

# CTSRD

CRASH-WORTHY  
TRUSTWORTHY  
SYSTEMS  
RESEARCH AND  
DEVELOPMENT

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

(\*but were afraid to ask)

Brooks Davis  
SRI International

September 24th, 2016  
EuroBSDCon 2016



Approved for public release; distribution is unlimited. This research is sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL), under contract FA8750-10-C-0237. The views, opinions, and/or findings contained in this article/presentation are those of the author(s)/presenter(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.



UNIVERSITY OF  
CAMBRIDGE

# 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, 13
    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/crteh.o /usr/lib/crtn.o
```



# Compiler runtime support

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

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



# Code and images online

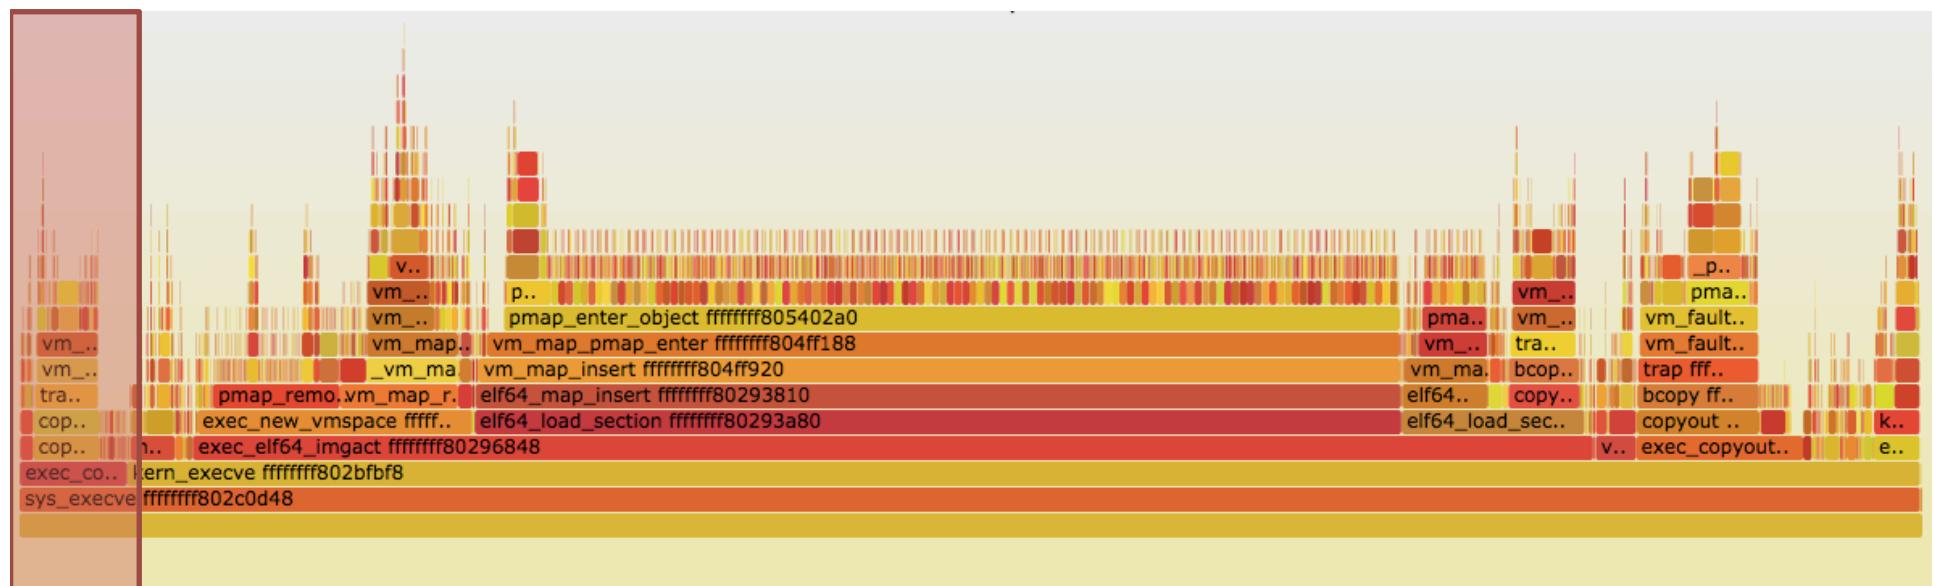
[https://people.freebsd.org/~brooks/talks/eurobsdcon2016-  
helloworld](https://people.freebsd.org/~brooks/talks/eurobsdcon2016-helloworld)

or

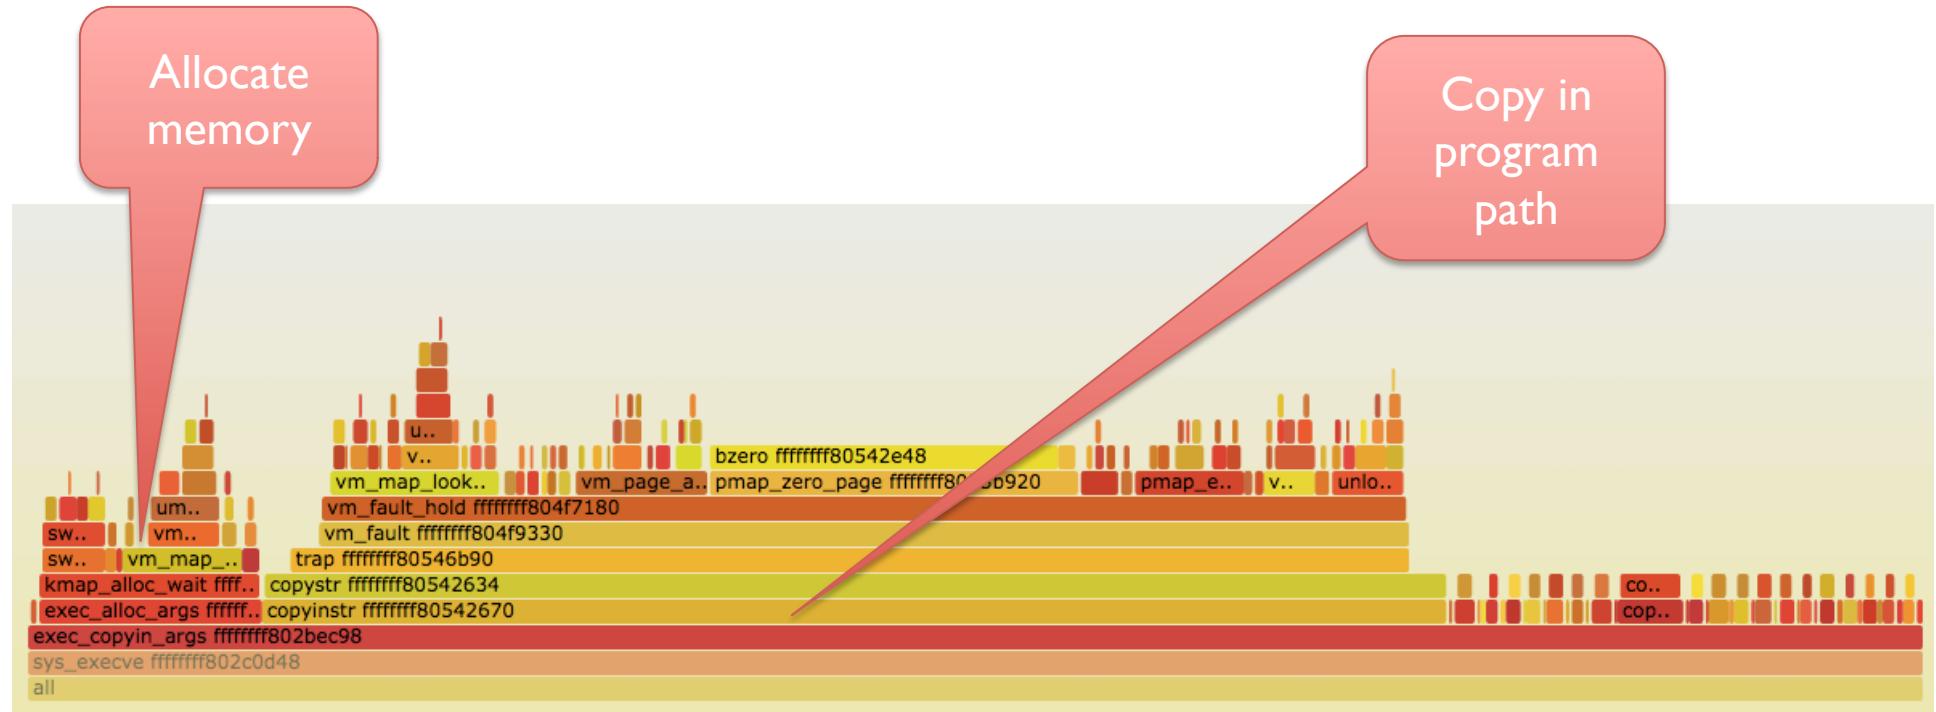
<http://bit.ly/helloworld-eurobsd>



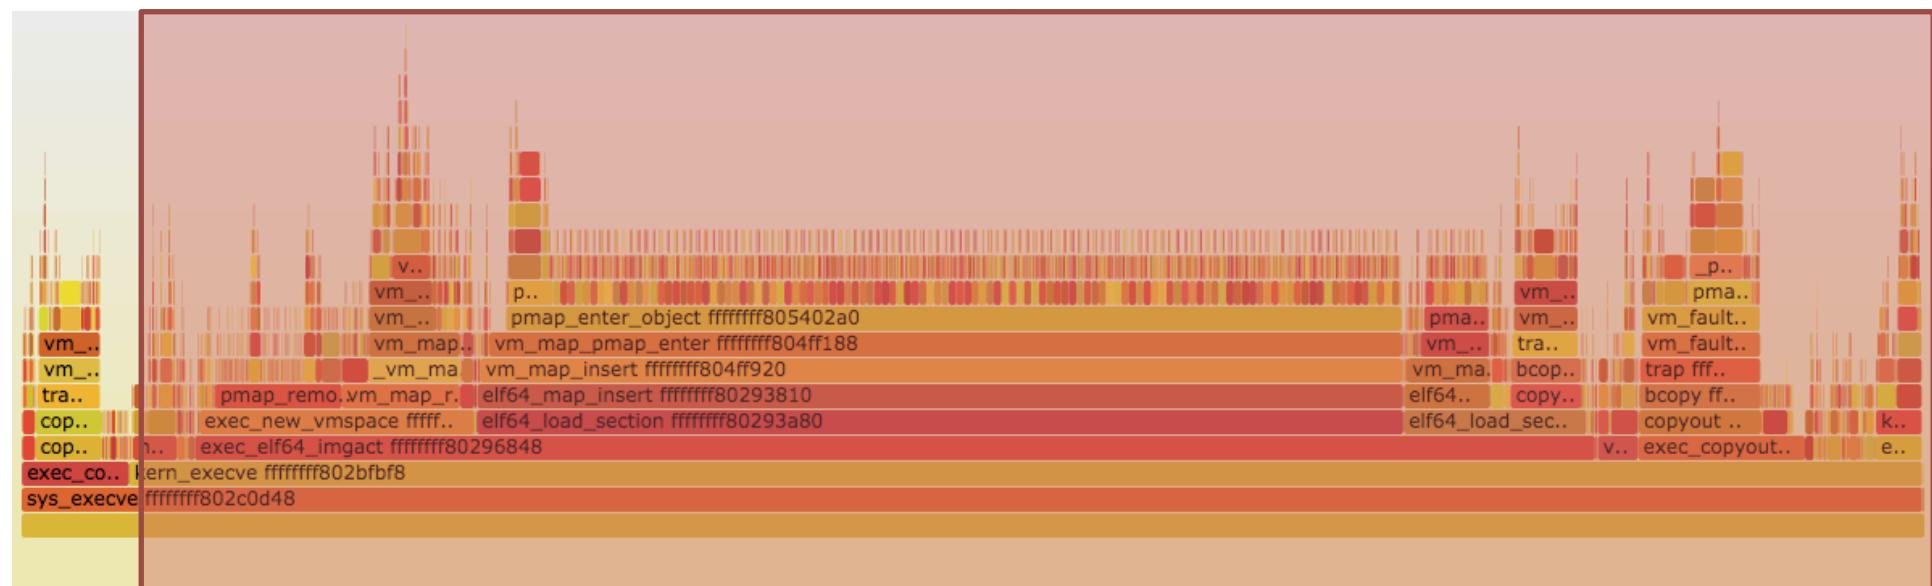
# execve()



# exec\_copyin\_args()

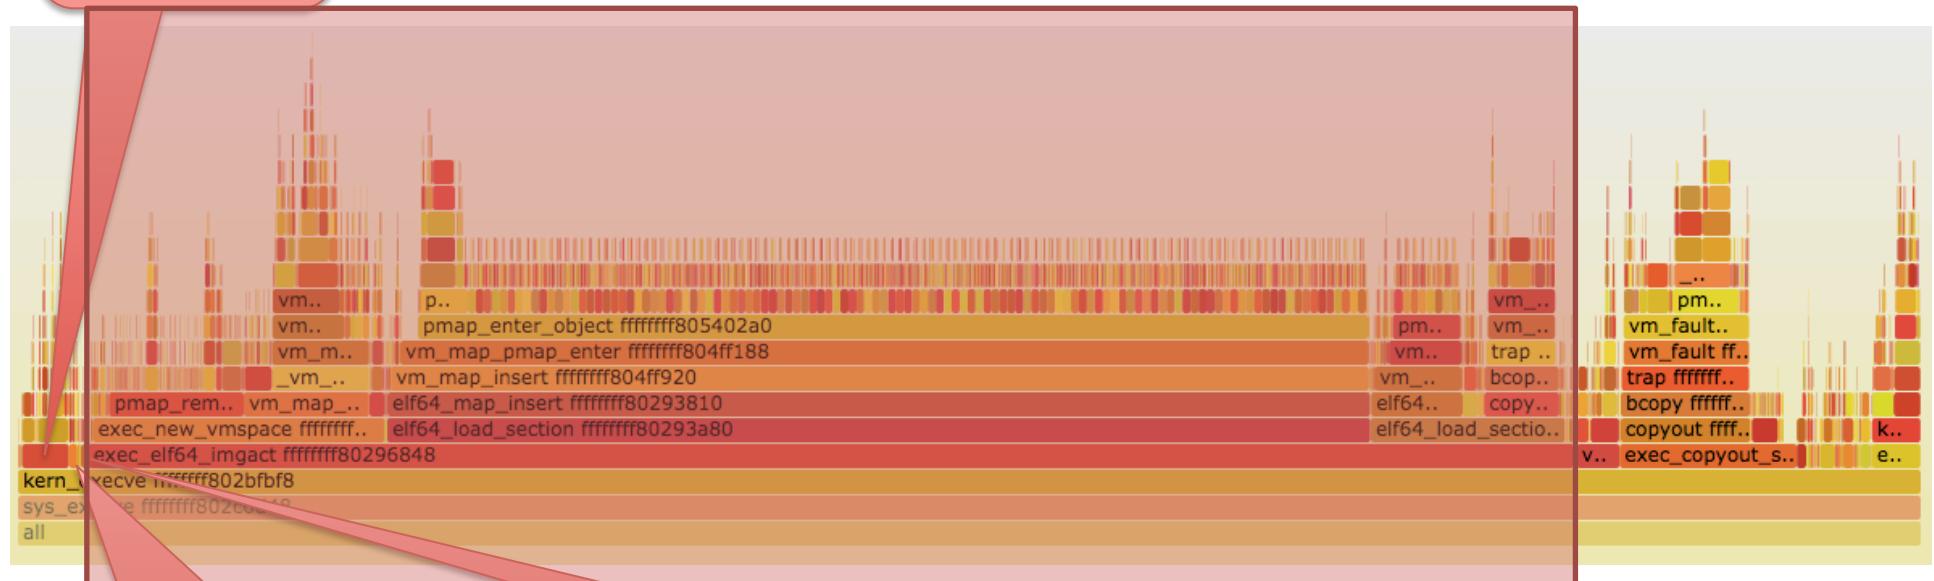


# sys\_execve()



# kern\_execve()

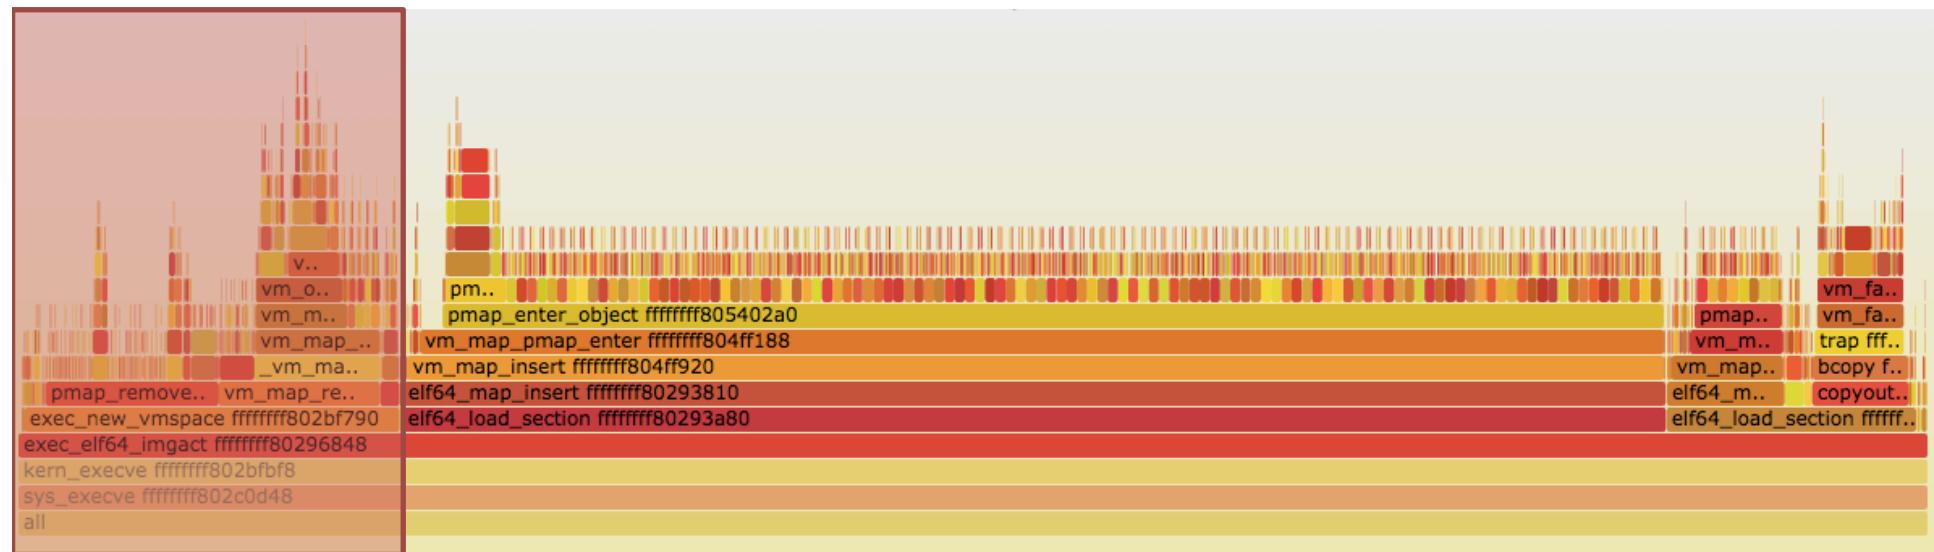
namei()  
Resolve  
path



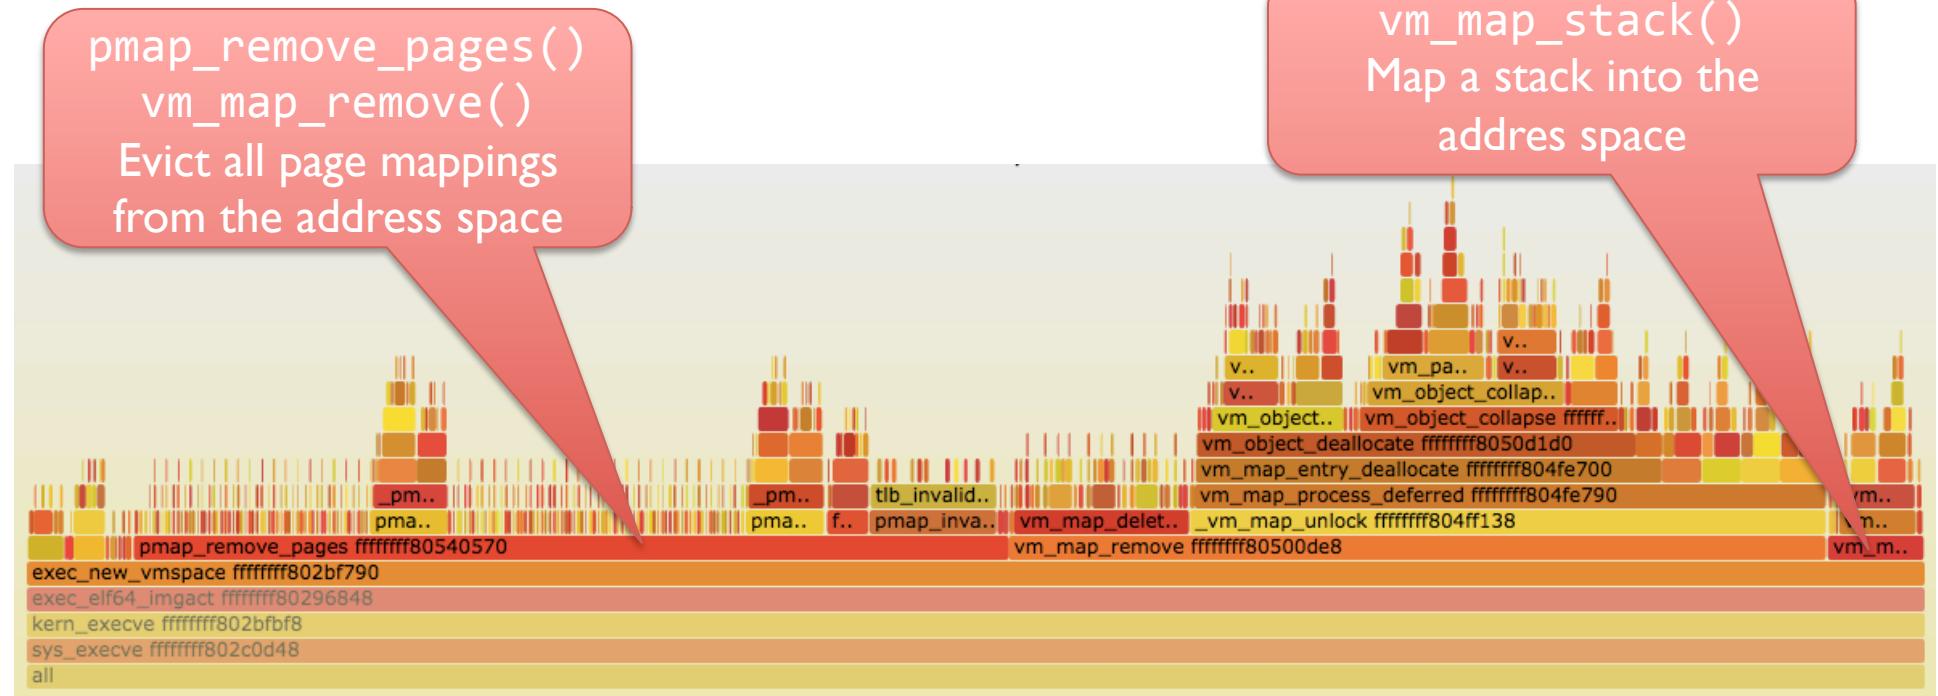
exec\_check\_permissions()  
Check that the file has the right  
permissions and open it.

exec\_map\_first\_page()  
Map the header into kernel  
memory.

