sup>

The launch configuration is a text file stored in `.vscode/launch.json`.


---

<sup>8</sup> Even though `db65816` is much closer to `gdb` than `lldb` at the command level, you need to specify `lldb` as the MI mode. The reason is that in the `gdb` mode, Visual Studio Code works around a bug on UNIX style platforms for the MI command `-exec-interrupt` (which interrupts execution). The work around is to send a signal to a real UNIX process to interrupt it. This does not work with `db65816` as the application being debugged is not a UNIX process at all. Instead you need to use the `lldb` mode, where the MI command `-exec-interrupt` works properly, just as it does in `db65816`.

Reference [https://sourceware.org/bugzilla/show\\_bug.cgi?id=20035](https://sourceware.org/bugzilla/show_bug.cgi?id=20035)

## 28.4.2 Running the application

Once configured, run your application using Run >> Start Debugging (F5). Consider inserting a breakpoint or interrupting execution once started. If debugger windows (variables, call stack, etc.) don't appear, click the bug icon to switch to the debugger view.





The debugger can be controlled and extended using the [Lua language](https://www.lua.org/)<sup>9</sup>. Scripting allows you to write programs that execute inside the debugger, from simple batch-like scripts to advanced plug-ins that add custom commands, listen to internal events, and control execution.

From a high-level perspective, Lua and the debugger enable various tasks:

- Perform setup tasks to prepare your debug session
- Simplify repetitive tasks
- Run fully automated tests
- Extend functionality; for instance, to introspect the state and behavior of an operating system or protocol stack. These commands integrate into the existing set and share a command-line interface.

You can choose any automation level, from running the entire session under script control to providing high-level functionality with commands.

### 29.1 Running a script

To get started, this simple Lua program just displays the registers:

```
local db = require('db')  
  
print(db.consoleCommand('info registers'))
```

If saved in `script-name.lua` (in the current directory), you can run the script with:

```
(db) source script-name.lua
```

Debugger functionality is provided in the `db` table which is created before the script is started. To access it, you need to use `require`. Here we store the debugger table in the local variable `db`.

---

<sup>9</sup> <https://www.lua.org/>

---

Note: Use the `--eval-command` (can also be used as `-e`) to invoke a Lua script from the command line and to make it execute automatically when the debugger starts. The `--eval-command` works with any command and can be specified multiple times on the command line.

---

## 29.2 Lua environments

Each time you run a Lua script, a new fresh Lua state is created. This means you will start out with the standard Lua and debugger libraries available imported, but not anything else. If you create a global variable, it will not normally be present the next time you run the script.

However, if a script provides new debugger functions (e.g., console commands), the Lua state is saved with the command. Subsequent invocations of that command reuse the same state, retaining global variables and closures. This allows commands to operate within their established state and depend on prior actions within that state, including previous uses of the same or other commands created by the script.

A Lua plug-in typically registers functions with the debugger for specific circumstances. These functions always execute in the state in which they were installed.

It is good practice to provide related functionality (e.g., a command set and plug-in behavior) in its own script. Multiple plug-ins can be installed via source invocation, allowing each to have its private shared state while co-existing separately.

## 29.3 Command interpreters

The debugger provides two sets of command interpreters, console commands and MI commands.

Console commands are intended for humans, provide command completion when being entered and formatted output that is easy to read.

The Machine Interface (or MI for short), was introduced by GDB to provide a more precise and stable command protocol, intended for (graphical) debugger front-ends. The GDB MI command set is supported by the debugger, allowing existing graphical debugger front-ends to be used.

Many commands can be accessed from either of these command sets, while some commands are only available in one of them.

The Lua interface allows both command sets to be used. In addition, adapted MI commands and additional functionality are also provided.

## 29.4 Console commands

The `db` table allows sending console commands to the debugger. These commands execute as if typed at the terminal, and their string results are passed back to Lua.

If you need to process the result beyond displaying or ignoring it, parsing is required. However, writing such parsers quickly becomes tedious, as the text output is primarily for display.

## 29.5 MI commands

MI commands offer precise interaction for debugger front-ends, making them well-suited for Lua scripts. Although commands are string-based, results are converted to nested Lua tables with basic types (booleans, numbers, strings).

