

# Paper Reading Group

## Ksplice: Automatic Rebootless Kernel Updates

Eurosys 2009 Paper  
by

Jeff Arnold and M. Frans Kaashoek

# Use Case

- Reboots are risky
  - Security updates are postponed
- Security patches often touch only few lines of code
- How about patching this code *while it runs*?

# Current Approaches

- require specially written software
- constrain mechanisms software can use
- can only update kernels / software that had this feature in mind to begin with
- are constrained to specific programming languages
- Need laborious adaptation work to make patches runtime applicable.

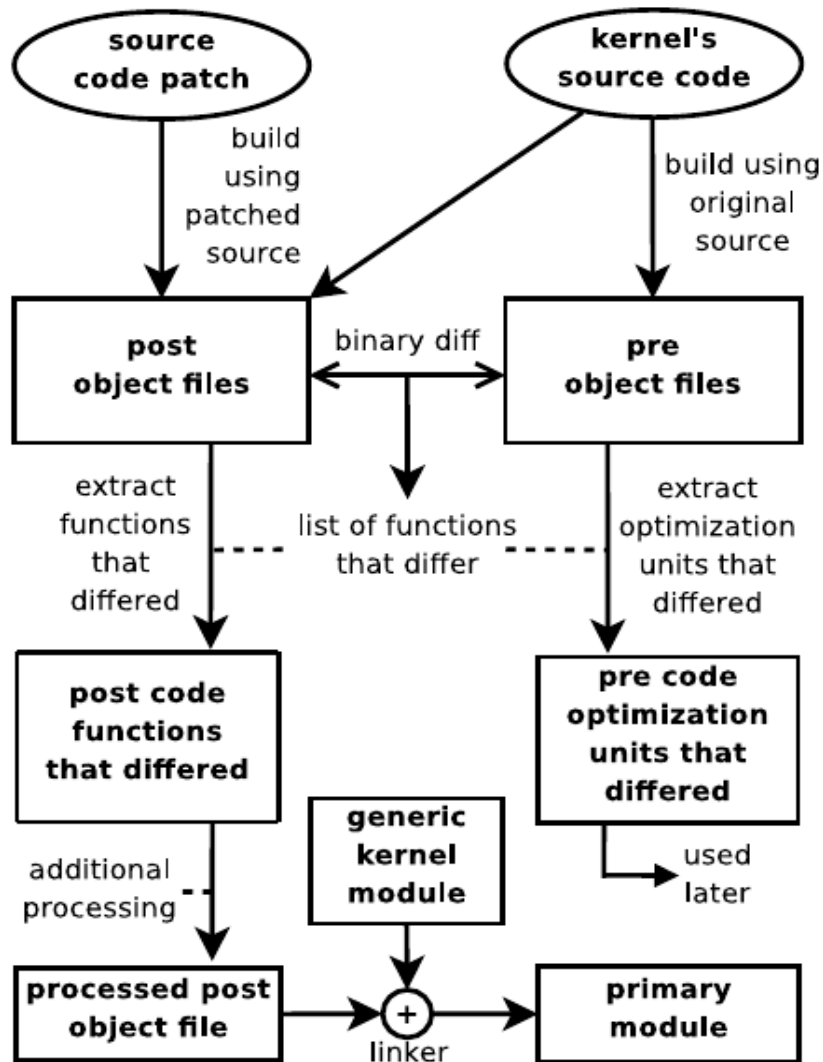
# How About Something New

- Binary diffs and patching of the changed code paths only!
  - Does not require changes in to original binary
  - Is not constrained on specific programming languages