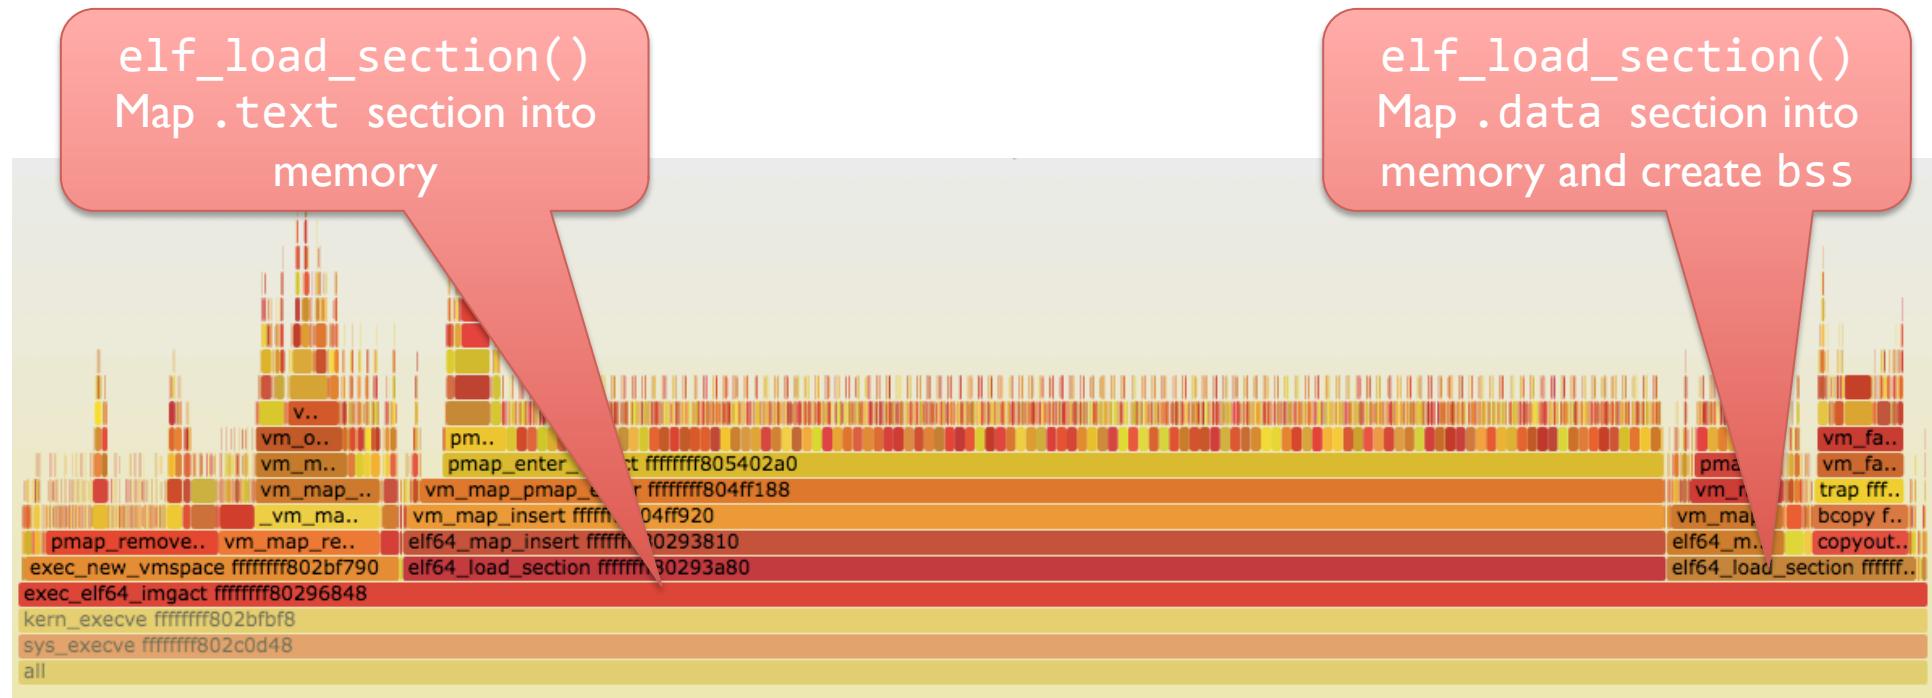
# exec\_elf64\_imgact()



# exec\_new\_vmspace()



# exec\_elf64\_imgact()



.text

.data

bss

Stack

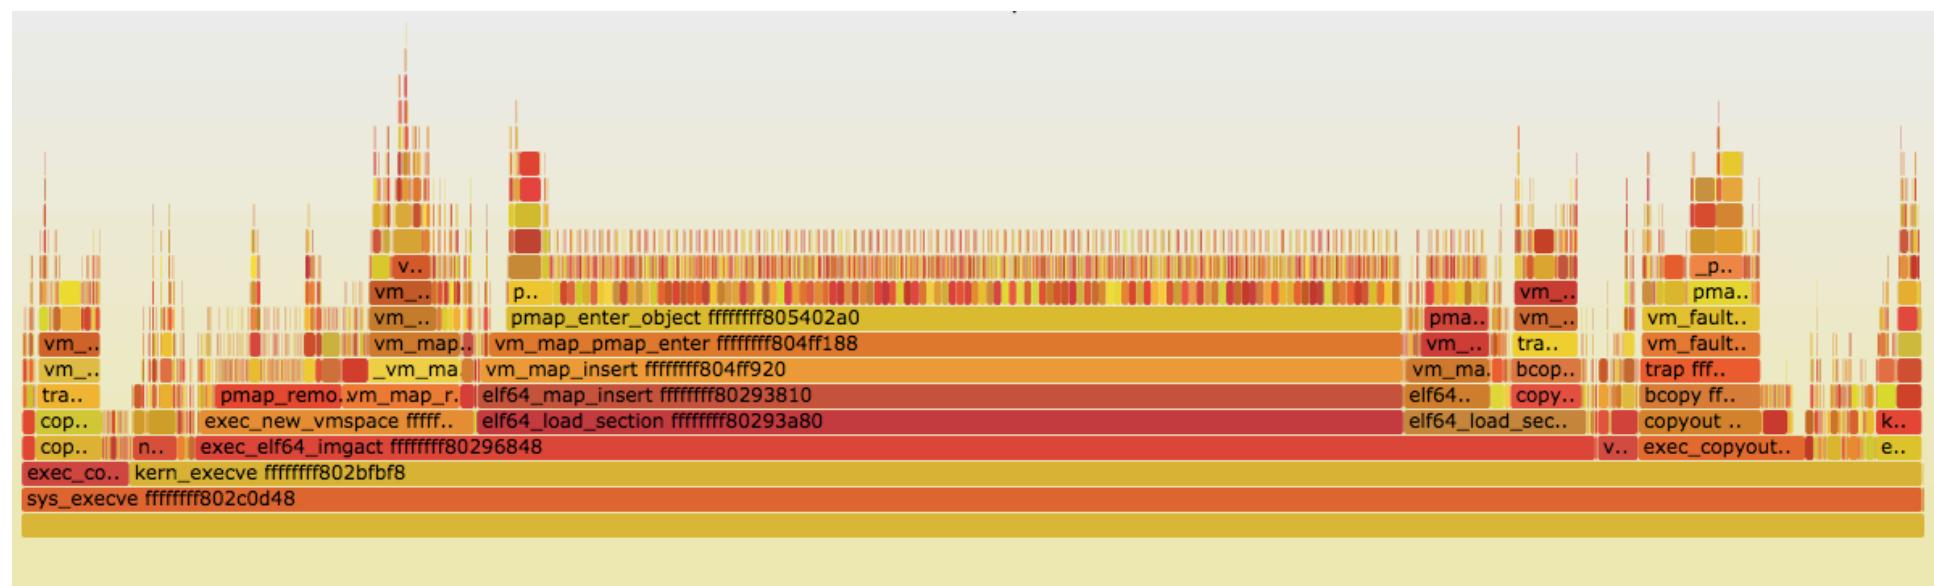
# kern\_execve()

`exec_copyout_strings()`  
`elf64_freebsd_fixup()`  
 Copy argv, envp, etc to the stack and adjust stack pointer.

`exec_setregs()`  
 Set initial register context to enter `__start()`.



# sys\_execve()



.text

.data

bss

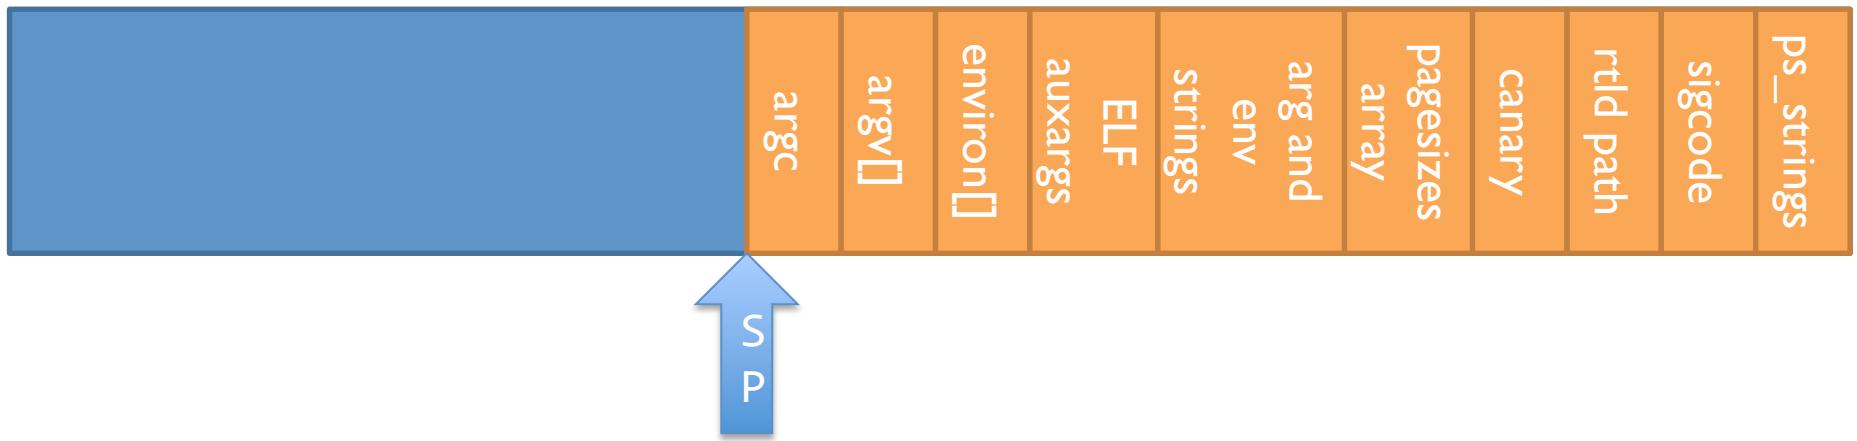
Stack

# 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()`

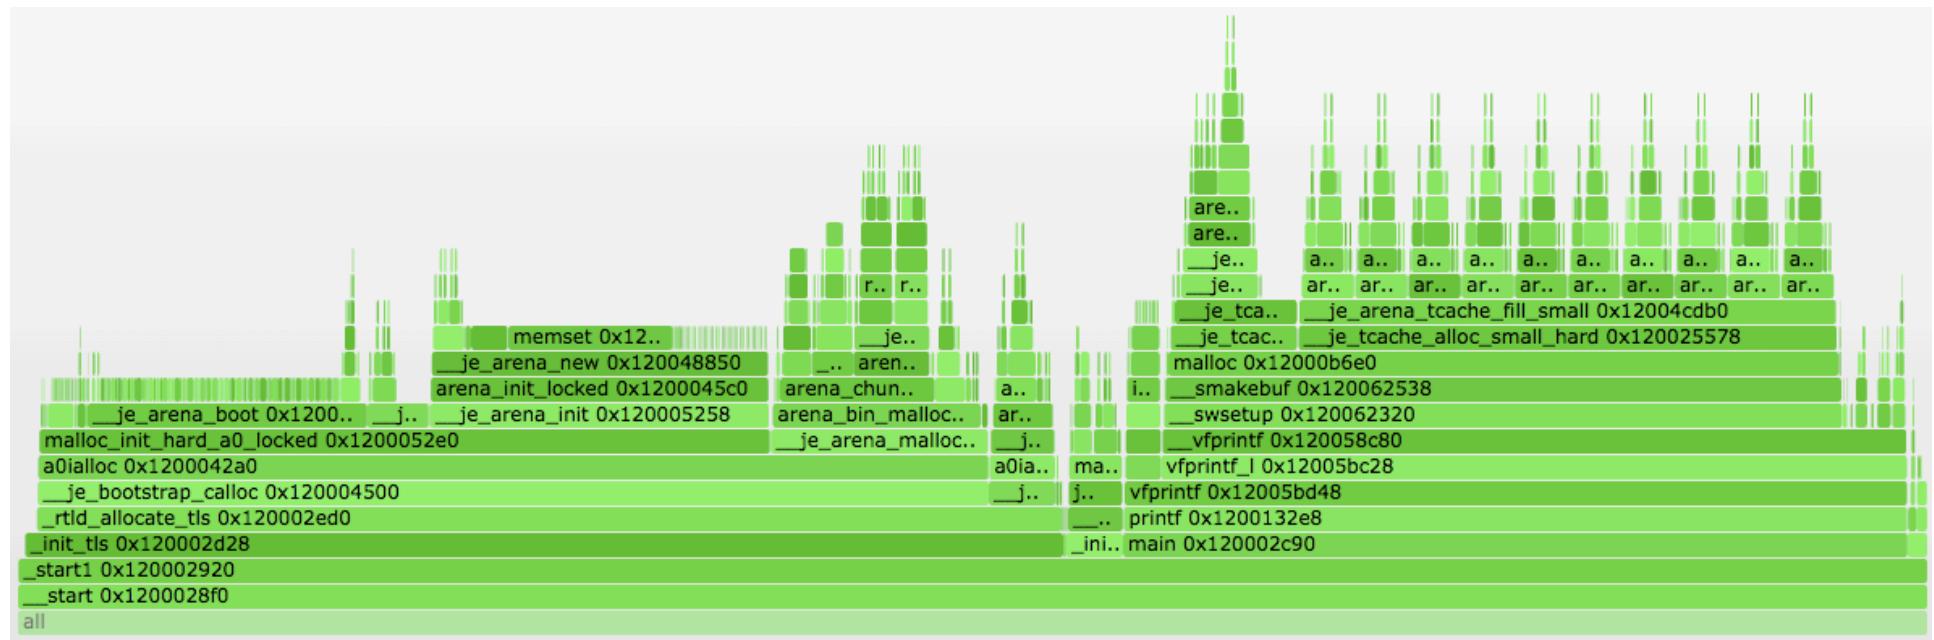


# SCO i386 ABI stack



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

# \_\_start()



## Most cycles spent in malloc()

# \_\_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

...