This significantly simplifies interpreting results from Lua-to-debugger calls.

## 29.6 Adapted MI commands

Adapted MI commands are basically MI commands converted to be a more natural fit for Lua. This is done in two ways. The actual commands take arguments like any ordinary function in Lua. A return values is the value returned or a failure, instead of the MI style result table hierarchy.

## 29.7 Choosing a command set

In most cases you will probably want to use the adapted MI command set as it provides the most natural Lua experience.

The console and MI command set may be useful if you have familiarity with either or prefer the alternative output form provided.

## 29.8 Arrays and Lua

Lua array indices start at 1, unlike C. This mismatch has implications when working with debugger C expressions from Lua.

To simplify working with C expressions, array results can start at 0 in Lua, mirroring C behavior. However, the `ipairs` iterator (designed for Lua's 1-based indexing) will skip the element at index 0.

Shifting the index in expressions is an alternative, but it would likely be more confusing and error-prone, especially given the context of C expressions.

## 29.9 Debugger API

### **`db:stepInstruction()`**

Single step once at instruction level. This function will quickly return an indication of whether it was successful in starting the target. Actual progress of execution will be posted as notifications

Corresponding MI instruction is `-exec-step-instrucion`.



This section describes the data structures used. Complex types (beyond booleans, numbers, or strings) are represented by tables.

## 30.1 Address

The debugger uses internal addresses, composed of a numeric value, memory type, and optional bank number. In Lua, addresses are represented by a table with these fields:

Table 30.1: Address type

| field     | type    | description                                                                                 |
|-----------|---------|---------------------------------------------------------------------------------------------|
| formatted | string  | The complete print value of the address location.                                           |
| address   | integer | The raw address bits of the location.                                                       |
| memory    | string  | The kind of the memory the address points to, can typically be “code” or “data” (string).   |
| bank      | integer | The bank number, or nil if this address location does not have any bank associated with it. |

The **formatted** field is set when the value is returned to Lua. If you provide an address as an argument to a function call in the API, you do not need to provide a **formatted** field.

## 30.2 Adapted MI functions

Many Lua functions are adapted from machine interface (MI) commands, used for communication between a debugger front end (like VS Code) and the debugger.

Functions are available in the `db` table. The Lua interface translates arguments for these functions and passes them to the debugger command handler.

Results are passed back and translated to Lua. As the underlying mechanism is a command protocol, the result table is a bit more elaborate than what might be expected.

A top-level result has the following fields:

| field | type    | description                                                                                                        |
|-------|---------|--------------------------------------------------------------------------------------------------------------------|
| class | string  | The result class (e.g., <code>done</code> for success, <code>error</code> for failures). A complete list is below. |
| value | integer | The result, which depends on the function called.                                                                  |
| type  | string  | The reply type. For function calls, this is <code>result</code> .                                                  |

This section covers functions provided by the debugger module, grouped by family for easy reference.

## 31.1 Execution

### 31.1.1 `db:stepInstruction()`

Steps once at the instruction level.

The result indicates only if the target was notified to start. An error result means the target is not in a state to start.

To monitor progress and result of execution, you need to listen to notifications.

## 31.2 Breakpoint

### 31.2.1 `db:breakDelete(number(s))`

Removes one or more breakpoints.

The argument is the breakpoint number(s) to remove (an integer or array of integers).

### 31.2.2 db:breakInsert(location, &callBack)

Inserts a new breakpoint.

The first argument, a Lua table, may contain breakpoint attributes:

Table 31.1: Breakpoint attributes

| field      | type    | description                             |
|------------|---------|-----------------------------------------|
| temporary  | boolean | set a temporary (one-time) breakpoint   |
| hardware   | boolean | Use a hardware breakpoint.              |
| pending    | boolean | Allow pending breakpoint.               |
| disabled   | boolean | Start with the breakpoint disabled.     |
| tracePoint | boolean | This is a tracepoint.                   |
| condition  | expr    | Set a conditional expression.           |
| ignore     | integer | Set initial ignore count.               |
| thread     | integer | Breakpoint is valid for a given thread. |

The second argument describes the breakpoint's location, which can be:

Table 31.2: Breakpoint location

| kind          | type   |                                 |
|---------------|--------|---------------------------------|
| address       | table  | an address table                |
| file and line | table  | { file : string, line : integer |
| function      | string | function name or symbol         |

The third if present should be a Lua function. This function will be called as followed when the breakpoint is hit:

```
cont = callBack(object, breakpointNumber)
```

The callBack's first argument is the object table used during breakpoint creation; the second is the breakpoint number. If callBack returns boolean false, the breakpoint is not taken, and execution resumes silently, without user interface feedback.