# Challenges

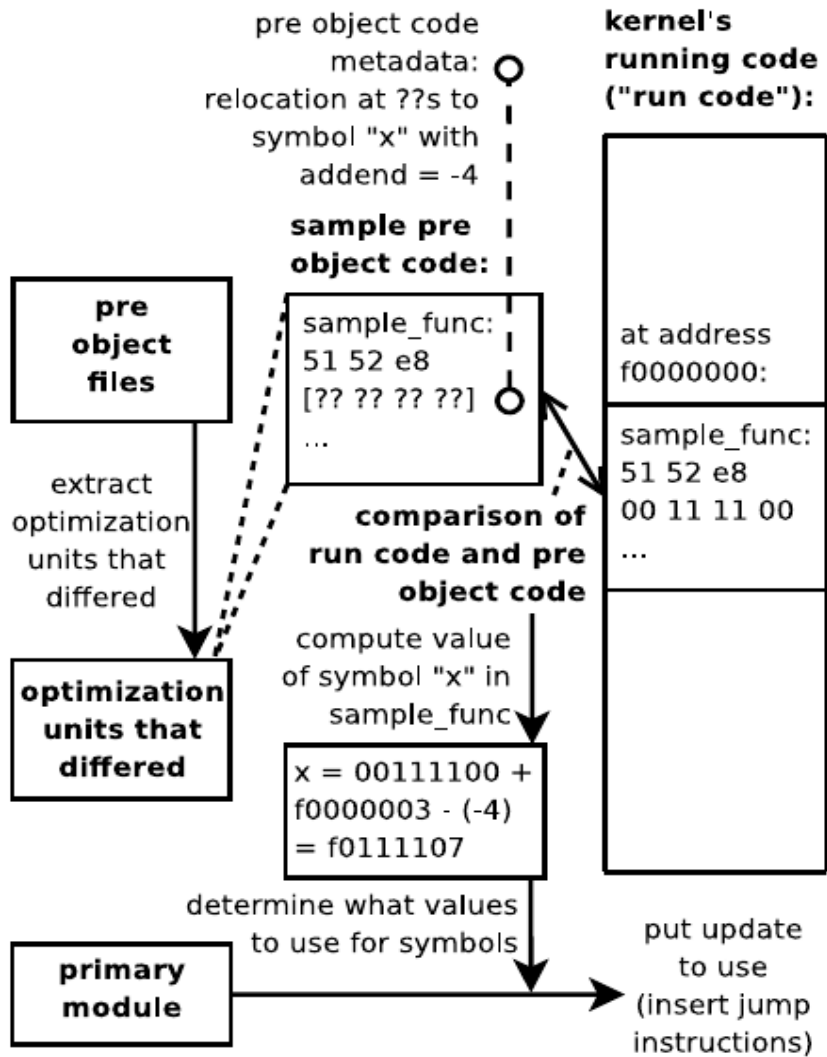
- resolve symbols in replacement code
- find changes resulting from patch in replacement object code
- avoid obfuscation by compiler optimization
- avoid inter code jumps making function isolation difficult
- ensure safety of code replacement (avoid crashes)
  - avoid missing inlined code leading to an inconsistent binary

# pre-post differencing



- Compile both version (patched and unpatched)
- each function within its own data and function section in the binary
- compare sections to find changed functions
- but what about symbol resolution?

# run-pre matching



- Use kallsyms symbol table
- problem for ambiguous symbols
- need to know no-ops (used as function alignment padding, does not imply difference!)
- must know instructions with relative addressing so that different rel. addresses pointing to the same location are not identified as differences.

# Ksplice Components

- Ksplice core kernel module (performs run-pre matching)
- pre/post object code generator in userspace
- helper module for loading pre-code (for comparison)
- primary module (loads new functions)

