

# Dynamic Context-Based Code Elimination

Florian Rommel Daniel Lohmann

Leibniz Universität Hannover

2021-09-22


*Remove unnecessary code from binaries*

Why?

- Smaller binaries
- Reduced attack surface

Well researched.

Can we go further?



### Executions


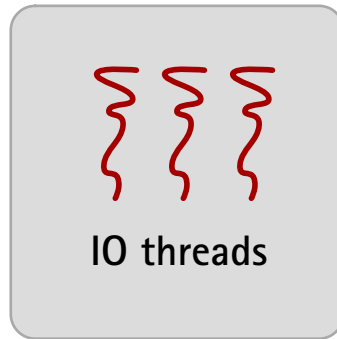



Fig. 1. Use of MUSL libc [16] functions by vsftpd [15].

A. Ziegler et al. 2019  
ACM Transactions on Embedded  
Computing Systems

# Context-Based Code Elimination

*Example: Redis (In-Memory Database)*



...

Communicate with  
Connected clients

**~82 % of the functions are  
not needed in IO threads  
(excl. libraries)**



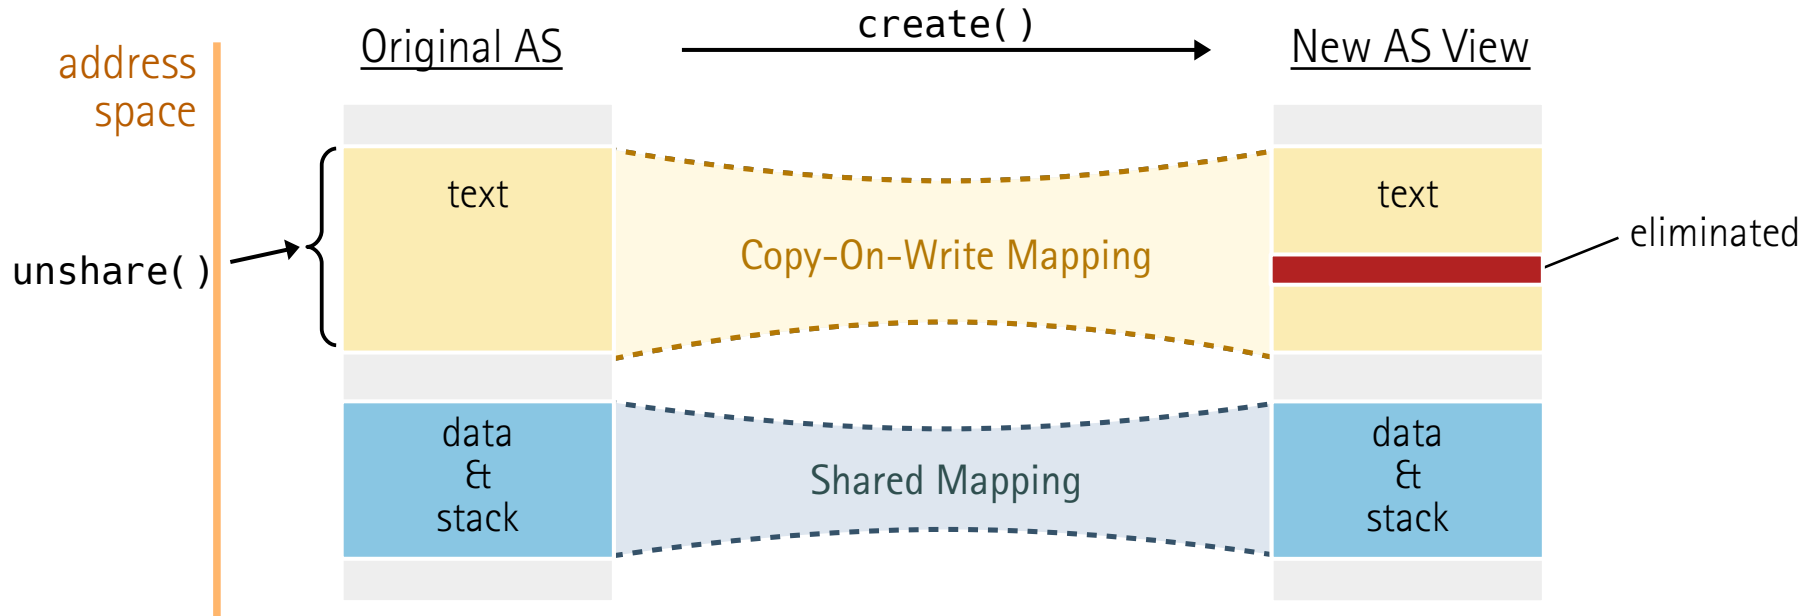
**Idea:**  
Elimination of these functions  
But only in IO threads