The callback occurs in the execution thread and blocks execution. If you intend to continue, return as quickly as possible.

If execution stops, you can process it within the callback or detect the stop via notifications. Be aware that processing in the callback will prevent user feedback until it returns.

## Non-alphabetical

`_DirectPageStart`  
symbol, 159  
`_NearBaseAddress`  
symbol, 159

## A

`abs:`  
modifier, 127  
`.address`  
directive, 135  
address space  
far, 21  
far24, 22  
huge, 22  
near, 21  
syntax, 22  
tiny, 21  
address spaces, 21, 22  
addressing mode  
specifying, 127  
`.align`  
directive, 136  
aligned  
attribute, 65  
alignment, 53, 62, 65  
of sections, 128  
allocated memory, 161  
Arch Linux  
installation, 5  
uninstalling, 5  
archive, 172  
`.ascii`  
directive, 135  
`.asciz`  
directive, 135  
assembler, 125  
constants, 129  
directives, 131  
expressions, 129  
invocation, 141  
labels, 126  
local labels, 130  
numbers, 129  
operators, 129  
symbols, 126  
syntax, 125  
assembly  
65816 specifics, 126

immediate size, 127  
assembly language, 125  
assembly source  
generated by compiler, 44  
assembly source file  
option, 49  
attribute  
aligned, 65  
far, 66  
far24, 66  
huge, 66  
interrupt, 66  
intrinsic, 67  
near, 65  
saveds, 67  
section, 65  
task, 67  
tiny, 65  
attributes, 63  
auto variables, 20  
auto\_type, 70

## B

back quoted symbols, 126  
bare metal, 153  
base address, 159  
definition, 159  
bit fields, 56  
volatile, 61  
block  
section in linker, 158  
board support  
enabling, 35  
bool type, 55  
bss section, 89, 128  
building  
C library, 108  
built-in, 75  
built-in functions, 75  
`.byte`  
directive, 135  
byte order  
macros, 84  
`.byte0`  
relocation operator, 137  
`.byte1`  
relocation operator, 137  
`.byte2`  
relocation operator, 137

## C

- C language, 11, 28
- C library, 98
  - building, 108
  - errno, 104
  - replace function, 106
- C startup, 106
  - replacing, 106
- C256, 36
  - target, 149
- C256 Foenix U
  - example, 8
- calling convention, 118
- cdata
  - section, 92
- cfar
  - section, 25, 90, 94
- character type, 56
- chuge
  - section, 25, 90, 94
- cnear
  - section, 25, 90, 93
- code
  - section, 25, 90, 92
- code model, 14
  - compact, 14
  - large, 14
  - small, 14
- command
  - structure, 181
- command completion, 181
- comments
  - assembly, 125
- compact
  - code model, 14
- compactcode
  - section, 25, 90, 92
- compiler
  - dependency file, 52
  - diagnostics, 44
  - errors, 44
  - generate assembly source, 44
  - invocation, 41
  - list file, 42
  - options, 46
  - warnings, 44
- conditional assembly, 126
- const
  - qualifier, 62
- constants
  - assembler, 129
- constraints
  - inline assembler, 121
- context
  - program, 182
- conversions
  - implicit, 32
- core selection
  - option, 53, 149
- cross call
  - optimization, 112
  - option, 51
- cross compiler, 11
- cross reference, 161
- customizing
  - system startup, 106

