

Everything you ever wanted to
know about “hello, world”*

(*but were afraid to ask)

Brooks Davis
SRI International

June 10, 2016
BSDCan 2016



K&R: *The C Programming Language*

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

```
main()  
{  
    printf("hello, world\n");  
}
```

K&R: *The C Programming Language*

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

```
void
```

```
main(void)
```

```
{
```

```
    printf("hello, world\n");
```

```
}
```

Today's version

```
int
main(void)
{
    const char hello[] =
        "Hello World!";

    printf("%s %d\n", hello, 123);

    return (0);
}
```

Minimal C version

```
void
main(void)
{
    const char *hello[] =
        "hello, world\n";

    write(1, hello, sizeof(hello));

    exit(0);
}
```

Minimal (MIPS) assembly version

```
        .text
        .global __start
        .ent __start
__start:
        li $a0, 1
        dla $a1, hello
        li $a2, 12
        li $v0, 4
        syscall          # write(1, "hello, world\n", 13)
        li $a0, 0
        li $v0, 1
        syscall          # exit(0)
        .end __start

        .data
hello:
        .ascii "hello, world\n"
```

Size comparison

- Assembly
 - Compiles to 9 instructions
 - Stripped binary less than 1K
 - Mostly ELF headers, MIPS ABI bits
- Minimal C
 - Stripped binary over 550K!
 - Mostly malloc() and localization

Program linkage

```
$ cc -static -o helloworld helloworld.o
```

```
$ ld -EB -melf64btsmip_fbsd -Bstatic \  
-o helloworld /usr/lib/crt1.o \  
/usr/lib/crti.o /usr/lib/crtbeginT.o \  
-L/usr/lib helloworld.o \  
--start-group -lgcc -lgcc_eh -lc \  
--end-group \  
/usr/lib/crtend.o /usr/lib/crtn.o
```


Compiler runtime support

File	Purpose
<code>crt1.o</code>	Contains <code>__start()</code> function which initializes process environment and calls <code>main()</code> .
<code>crti.o</code>	Entry points for old style <code>_init()</code> and <code>_fini()</code> functions.
<code>crtbegin.o</code> <code>crtbeginS.o</code> <code>crtbeginT.o</code>	Declares <code>.ctor</code> and <code>.dtor</code> constructor and destructor sections. Declares functions to call constructors and destructors.
<code>crtend.o</code>	NULL terminates <code>.ctor</code> and <code>.dtor</code> sections.
<code>crtn.o</code>	Trailers for <code>_init()</code> and <code>_fini()</code> functions.

Built in `gnu/lib/csu` and `lib/csu/ARCH`.


Code and images online

<https://people.freebsd.org/~brooks/talks/bsdcan2016-helloworld>


or

<http://bit.ly/helloworld-talk>

execve()




sys_execve()



kern_execve()


namei()
Resolve
path



exec_check_permissions()
Check that the file has the right
permissions and open it.

exec_map_first_page()
Map the header into kernel
memory.


exec_elf64_imgact()



exec_elf64_imgact()