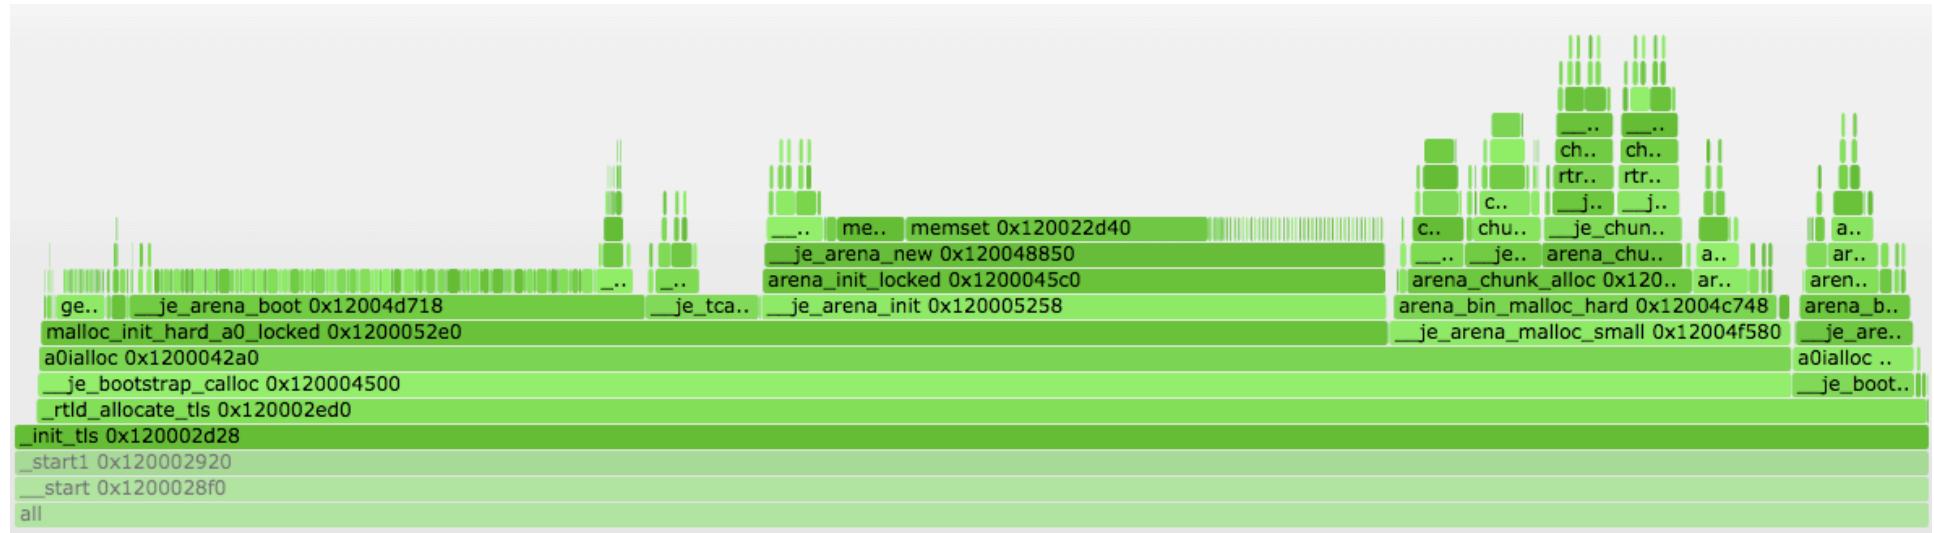
Set environ and  
\_\_progname variables.

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

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



# \_init\_tls()



Most cycles spent in malloc()

# \_init\_tls()

- Find the ELF auxargs vector

```
Elf_Addr *sp;
sp = (Elf_Addr *) environ;
while (*sp++ != 0)