## D

- data
  - section, 92
- data bank
  - start, 159
- data initialization, 27
  - omitting, 168
- data model, 14
  - huge, 15
  - large, 15
  - medium, 14
  - small, 14
- data pointer, 57
- data section, 89, 128
  - initialization, 90
- data size, 119
- data\_init\_table
  - section, 25, 90, 95
- Debian
  - installation, 3
  - uninstalling, 4
- debug library, 98
- debugger
  - command interpreter, 177
  - logging, 178
  - starting, 175
- debugging information
  - macros, 84
  - option, 49, 148
- dependency file
  - compiler, 52
- devices
  - supported, 11
- diagnostics
  - compiler, 44
  - pragma directive, 74
- direct page, 19, 65
  - assembly, 126
  - pseudo registers, 19
  - start, 159
- directive
  - .address, 135
  - .align, 136
  - .ascii, 135
  - .asciz, 135
  - .byte, 135
  - .equ, 132, 133
  - .equlab, 133
  - .extern, 133
  - .fillto, 135
  - .global, 133
  - .globl, 133
  - include binary, 134
  - .long, 135
  - .macro, 136, 139
  - .public, 133
  - .quad, 135
  - runtime model, 134
  - .section, 127, 132
  - .space, 135
  - .word, 135
- directives
  - assembler, 131
- distribution
  - installation, 2
- dp:
  - modifier, 127
- DWARF

output, 165

## E

ELF

output, 165

embedded systems, 153

enabling

board support, 35

endian

macros, 84

enumeration type, 59

.equ directive, 132, 133

.equlab directive, 133

errno

C library, 104

error

pragma directive, 74

errors

compiler, 44

internal, 44

example

C256 Foenix U, 8

Hello World, 6

exit, 105

expressions

assembler, 129

extended attributes, 63

.extern directive, 133

## F

F256, 38

target, 149

far

address space, 21

attribute, 66

section, 25, 90, 94

far24

address space, 22

attribute, 66

farcode

section, 25, 90, 92

Fedora

installation, 4

uninstalling, 5

file extensions, 12

file inclusion, 126

file streams

initialization, 27

fill word, 157

.fillto

directive, 135

floating point

precision, 57

setting, 15

types, 15, 57

floating point types, 31

Foenix

target, 53

Foenix C256, 36

Foenix C256 U

example, 8

Foenix F256, 38

formatter

printf, 105

scanf, 105

function prototypes, 32

functions, 75

## G

global variables, 20

## H

headers

search path, 42

heap

initialization, 27

specify size, 158

Hello World

example, 6

host based systems. operating system, 153

huge

address space, 22

attribute, 66

data model, 15

section, 25, 90, 94

## I

idata

section, 25, 90, 92, 93, 95

IDE

Visual Studio Code, 188

IEEE 754, 57

ifar

section, 25, 90, 93, 94

ihuge

section, 25, 90, 94

immediate size

assembly, 127

include binary

directive, 134

include files, 126

include search

option, 48, 148

include search path, 42

index register

size, 119

index type

pointer, 57

inear

section, 25, 90

initialization

of data, 27, 90

inline assembler, 120

constraints, 121

substitutions, 122

volatile, 121

inlining

optimization, 110

inlining functions

not inlined, 112

option, 51

installation, 2

Arch Linux, 5

Debian, 3

Fedora, 4

Linux, 35

macOS, 3

Manjaro, 5

multiple, 6

multiple versions, 6

Ubuntu, 3

Windows, 5

integer types, 31, 55

Intel hex

- output, 165
- internal errors, 44
- interrupt
  - attribute, 66
  - vector, 66
- intrinsic, 75
  - attribute, 67
- intrinsic functions, 75
- intrinsic, 17
- invocation
  - assembler, 141
  - compiler, 41

## K

- .kbank
  - relocation operator, 138
- keywords
  - address spaces, 22

## L

- labels
  - assembler, 126
  - global, 133
  - location counter, 130
  - weak, 133
- language extensions
  - auto\_type, 70
  - overloaded functions, 69
  - statement expression, 69
  - typeof, 70
- large
  - code model, 14
  - data model, 15
- librarian, 172
- library
  - building, 108
  - files, 98
  - replace function, 106
- linker, 149
  - cross reference, 161
  - scatter, 160
- linker rules, 27, 152, 157
- Linux
  - installation, 35
  - uninstalling, 4, 5
- list file
  - compiler, 42
  - option, 48, 147
- list files
  - linker, 161
- load address
  - option, 171
- local labels
  - assembler, 130
- local labels inside macro, 139
- local variables, 20
- location counter, 130
- .long
  - directive, 135
- long:
  - modifier, 127
- Lua, 193
- lua
  - adapted MI functions, 199
  - address type, 199
  - breakpoints, 201
  - execution control, 201