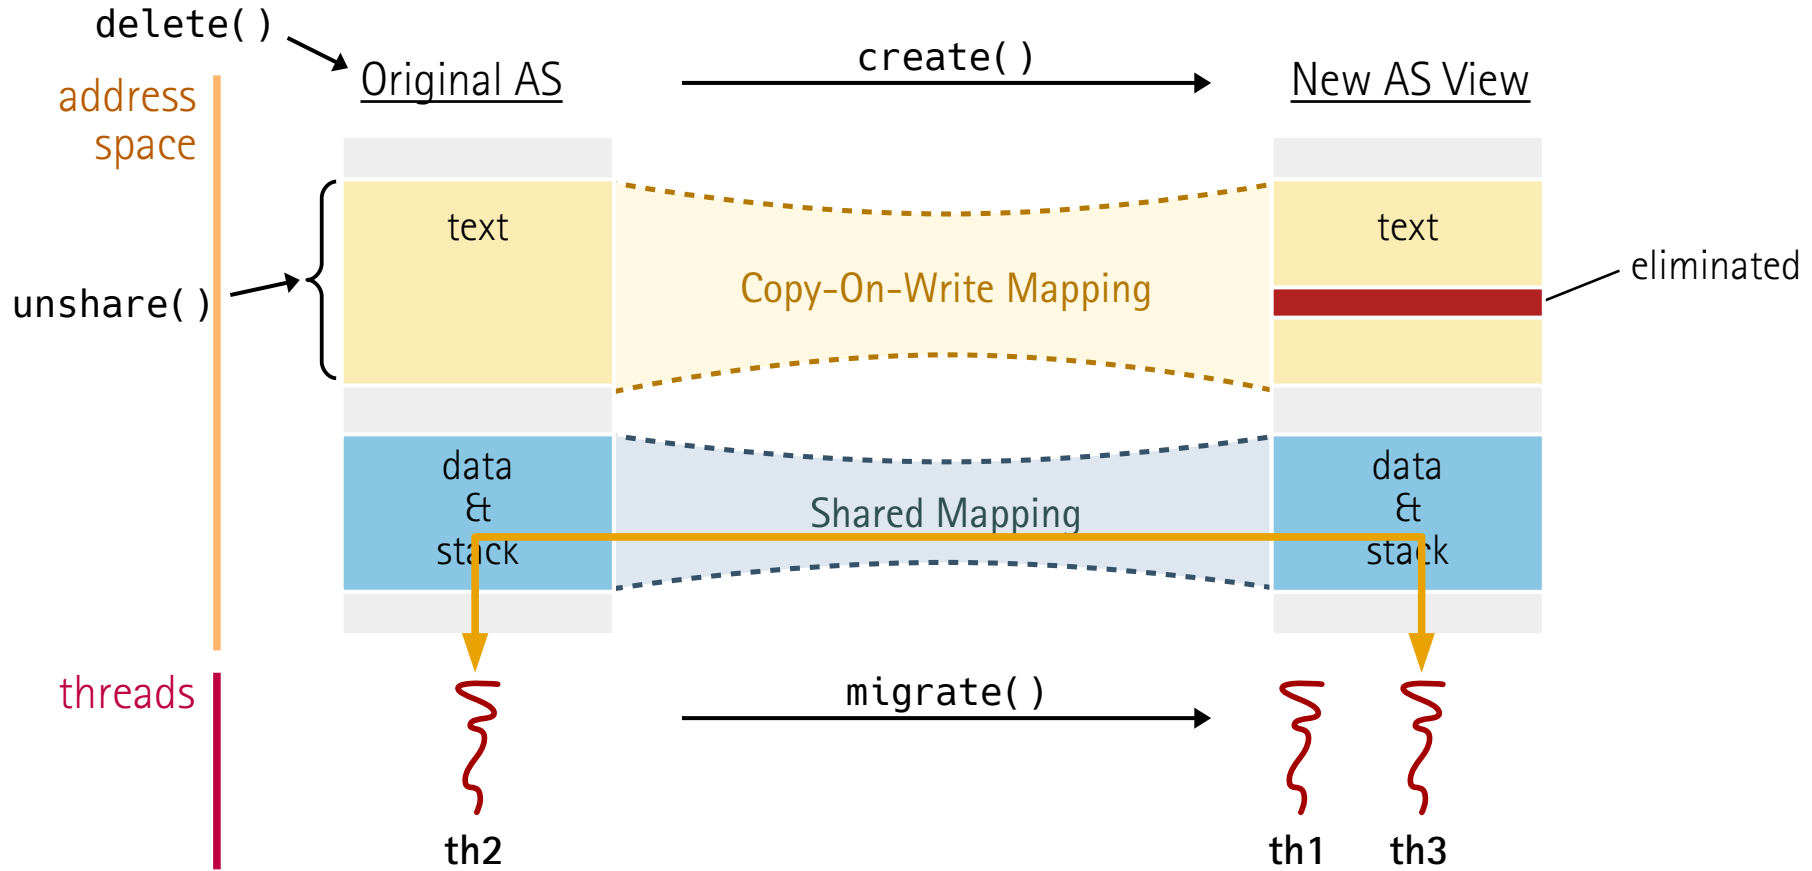
How can we eliminate code on the basis of threads?