    ;
aux = (Elf_Auxinfo *) sp;
```

# \_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 `_rtld_allocate_tls()`)
  - Allocates space
  - Copies initial values
- Set the TLS pointer

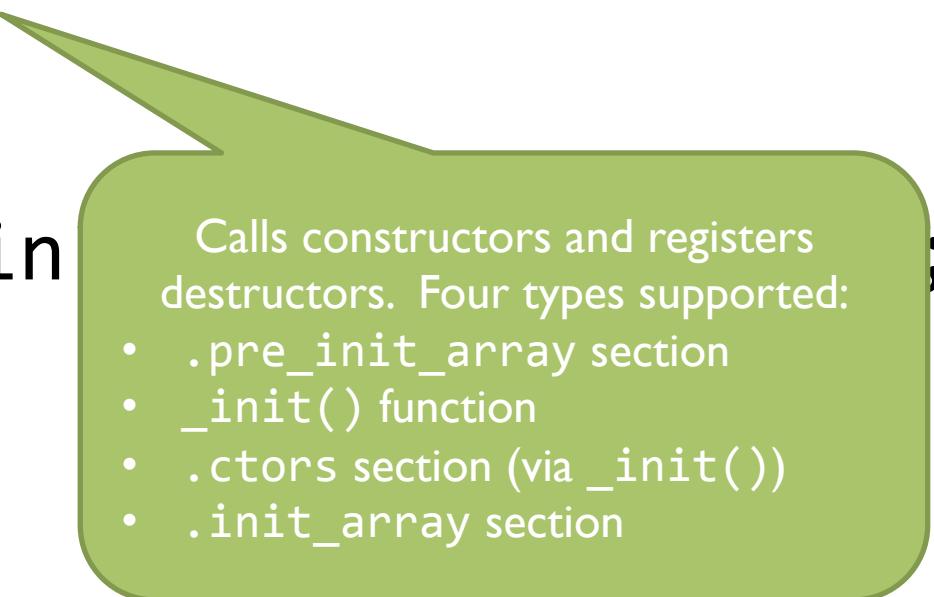


Uses JEMalloc, but  
JEMalloc uses TLS!

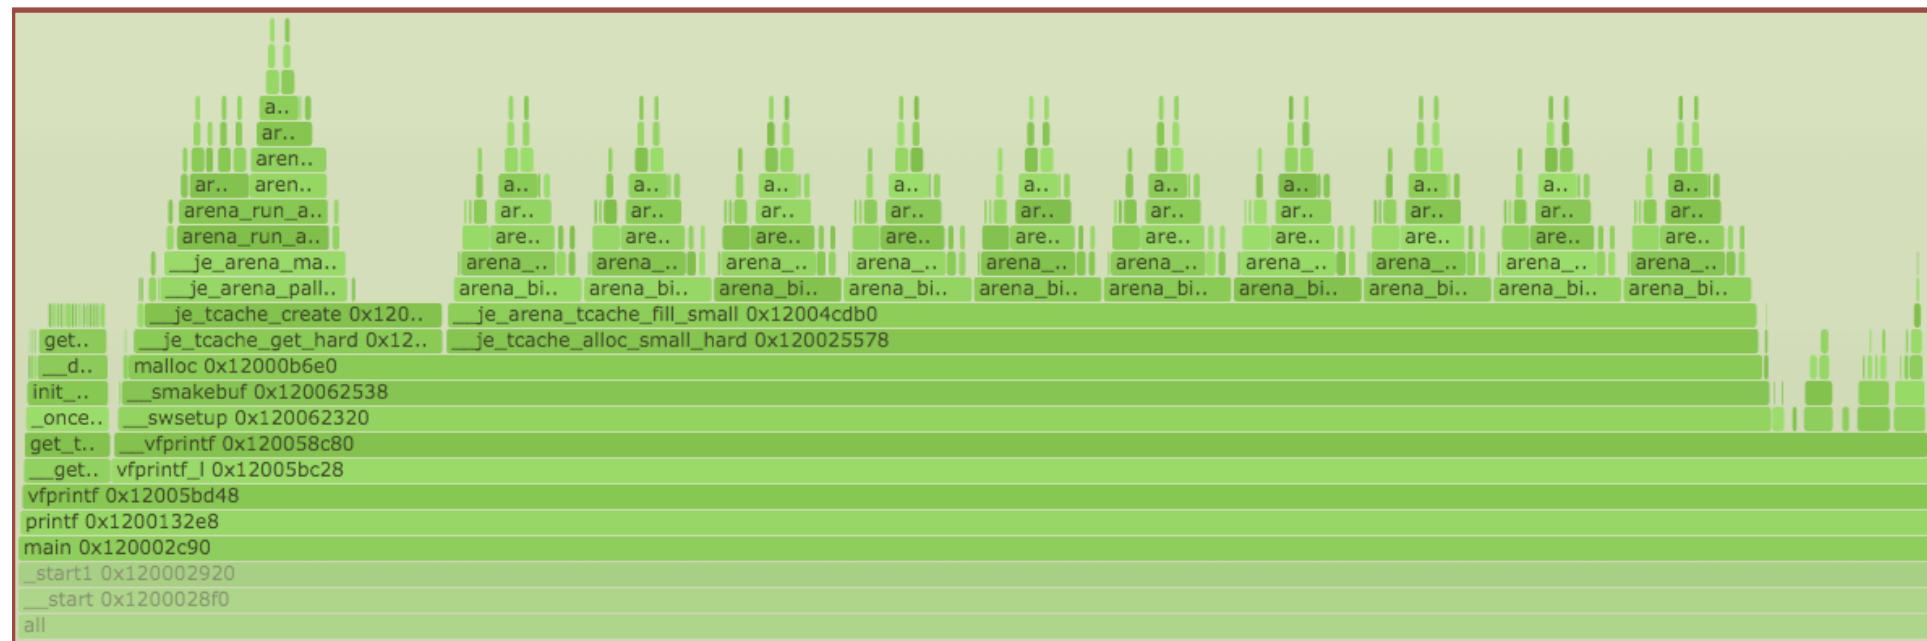
# \_\_start() 2/2

...