elf_load_section()
Map .text section into
memory

elf_load_section()
Map .data section into
memory and create bss

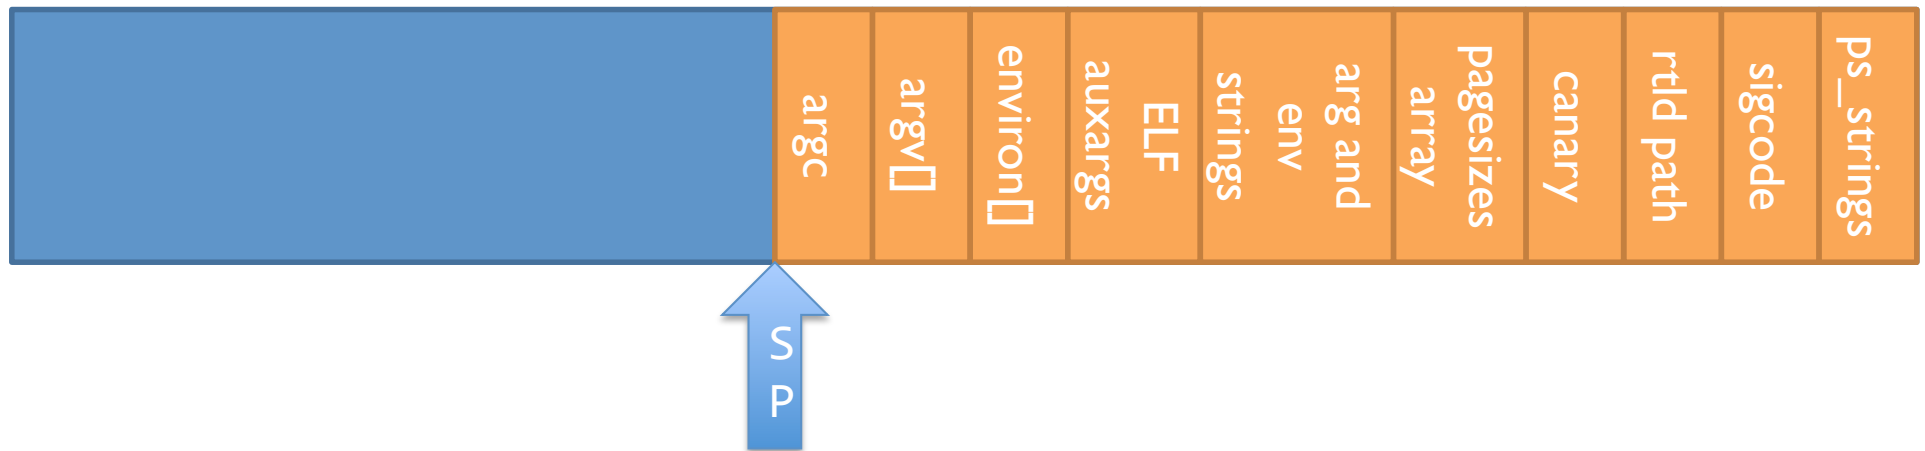


Returning to userspace

- Stack is mapped into address space
- Program is mapped into address space
- Strings, argv, envp, signal handler, etc are on the top of the stack
- Register state is set up to call `__start()`



SCO i386 ABI stack



```

__start(char **ap, ...) {
    ...
    argc = * (long *) ap;
    argv = ap + 1;
    env  = ap + 2 + argc;
    ...
}

```


__start() 1/2

```
void __start(char **ap)
{
    int argc;
    char **argv, **env;

    argc = * (long *) ap;
    argv = ap + 1;
    env = ap + 2 + argc;

```

...

__start() 2/2

...

Set environ and
__progname variables.

```
handle_argv(argc, argv, env);  
_init_tls();  
handle_static_init(argc, argv,  
env);
```

```
exit(main(argc, argv, env));  
}
```


_init_tls()

- Find the ELF auxargs vector

- Use that to find the program headers
- Use those to find the PT_TLS section (initial values)
- Call `libc_allocate_tls()` (as `while (*sp++ != 0) _rtld_allocate_tls()`)
 - Allocates space.
- Call `aux = (Elf_Auxinfo *)`
- Copies initial values

Uses JEMalloc, but JEMalloc uses TLS!

- Set the TLS pointer

__start() 2/2

...