## M

- macOS
  - installation, 3
  - uninstalling, 3
- .macro, 139
  - directive, 136
- macro, 139
  - local label, 139
- macro definition
  - option, 49, 148
- macro definitions
  - showing, 50
- macro language, 139
- macros
  - byte order, 84
  - debugging information, 84
  - endian, 84
  - predefined, 83
  - version, 84
- Manjaro
  - installation, 5
  - uninstalling, 5
- medium
  - data model, 14
- memory, 156
  - address range, 156
  - fill, 157
  - name, 156
  - placement group, 159
  - size, 161
- message
  - pragma directive, 74
- Motorola S-record
  - output, 165

## N

- .near
  - relocation operator, 138
- near
  - address space, 21
  - attribute, 65
  - base address, 159
  - section, 25, 90, 93
- near bank, 19
- no-init
  - section, 90
- numbers
  - assembler, 129

## O

- omit data initialization, 168
- operator
  - .sectionEnd, 138
  - .sectionSize, 138
  - .sectionStart, 138
- operators
  - assembler, 129
  - relocation, 136
  - section, 138
- optimization
  - cross call, 112
  - inlining, 110
- optimizer, 16
  - option, 50
  - tuning, 50, 51
- option

- core selection, 53, 149
- cross call, 51
- debugging information, 49, 148
- function inlining, 51
- include search, 48, 148
- list file, 48, 147
- load address, 171
- macro definition, 49, 148
- optimizer level, 50
- output file, 48, 147
- runtime attribute, 51, 149
- signed char, 50
- size of double, 50
- switch control, 52
- system include search, 49, 148
- target selection, 53, 149, 171
- unsigned char, 50
- vector sections, 51
- version, 48, 147
- weak symbols, 51, 149
- option: assembly source file, 49
- options
  - compiler, 46
- output
  - DWARF, 165
  - ELF, 165
  - Intel hex, 165
  - Motorola S, 165
  - raw, 166
  - S19, 165
  - S28, 165
  - S37, 165
  - SREC, 165
- output file
  - option, 48, 147
- output format
  - ELF/DWARF, 165
  - Intel hex, 165
  - Motorola S-record, 165
  - raw, 165
- overloaded functions, 69

## P

- placement control, 27
- placement group
  - memory, 159
- pointer
  - index type, 57
  - types, 57
- pointers, 22
- pragma directive
  - diagnostics, 74
  - error, 74
  - message, 74
  - require, 75
  - rtattr, 75
  - runtime attribute, 75
  - section, 74
  - warning, 74
- precision
  - floating point, 57
- predefined
  - macros, 83
- preprocess, 50
- preprocessor, 81
- print formatter, 105
- program context, 182

- prototypes
  - function, 32
- pseudo registers, 19
- ptrdiff\_t type, 58
- .public directive, 133

## Q

- .quad
  - directive, 135
- quoted symbols, 126

## R

- raw
  - output, 166
- read only
  - data section, 128
- read-only
  - section, 90
- record type, 58
- registers
  - pseudo, 19
- relocation operators, 136
- relocations, 136
- replace function
  - in C library, 106
- require
  - pragma directive, 75
- required symbols, 134
- rodata section, 90
- ROM based systems, 153
- rtattr
  - pragma directive, 75
- rules for linking, 152, 157
- runtime attribute
  - option, 51, 149
  - pragma directive, 75
- runtime library, 98
- runtime model directive, 134

## S

- S19
  - output, 165
- S28
  - output, 165
- S37
  - output, 165
- saves
  - attribute, 67
- scan formatter, 105
- scatter
  - linker, 160
- search path
  - headers, 42
- .section
  - directive, 127
- section, 25, 127
  - alignment, 128
  - attribute, 65
  - block, 158
  - bss, 89, 128
  - cdata, 92
  - cfar, 25, 90, 94
  - chuge, 25, 90, 94
  - cnear, 25, 90, 93
  - code, 25, 90, 92
  - compactcode, 25, 90, 92