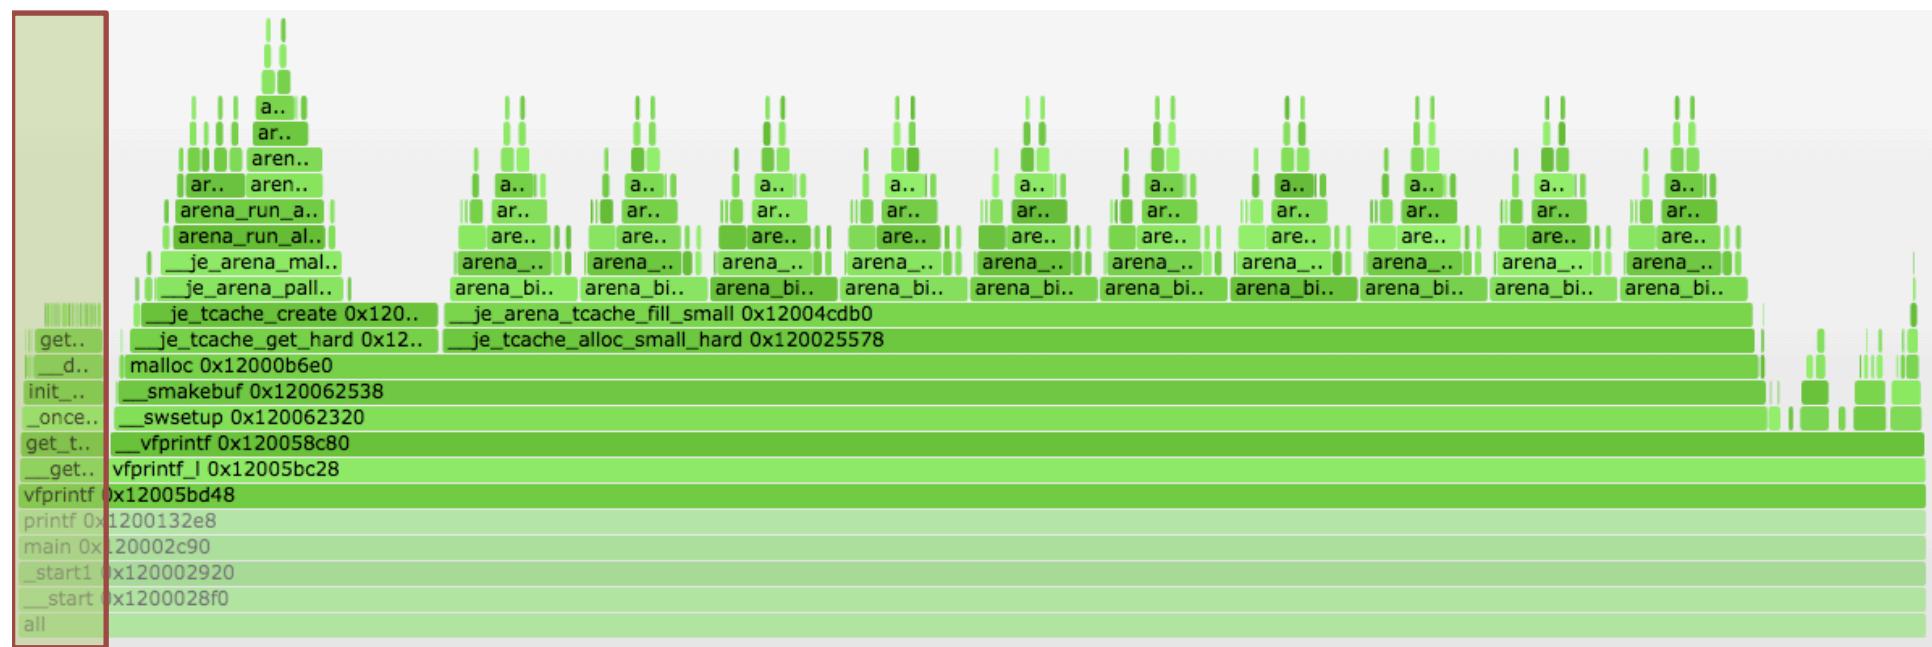
```
handle_argv(argc, argv, env);
_init_tls();
handle_static_init(argc, argv,
env);
exit(main
}
```



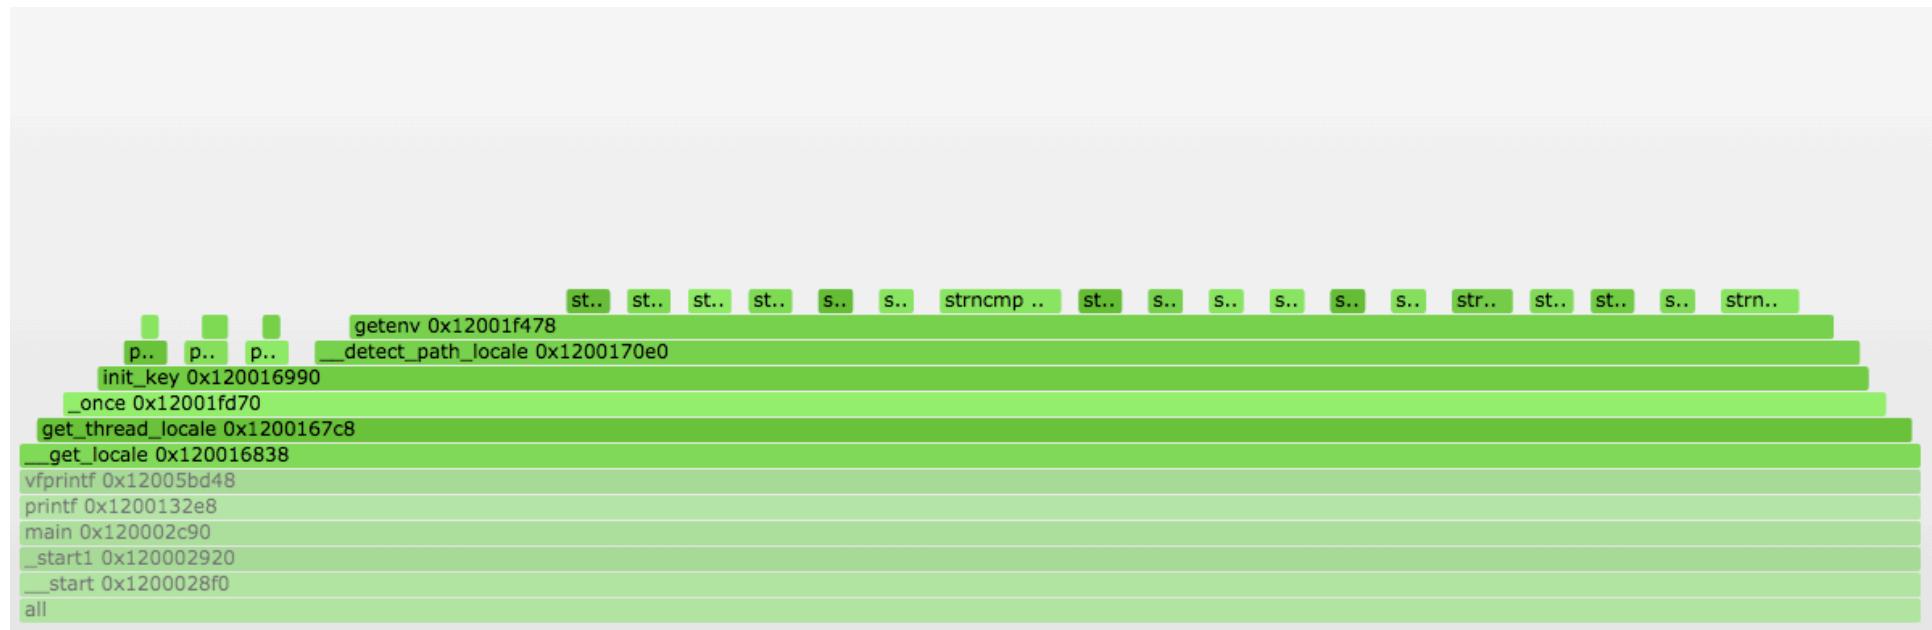
# main()



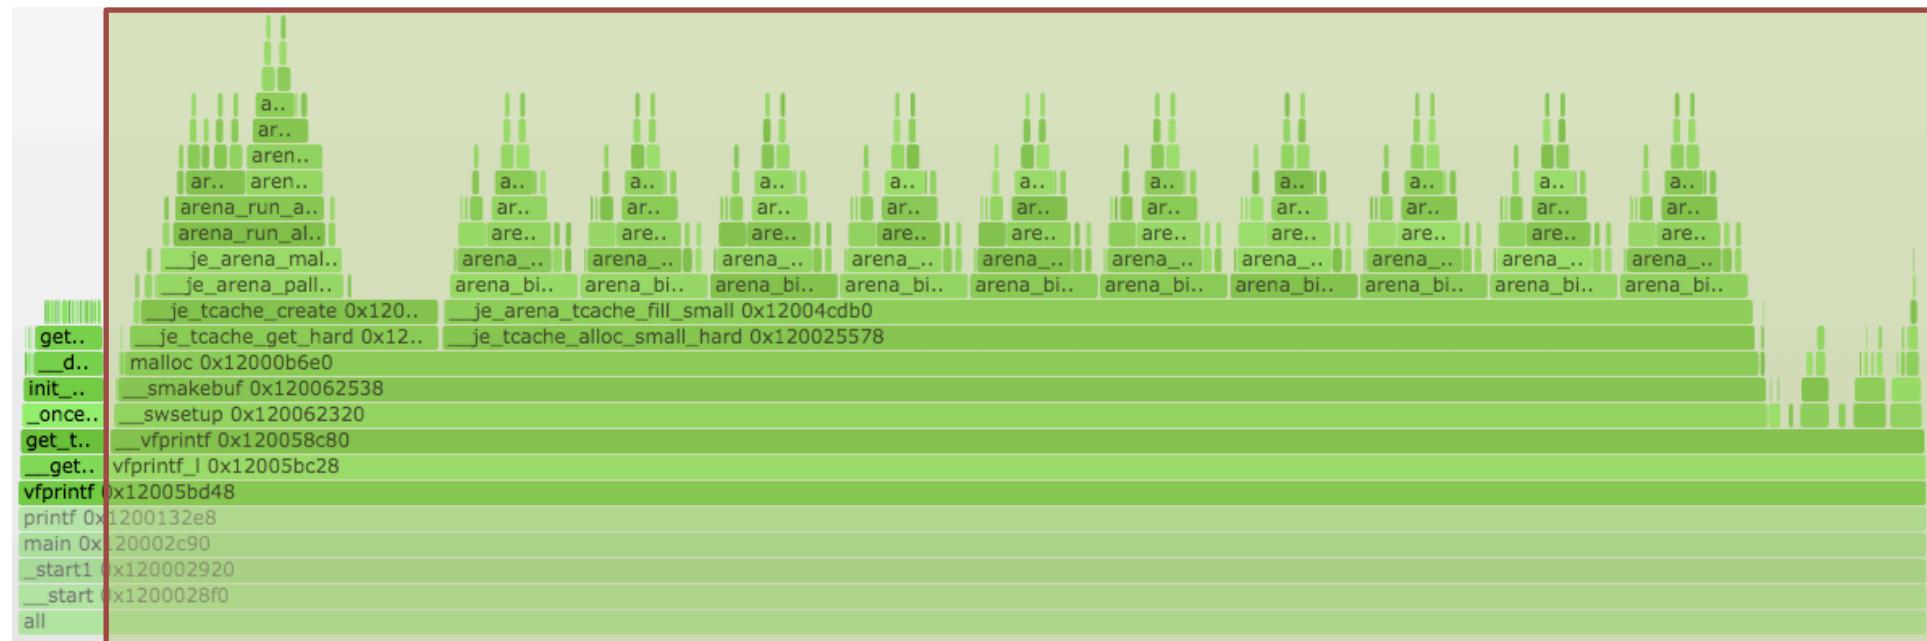
# vfprintf()



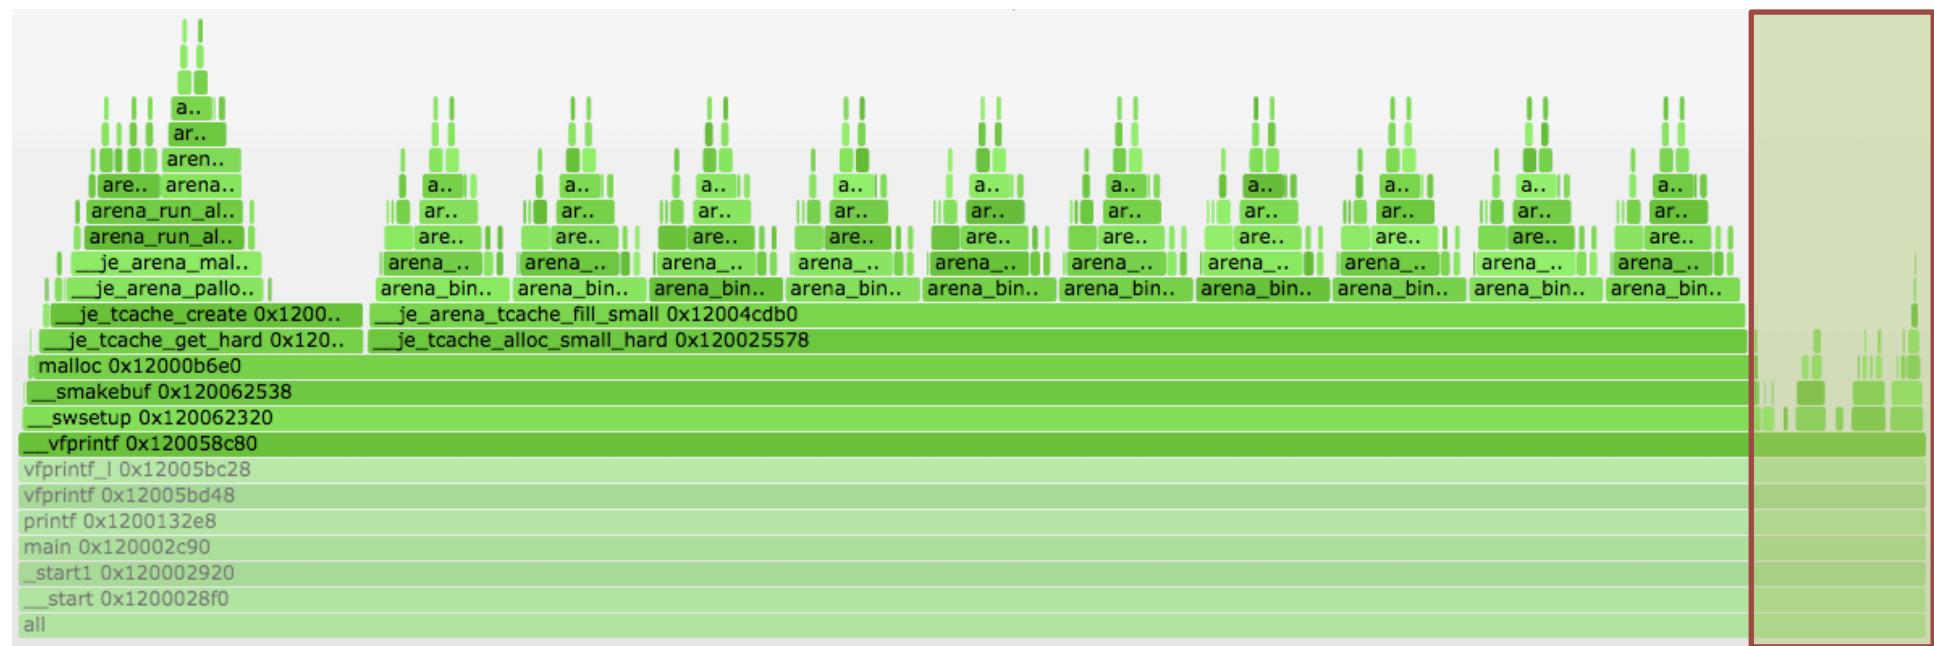
# \_\_get\_locale()



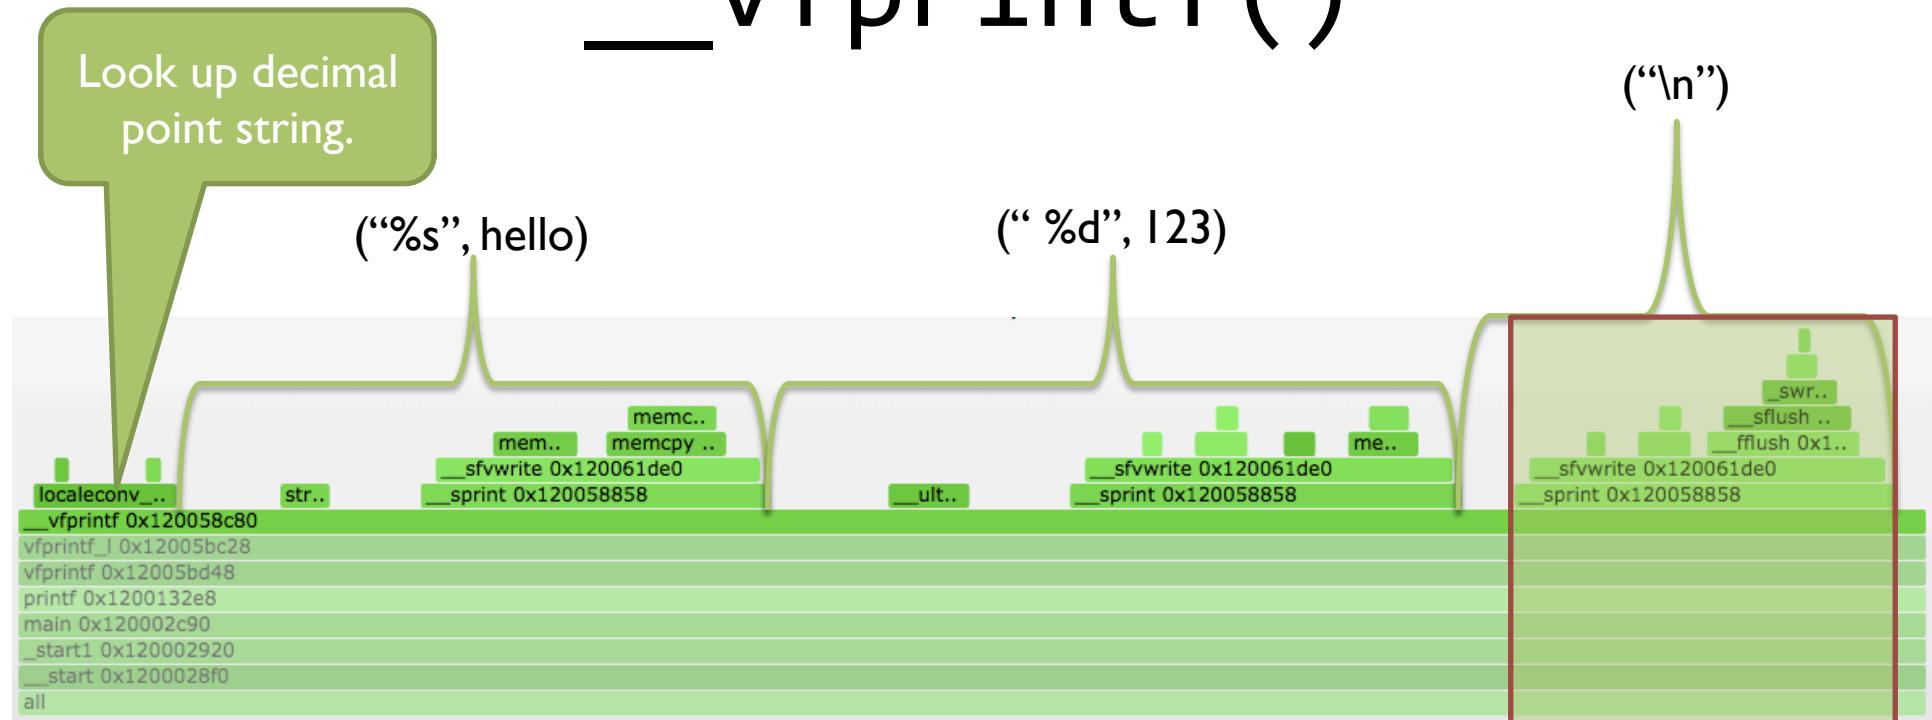