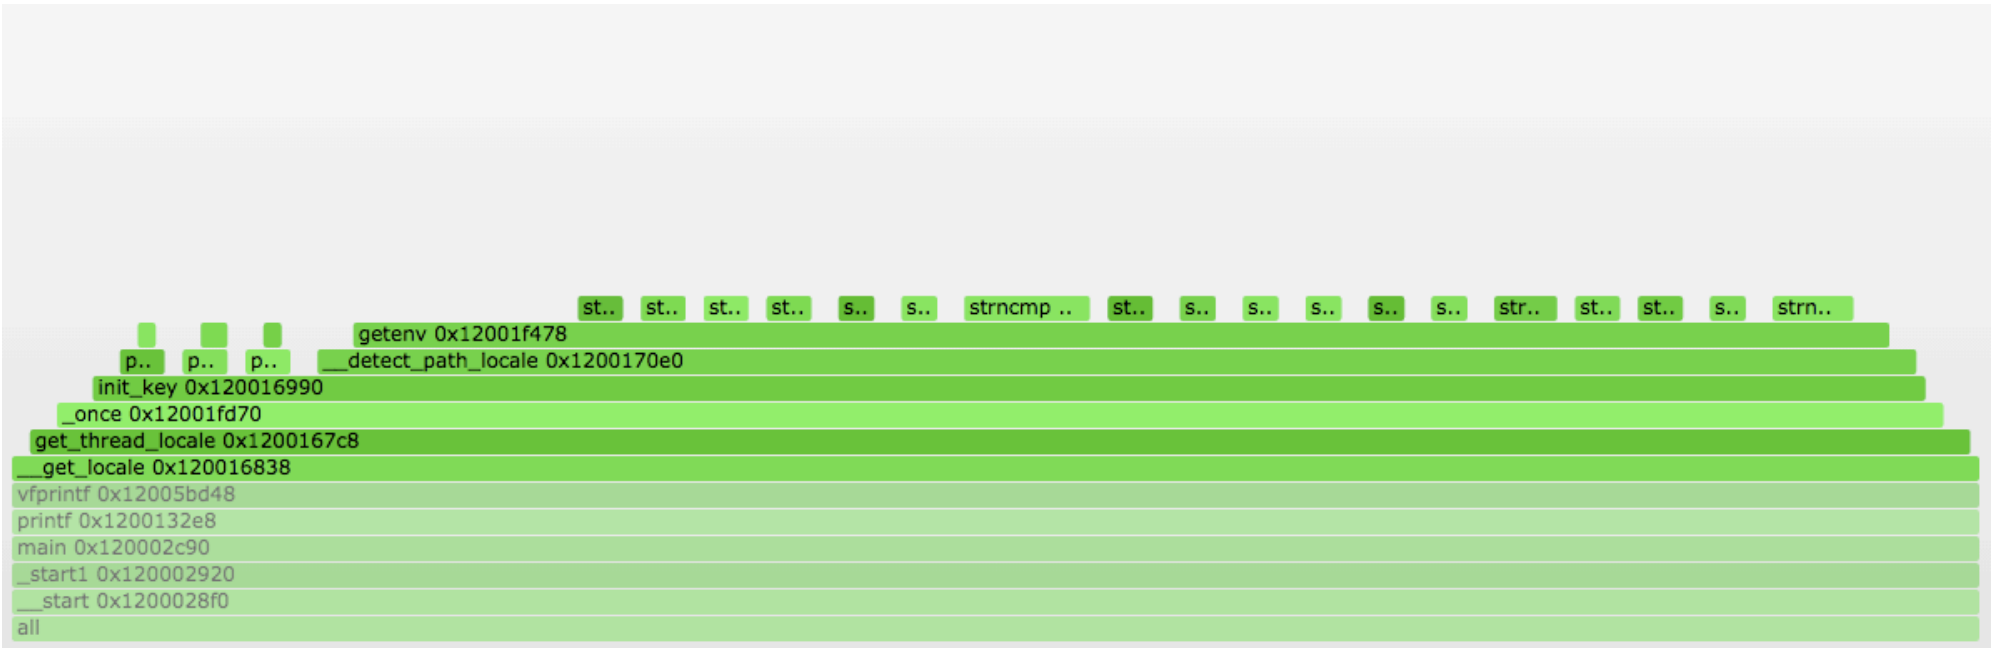
```
handle_argv(argc, argv, env);  
_init_tls();  
handle_static_init(argc, argv,  
env);
```

```
exit(main  
}
```


Calls constructors and registers destructors. Four types supported:

- .pre_init_array section
- _init() function
- .ctors section (via _init())
- .init_array section

__get_locale()

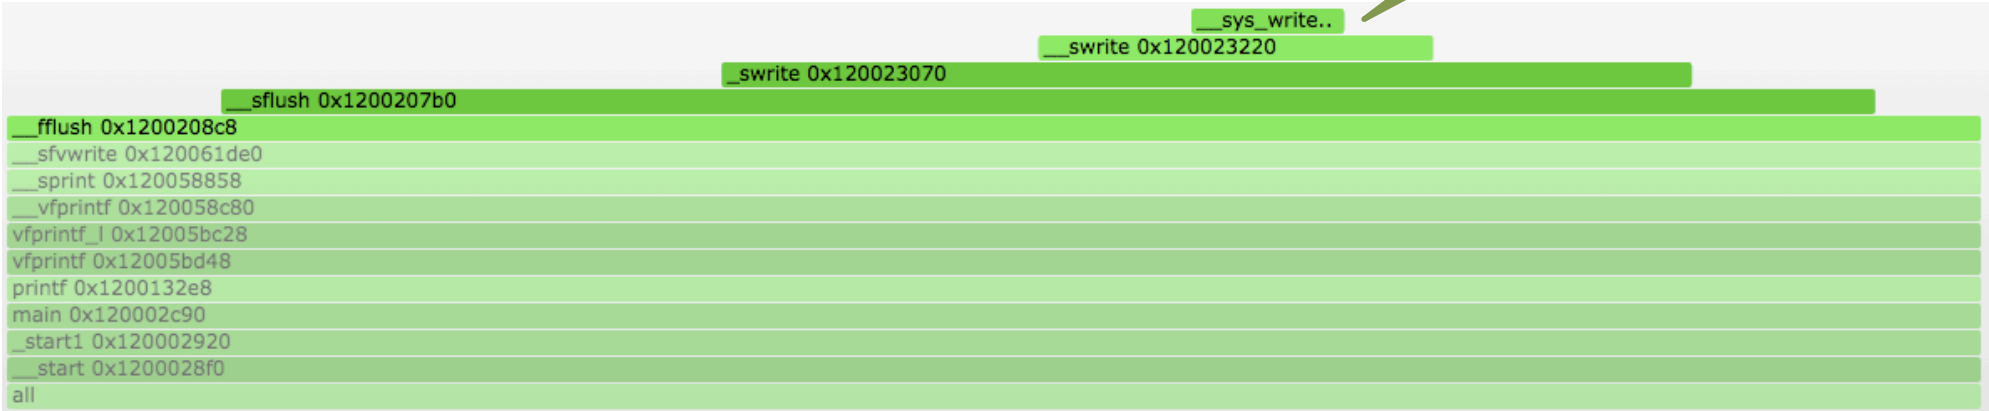


__vfprintf()



__flush()

The actual call
to write()