## Dynamic Context-based Code Elimination via Address-Space Views

### Address-space views:

- Synchronized clones of the process's address space that differ can differ areas
- Threads can move between Views
- Implemented in the Linux Kernel






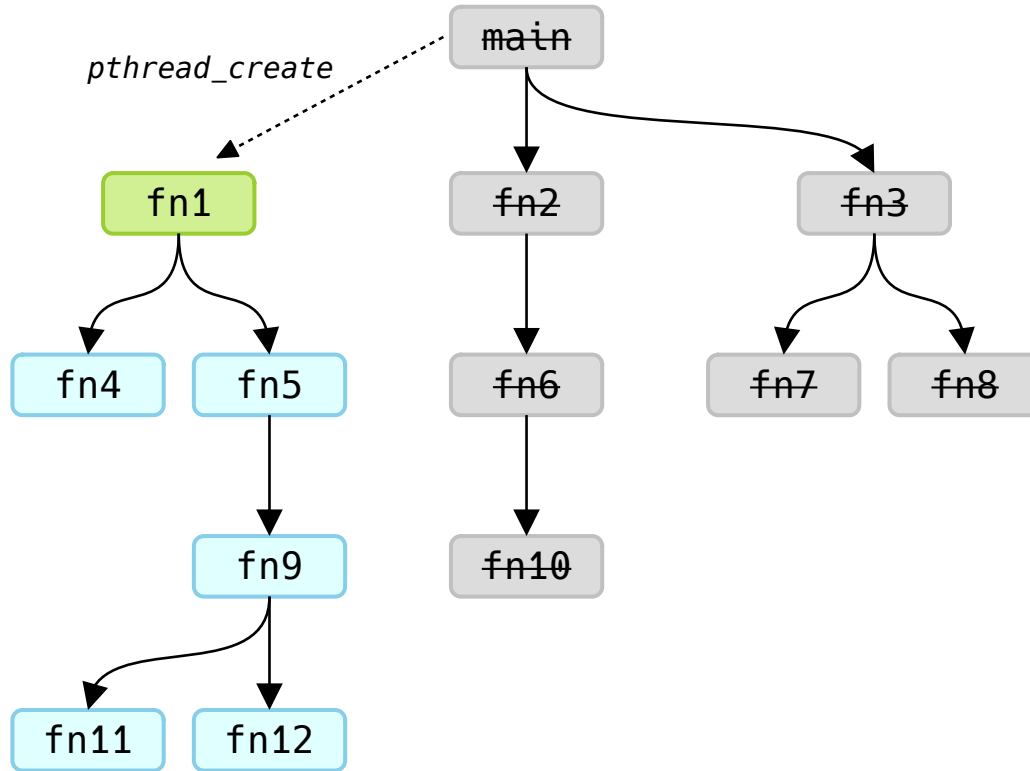
- Compiler Plugin (GCC)
  - Captures static call-graph information
  - Embeds information into the binary object (metadata section)
- Runtime Library
  - Consolidates metadata
  - Allows elimination of unused functions → replaces code with Invalid Opcodes

API:

```
void cte_init(void);  
void cte_eliminate(void *keep[], long keepc, void *nokeep[], long nokeepc);  
void cte_eliminate_self(void);
```