# vfprintf()



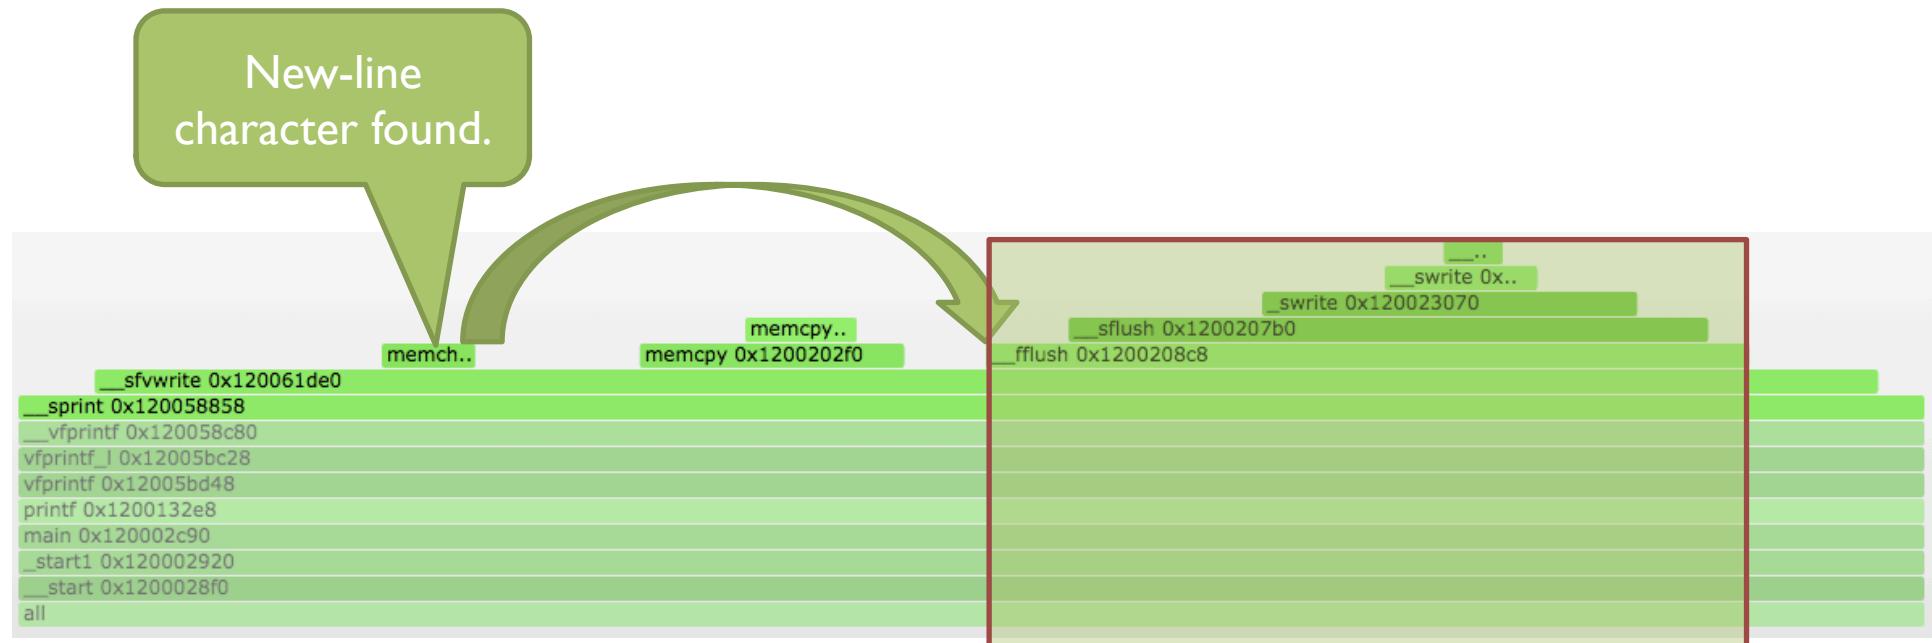
# \_vfprintf()



# \_vfprintf()

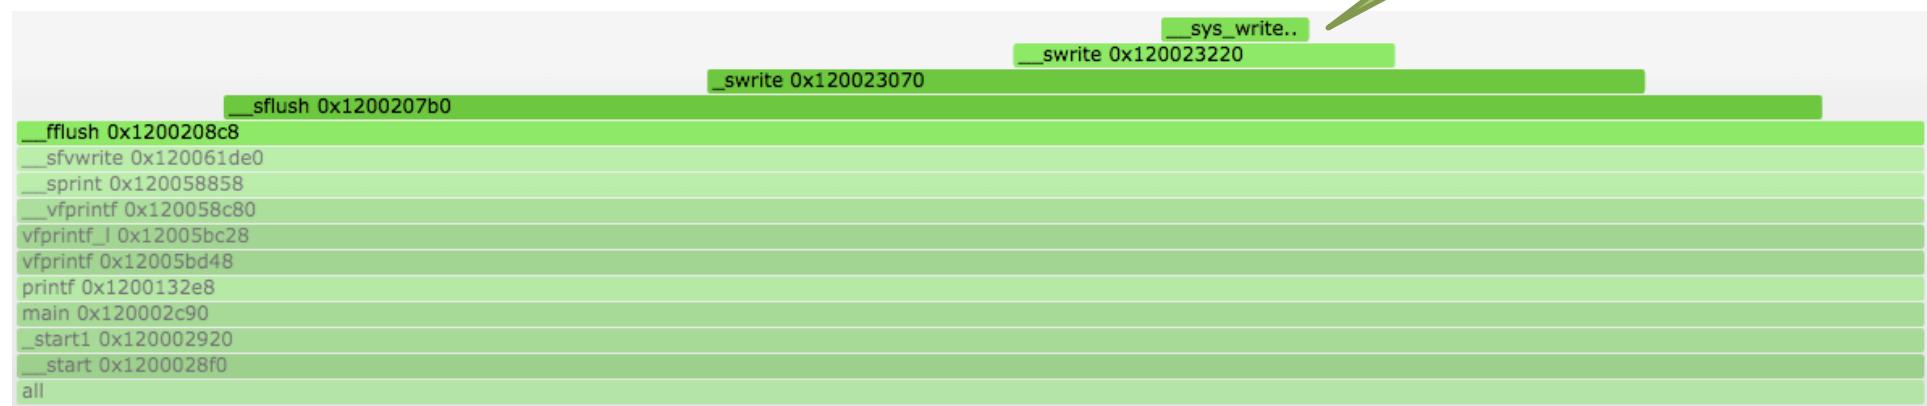


# \_\_sprint()



# \_\_flush()

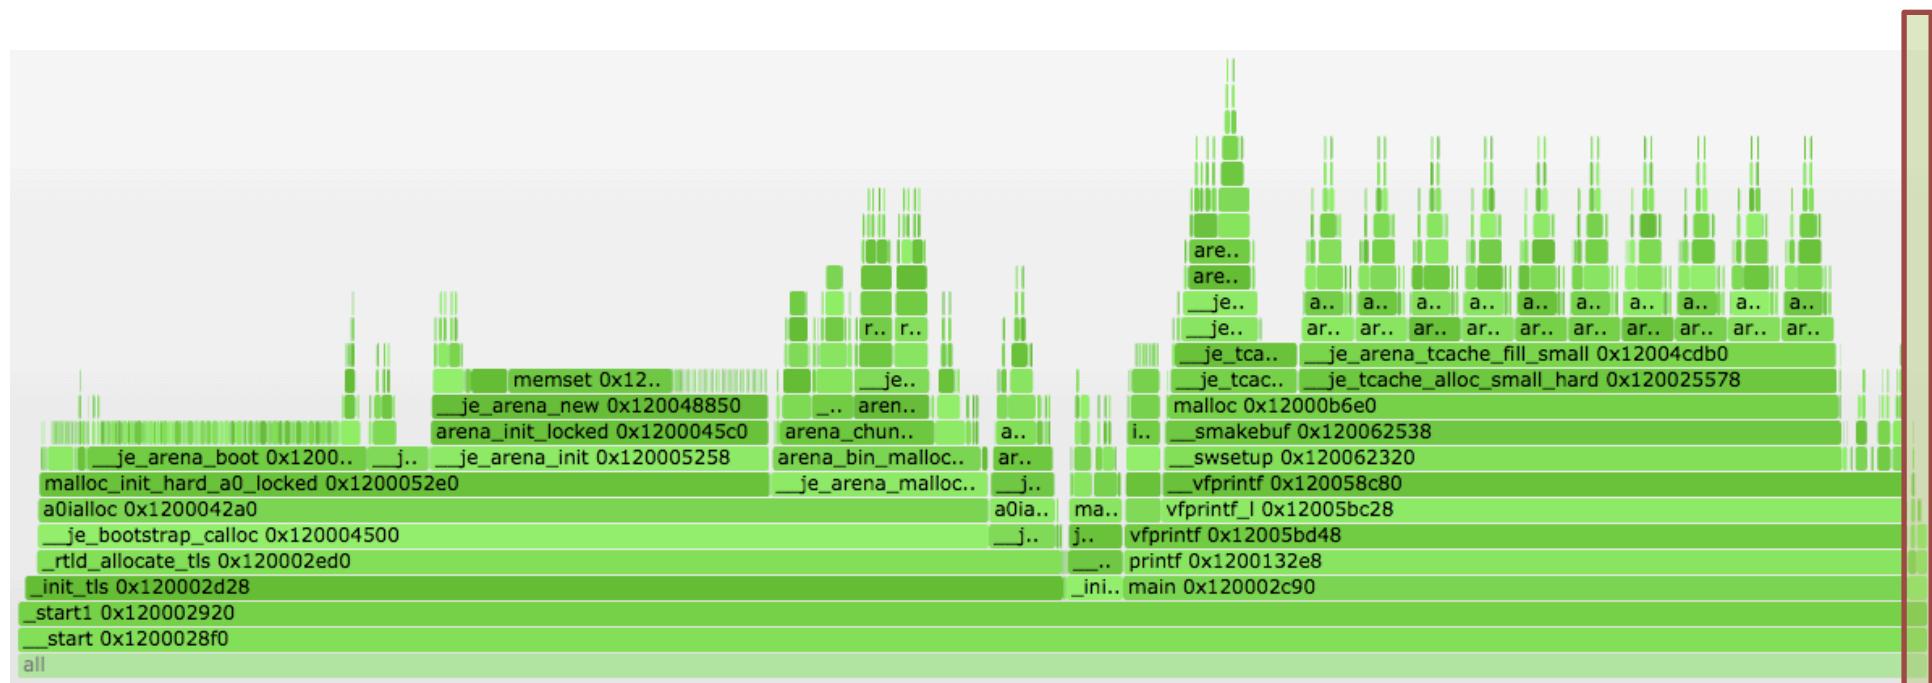
The actual call  
to write()



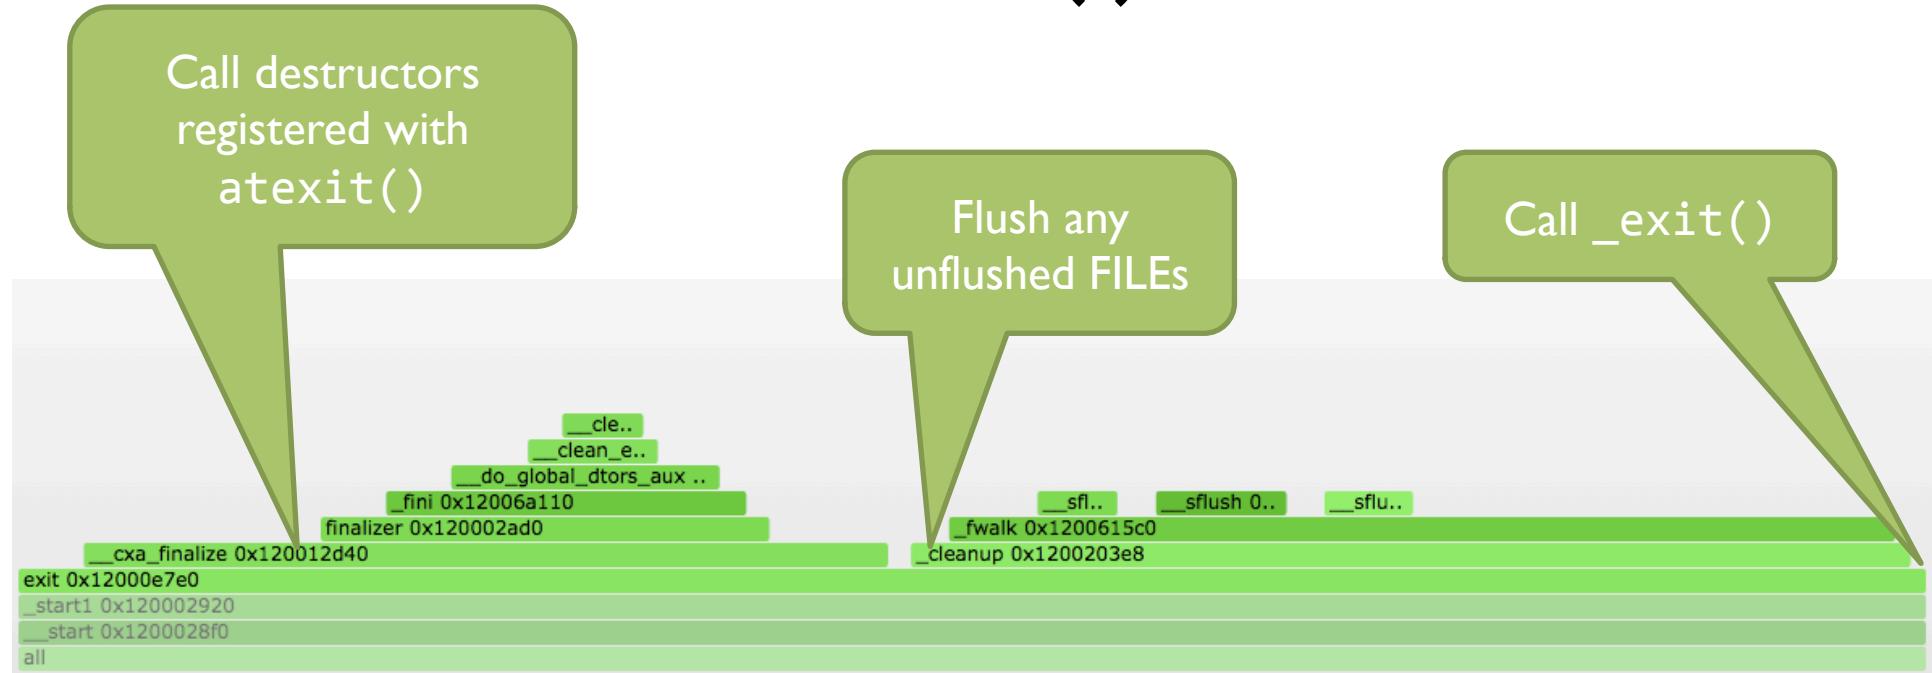
# Hello World! 123



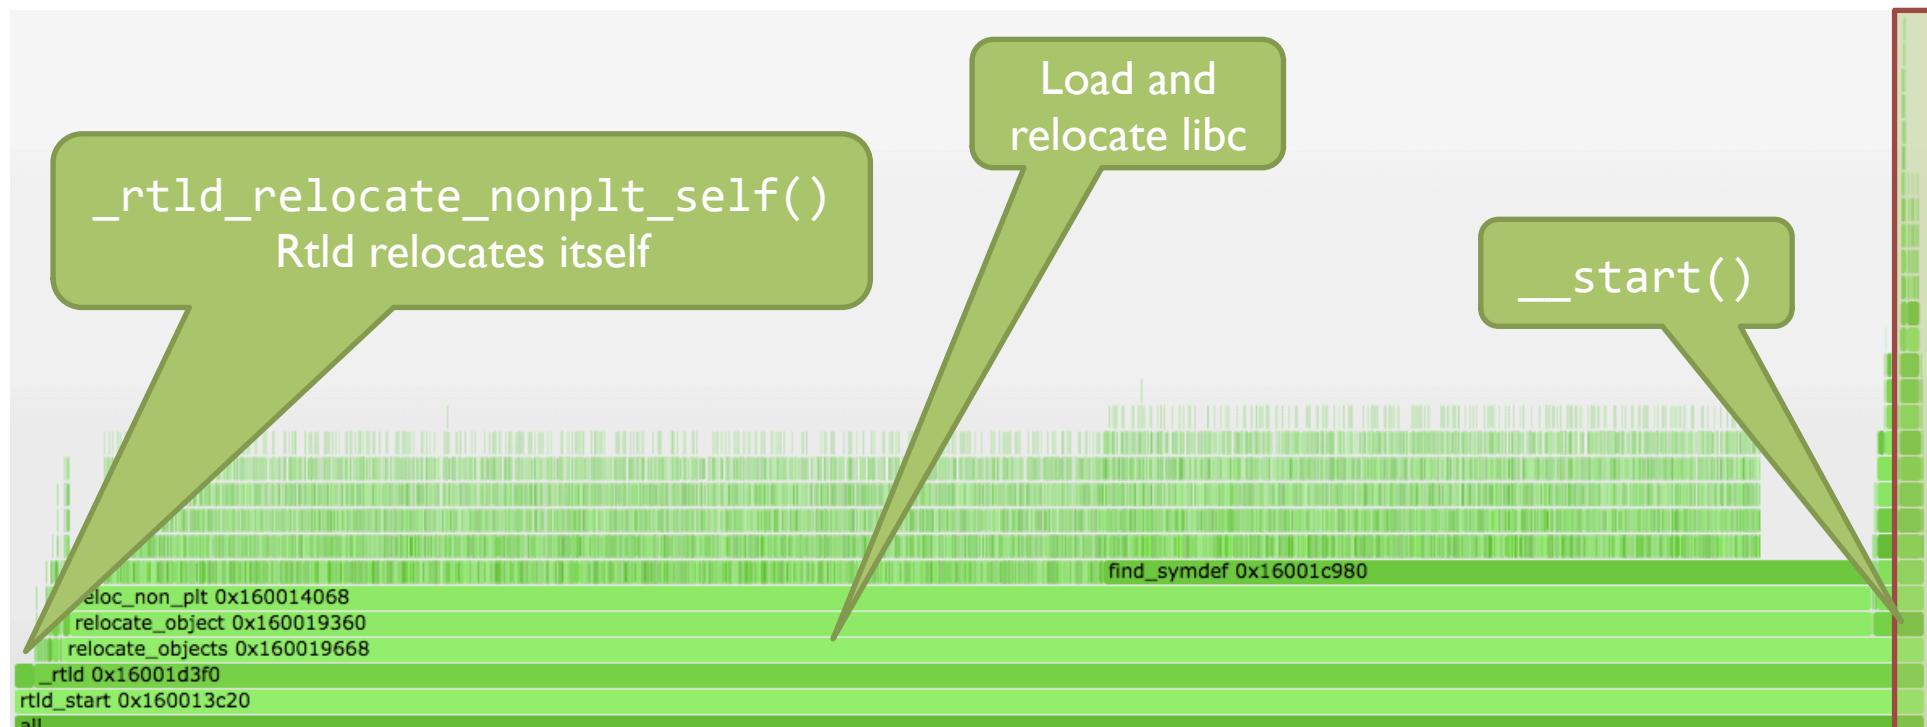