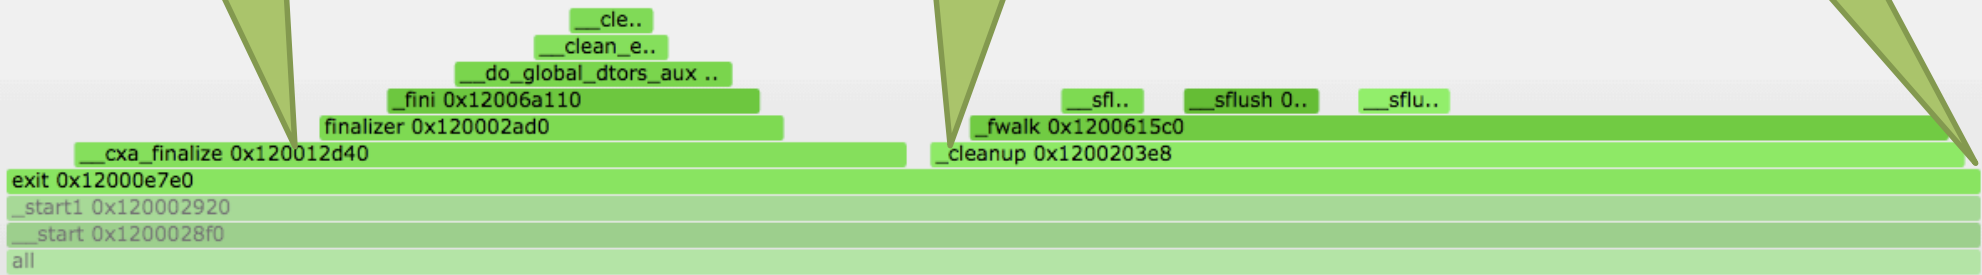
Hello World! 123

exit()