- data, 89, 90, 92, 128
- data\_init\_table, 25, 90, 95
- directive, 132
- far, 25, 90, 94
- farcode, 25, 90, 92
- fragments, 127, 161
- huge, 25, 90, 94
- idata, 25, 90, 92, 93, 95
- ifar, 25, 90, 93, 94
- ihuge, 25, 90, 94
- inear, 25, 90
- kinds, 128
- modifiers, 128
- near, 25, 90, 93
- no-init, 90
- operators, 138
- placement, 27, 157
- pragma directive, 74
- read only data, 128
- read-only, 90
- stack, 158
- switch, 25, 90, 95
- text, 89, 128
- tiny, 25, 90, 93
- types, 90
- zdata, 92
- zfar, 25, 90, 93
- zhuge, 25, 90, 94
- znear, 25, 90, 93
- ztiny, 25, 90, 93
- .section directive, 132
- .sectionEnd
  - operator, 138
- .sectionSize
  - operator, 138
- .sectionStart
  - operator, 138
- semi-hosted, 99
- sequence point, 60
- signed char
  - option, 50
- simulated I/O, 99
- size
  - data access, 119
  - index register, 119
  - of double, 15
  - of float, 15
- size of double
  - option, 50
- size\_t type, 58
- small
  - code model, 14
  - data model, 14
- SNES
  - target, 53, 149
- .space
  - directive, 135
- SREC
  - output, 165
- stack
  - section, 158
  - size considerations, 28
  - specify size, 158
  - usage, 28
- startup, 106
  - custom, 106
- statement expression, 69
- static variables, 20

- streams, 104
- structure type, 58
- stubs interface, 99
- substitutions
  - inline assembler, 122
- supported devices, 11
- switch
  - section, 25, 90, 95
- switch control
  - option, 52
- symbol
  - \_DirectPageStart, 159
  - \_NearBaseAddress, 159
- symbols
  - assembler, 126
  - global, 133
  - quoted, 126
  - required, 134
  - syntax, 126
  - weak, 133
- syntax
  - assembler, 125
- system include search
  - option, 49, 148
- system startup, 106
  - customizing, 106

## T

- target
  - C256, 149
  - F256, 149
  - Foenix, 53
  - option, 53, 149
  - SNES, 53, 149
- target selection
  - option, 171
- task
  - attribute, 67
- termination
  - simplification, 105
- text section, 89, 128
- time
  - functions; date; functions, 107
- .tiny
  - relocation operator, 137
- tiny, 159
  - address space, 21
  - attribute, 65
  - base address, 159
  - section, 25, 90, 93
- type
  - enumeration, 59
  - ptrdiff\_t, 58
  - size\_t, 58
  - structure, 58
  - union, 58
- type definitions, 62
- type qualifier, 22
- type qualifiers, 60
- typedef, 62
- typeof, 70
- types
  - bool, 55
  - character, 56
  - integer, 55
  - wide character, 56

## U

- Ubuntu
  - installation, 3
  - uninstalling, 4
- uninstalling
  - Arch Linux, 5
  - Debian, 4
  - Fedora, 5
  - Linux, 4, 5
  - macOS, 3
  - Manjaro, 5
  - Ubuntu, 4
  - Windows, 6
- union type, 58
- unsigned char
  - option, 50
- used memory, 161

## V

- variables
  - auto, 20
  - dynamic, 20
  - global, 20
  - local, 20
  - static, 20
- vector sections
  - option, 51
- version
  - macros, 84
  - option, 48, 147
- Visual Studio Code, 188
- volatile
  - access size, 60
  - inline assembler, 121
- volatile objects, 60

## W

- warning
  - pragma directive, 74
- warnings
  - compiler, 44
  - controlling, 45
- weak symbols, 133
  - option, 51, 149
- wide character type, 56
- Windows
  - installation, 5
  - uninstalling, 6
- .word
  - directive, 135
- .word0
  - relocation operator, 137
- .word2
  - relocation operator, 137

## Z

- zdata
  - section, 92
- zfar
  - section, 25, 90, 93
- zhuge
  - section, 25, 90, 94
- znear
  - section, 25, 90, 93
- ztiny
  - section, 25, 90, 93