# \_\_start()



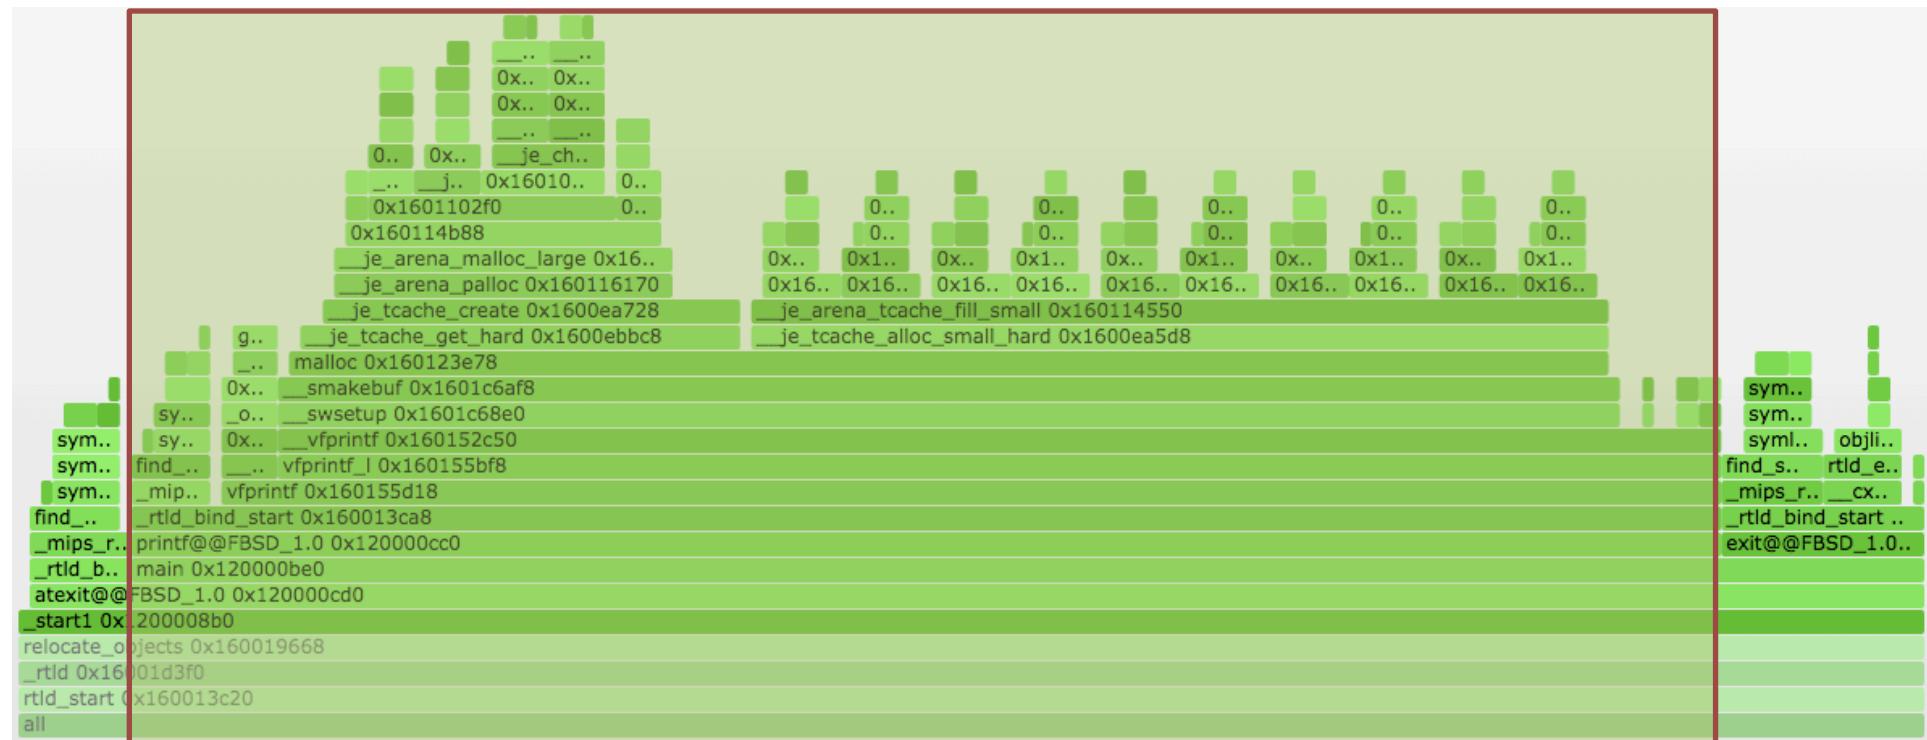
# exit()



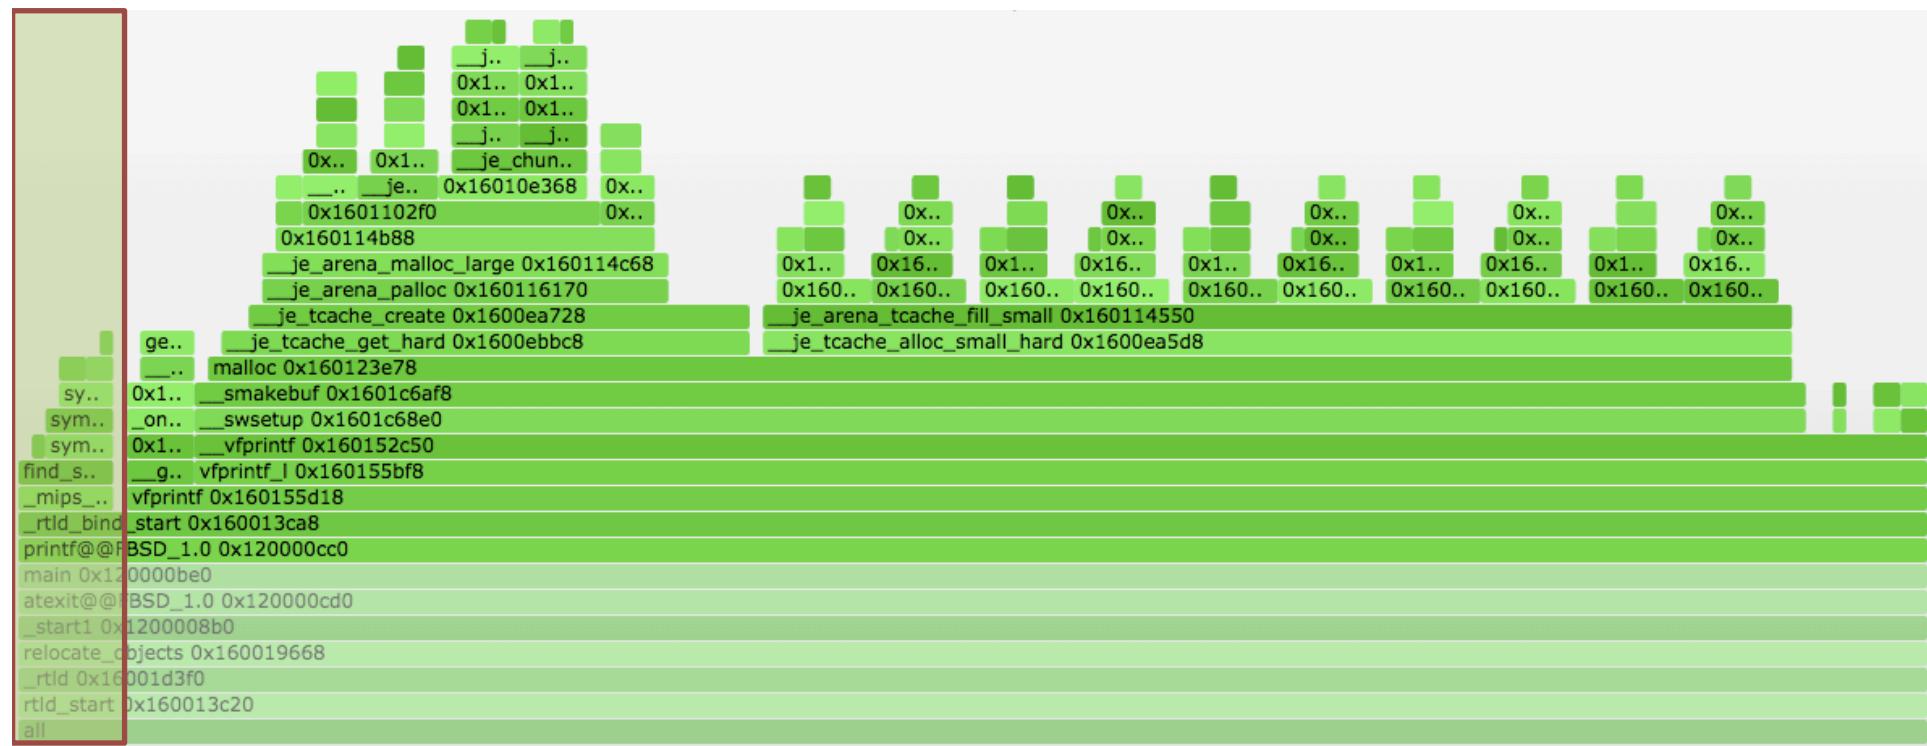
# Dynamic binary



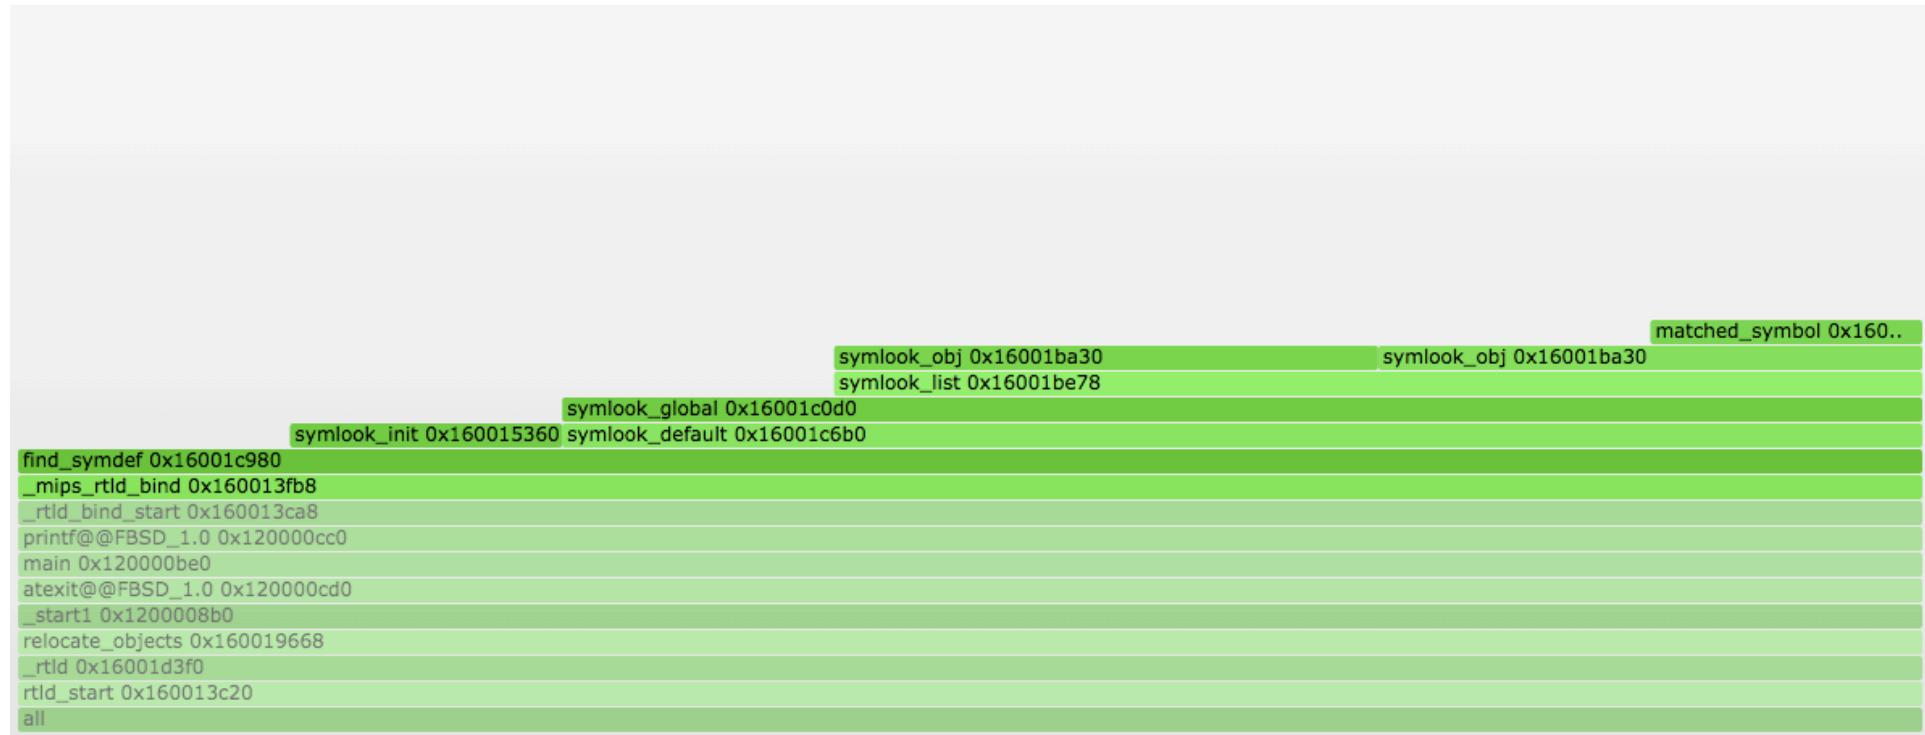
# \_start()



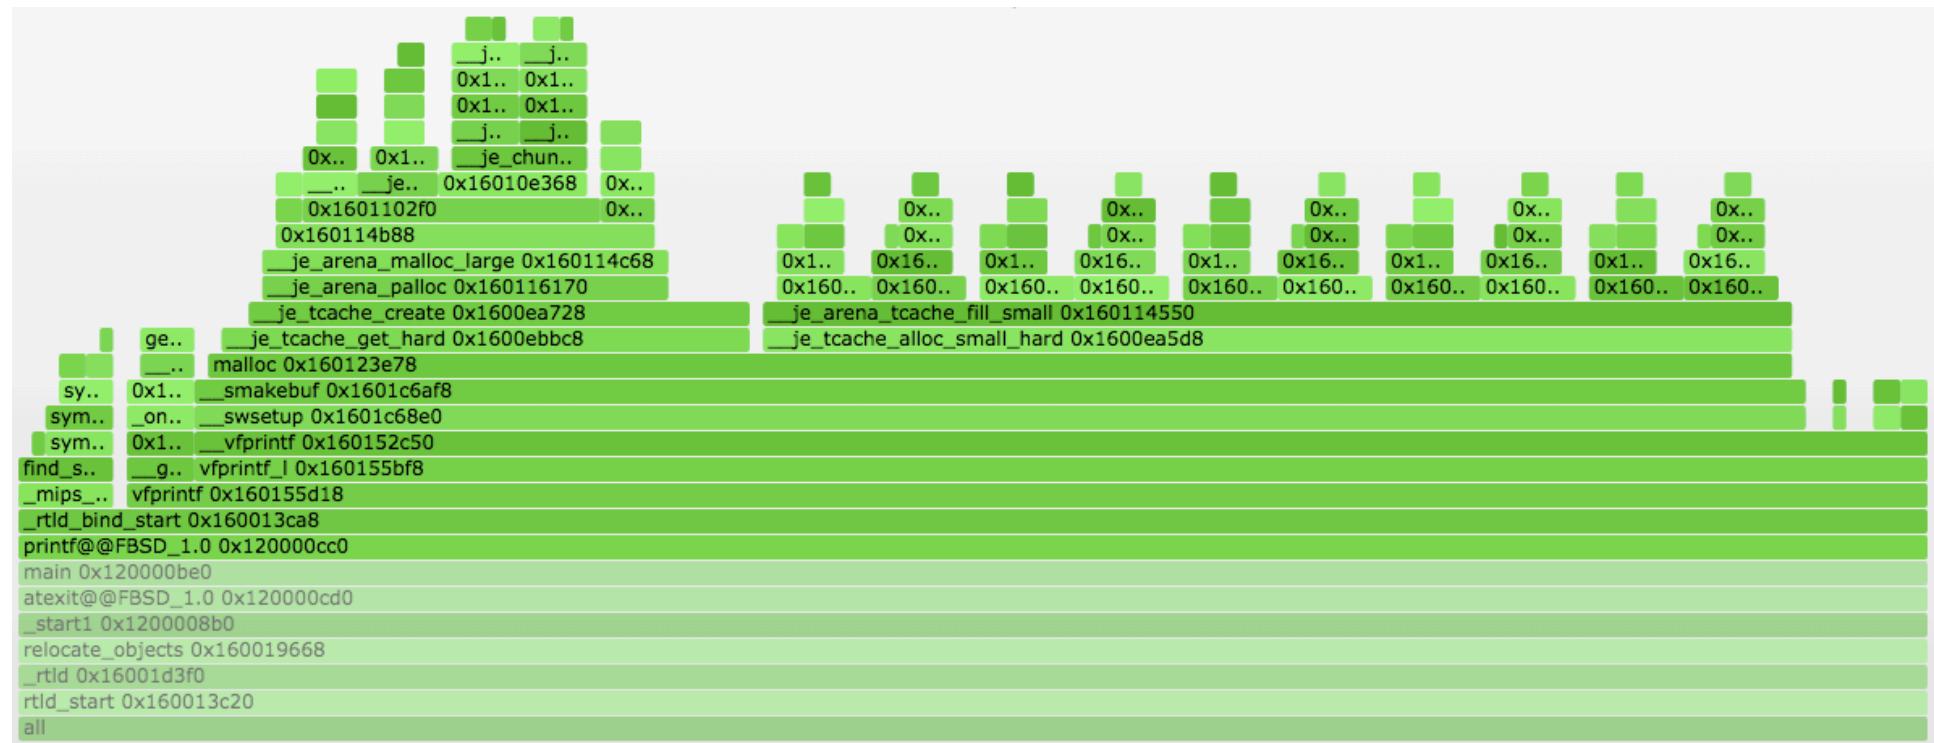
# printf()



# \_mips\_rtld\_bind()



# 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?
- [brooks.davis@sri.com](mailto:brooks.davis@sri.com)