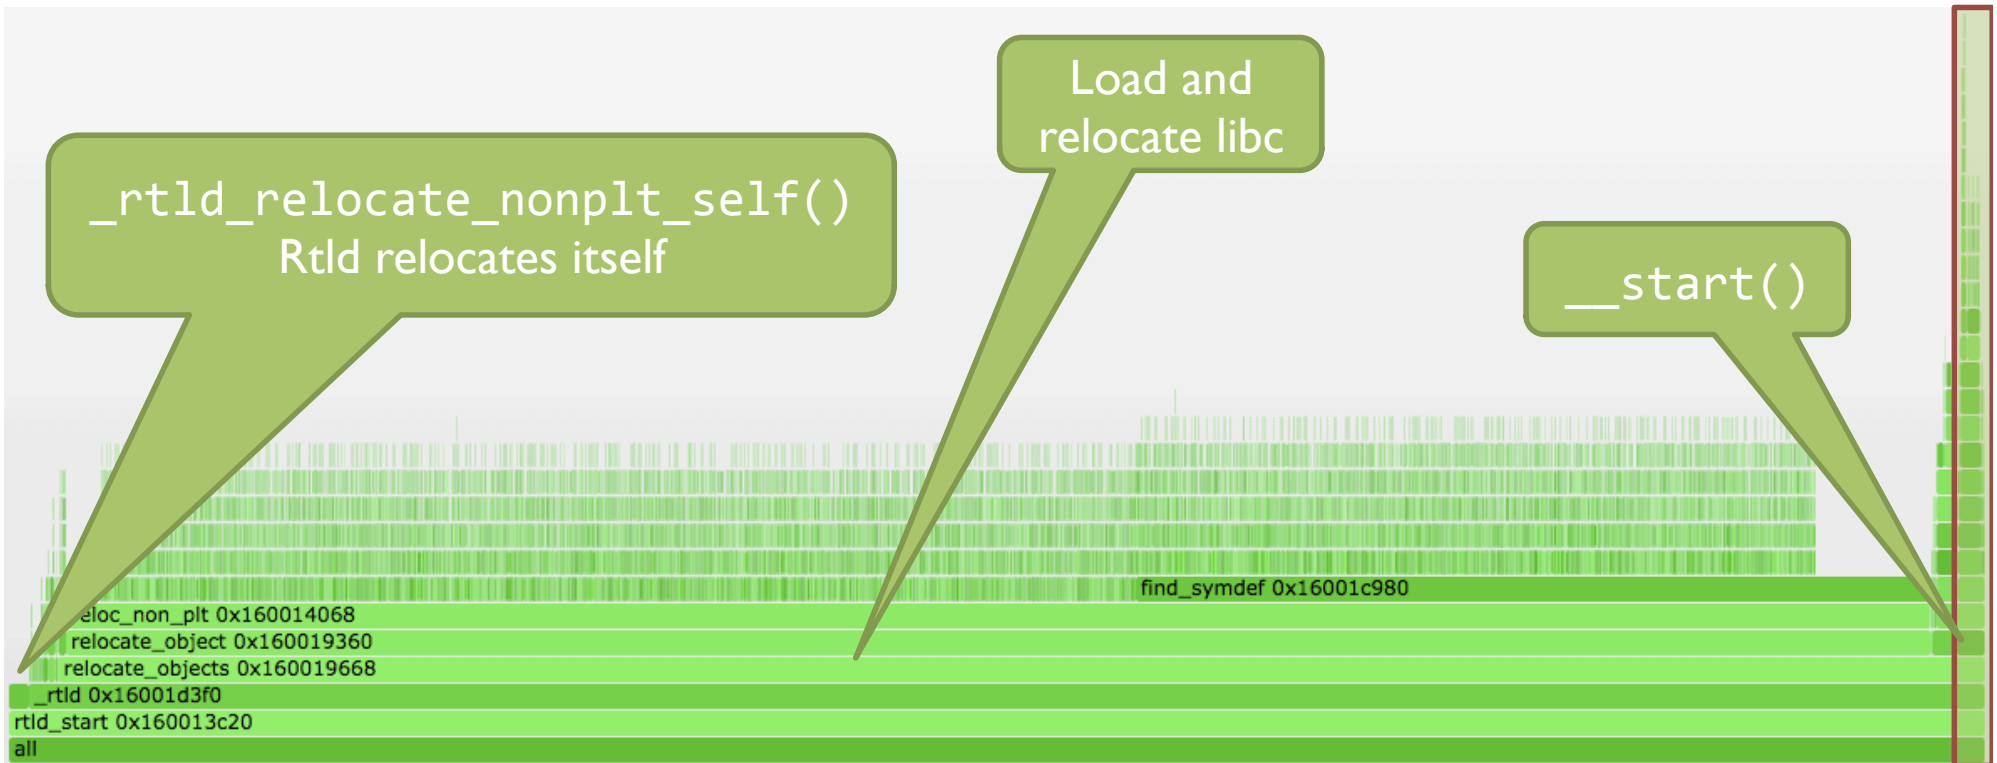
Call destructors registered with atexit()