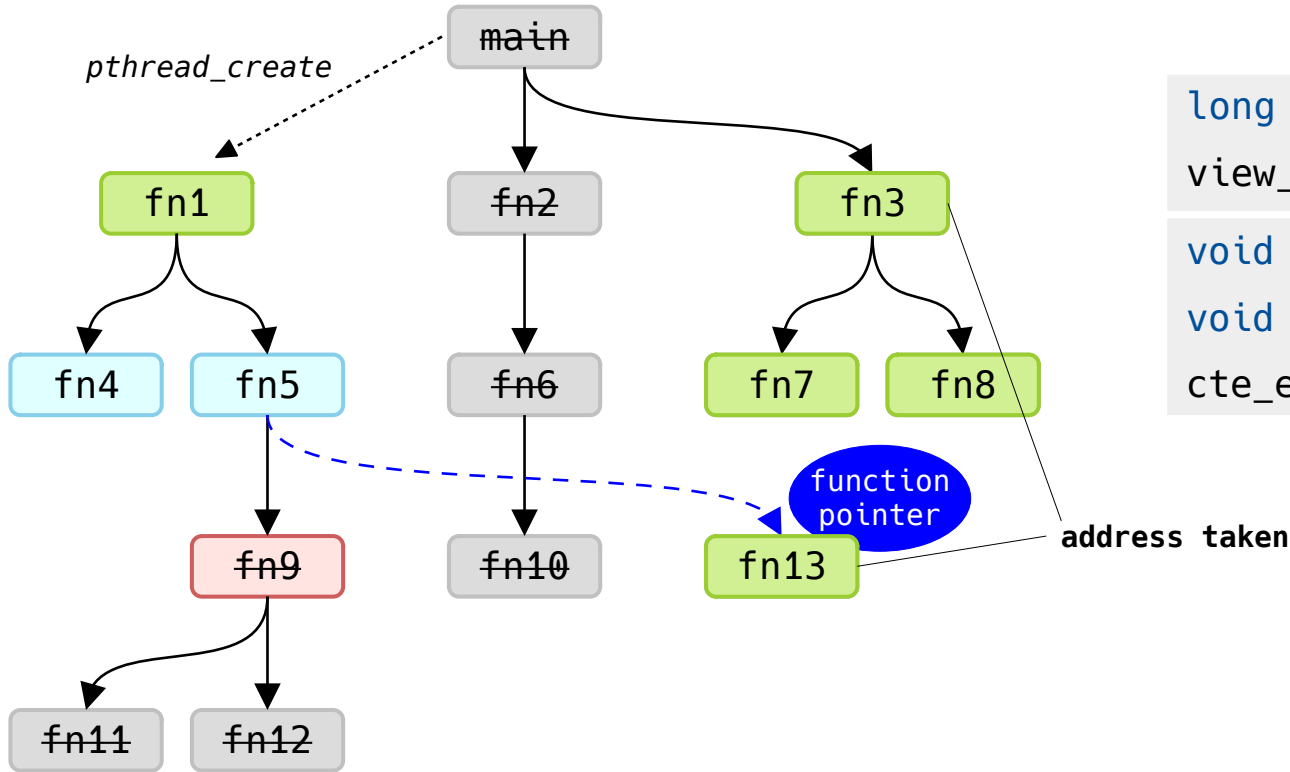


```
long view = view_create();
```

```
view_migrate(view);
```

```
void *keep[] = { fn1 };
```

```
cte_eliminate(keep, 1, NULL, 0);
```



```
long view = view_create();
view_migrate(view);
```

```
void *keep[] = { fn1 };
void *nokeep[] = { fn9 };
cte_eliminate(keep, 1, nokeep, 1);
```

→ Eliminate unnecessary functions for IO threads

- Functions eliminated
  - while preserving all "address-taken" functions:  
Removed **1738** of **2717** functions (~72 % code size [bytes])
  - While preserved only hand-selected "address-taken" functions:  
Removed **2227** of **2717** functions (~82 % code size [bytes])
- No measurable performance impact

- Separate processes / fork
  - Address spaces diverge
  - No more thread-like communication between contexts
  - Switching between contexts is not possible
- Intel Protection Keys
  - Available since Skylake in server CPUs
  - 16 protection domains per process
  - Restricted to page granularity

- Improve call-graph analysis
- Context-based elimination for data
- More areas of application (browser, web services)

## Dynamic Context-based Code Elimination

- Goal:  
Dynamically eliminate unreachable code on the basis of user-defined contexts.
- Approach:  
Use address space views to give each context its own view of the text segment.

*Thank you for your attention.*