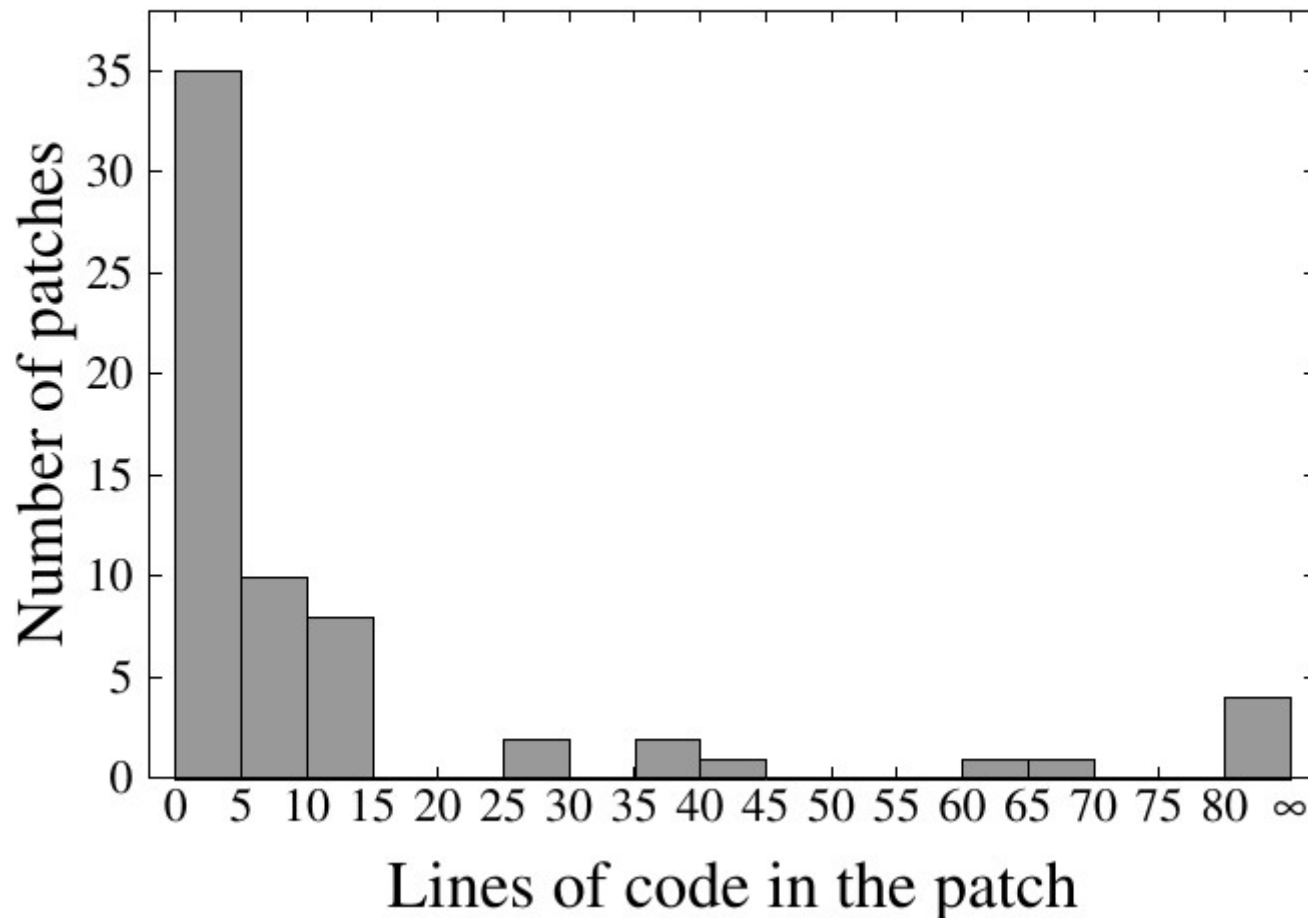


# Quiescent State

- Functions can only be replaced when:
  - they are on no threads kernel stack
  - no threads fp is in the function
- Retry on failure, give up after threshold
- stop\_machine takes about 0.7 ms to execute

# Evaluation

Figure 3: Number of patches by patch length



# Conclusion

- Found a way to patch arbitrary software during runtime
- With few limitations
- independent of programming language
- Safe (checks performed, does abort if assumptions not met)
- Can even cope with changes to data structures
- Eliminates reboots for kernel security updates

# Discussion Points

- How about jumps into functions? How are they covered?
- How do they compare *pre* and *run* code?
- Evaluation: 0.7 ms for how many threads?