Flush any unflushed FILEs

Call _exit()



Dynamic binary



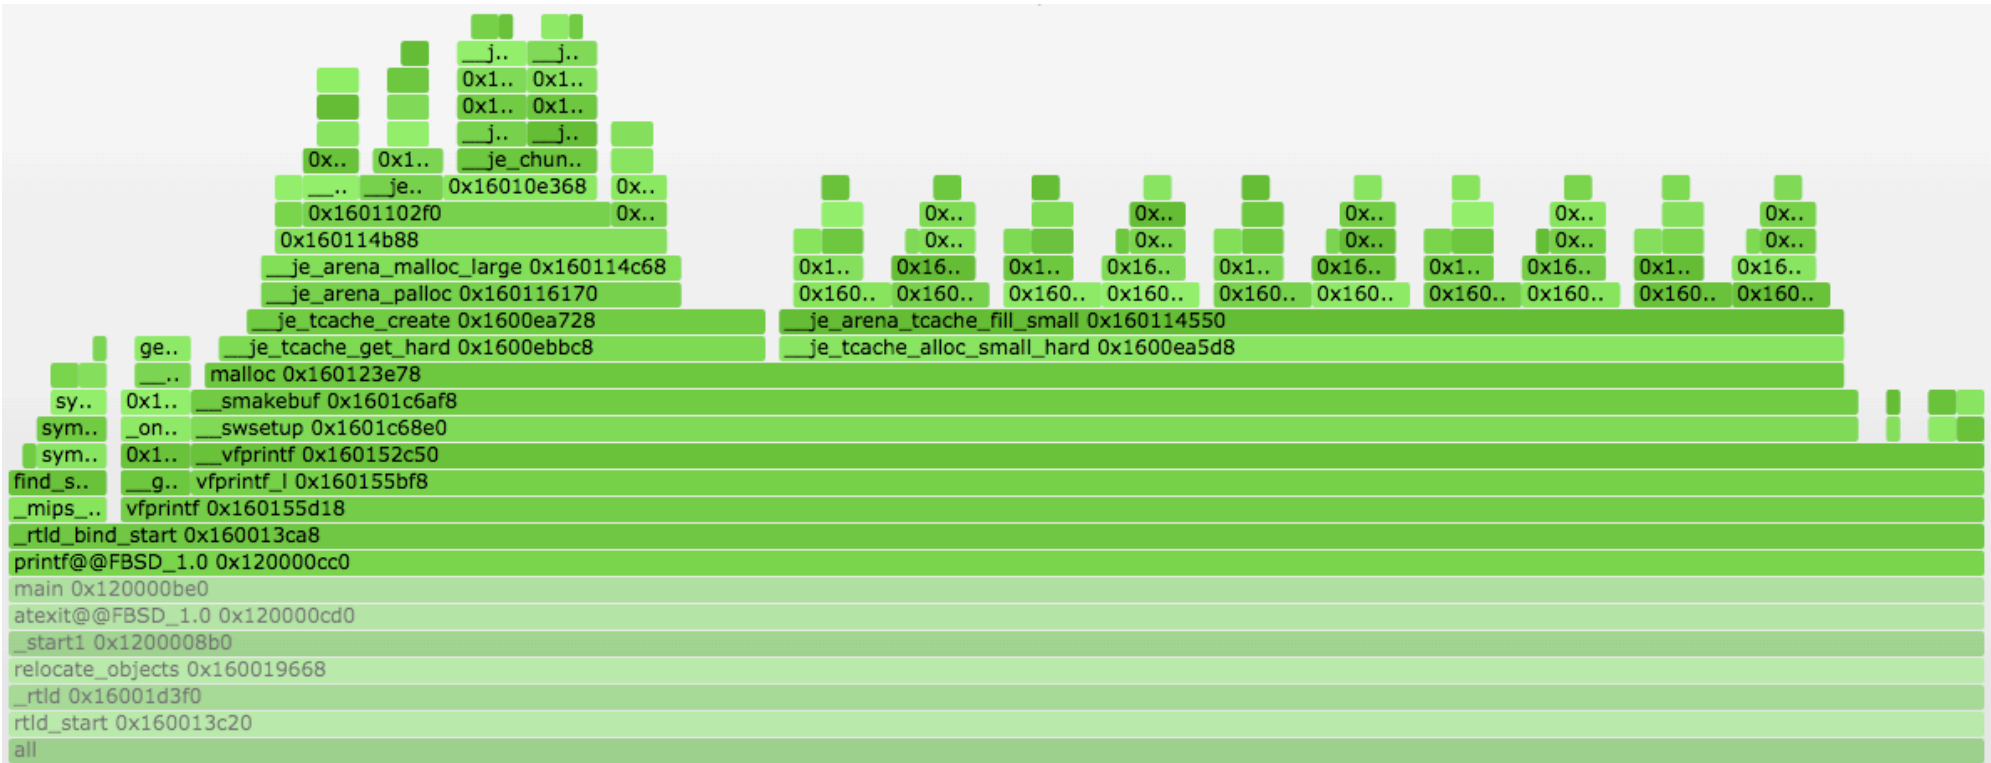
_mips_rtld_bind()

```

matched_symbol 0x16001c00
symlook_obj 0x16001ba30 symlook_obj 0x16001ba30
symlook_list 0x16001be78
symlook_global 0x16001c0d0
symlook_init 0x160015360 symlook_default 0x16001c6b0
find_symdef 0x16001c980
__mips_rtld_bind 0x160013fb8
__rtld_bind_start 0x160013ca8
printf@@@FBSD_1.0 0x120000cc0
main 0x120000be0
atexit@@@FBSD_1.0 0x120000cd0
_start1 0x1200008b0
relocate_objects 0x160019668
__rtld 0x16001d3f0
rtld_start 0x160013c20
all

```

printf()



Feedback requested

- Was the talk interesting and/or helpful?
- What didn't make sense?
- What would you like have learned more (or less) about?
- <http://bit.ly/bsdcan16-helloworld>
- brooks.davis@sri.com

