

HCSS 2012



Static Previrtualization¹

N. Shankar

Computer Science Laboratory
SRI International
Menlo Park, CA

May 9, 2012

¹Funded by the Maryland Procurement Office Contract No. R21.002.10 and NASA Cooperative Agreement NNX08AY53A. Joint work with Drew Dean, Bruno Dutertre, Ashish Gehani, Gregory Malecha, Ian Mason, Sam Owre, Hassen Saïdi.

The View from the Top

- The modern hardware/software/communication stack is an unprecedented success.
- Complex applications are now easier to write.
- But this simplicity comes at a cost
 - **MiniBlog** – “Simple” PHP blogging application
 - 683 lines of PHP code
 - Depends on PHP & MySQL
 - **PHP** – Programming language interpreter
 - 625,000 lines of C
 - Depends on LibC
 - **LibC** – C standard runtime library
 - 650,000 lines of C
 - Depends on Linux kernel
 - **Linux Kernel** – Operating System
 - 15 million lines of code!



The Problem

- Software stack growing in size
 - More functionality
 - More hardware supported
 - Diverse application needs
 - Complex and numerous device drivers
 - Less robust to change
- Adverse security implications
 - More code to analyze
 - Increasingly complex interactions
 - Same time-to-market
 - - Reduced relative code coverage
 - Greater diversity of exploits
 - Wider attack surface
 - Larger consequences

Static Previrtualization: A Principled Approach

- Shrink the software by specializing it to a specific deployment
- Target full software stack and address all code growth.
- Specialize software stack to
 - Specific applications
 - Specific hardware/deployment configuration
 - Kernel modules trimmed
 - Device drivers trimmed
 - Libraries trimmed
- Applicable to
 - Specialized application servers
 - Custom libraries



Comparison




Figure 1: Single OS systems, Virtualized Systems, and Previrtualized Systems

- A practical application of Kleene's s-m-n theorem for optimizing code by specializing it to specific arguments

- ```
int atoi_strcat(char *str1, char *str2){
 return atoi(strcat(str1, str2));
}
...
char *pad = "000";
i = atoi_strcat(mystr, pad);
```

- The `atoi_strcat` invocation can be replaced by  
`atoi_strcat_1000(argv[1])`

where

```
int atoi_strcat_1000(char *str1){
 return 1000 * atoi(str1);
}
```

- Fixes specific parameters
  - Interpreter + program → object code
  - Specializer + interpreter → compiler
  - Specializer + specializer → compiler-compiler
- Partial evaluation of common system calls
  - Improved performance (*McNamee, 2001*)
- Customized runtimes via partial evaluation
  - Application developer / user knows needs
  - Operating system designer does not (*Howell, 1998*)

- Application system calls partially evaluated
- Resulting code is
  - Compact
  - Efficient
  - Portable
  - Isolatable
  - Less vulnerable to attack
  - More amenable to static analysis
- LiveCD or virtual appliance, but without redundant software

# Project Background

- Tried package minimization: Reduces footprint but granularity is too coarse
- Evaluated existing partial evaluation technology:
  - CMix-II (Henning Makholm)
  - Tempo (Charles Conzel)
- The technology is impressive but not current with language/architecture changes (C99 or 64-bit machines)
- More recently, we switched to LLVM
- We have developed a previrtualization/monitoring toolchain called **Occam** based on LLVM
- It has been applied to web servers and PHP/MiniBlog



# Reducing Functionality with Occam

- **Program:** thttpd
- **Size:** 11,322 lines
- **Problems**
  - Uses potentially dangerous functions like listen , connect, etc.
  - Reads configuration data from the command line.
- **Solutions**
  - Limit the ways that dangerous functions can be called.
  - Compile configuration data into the program.



# Partial Evaluation of LLVM Bitcode

- Low Level Virtual Machine (LLVM) is a typed, machine-independent intermediate format (IF, called bitcode) due to Adve and Lattner.
- The IF uses static single assignment on typed registers.
- Many languages have front-ends to generate LLVM, e.g., Ada, C, C++, Objective C, Haskell.
- Analyzers and code-generators can be driven from LLVM.
- `glibc` could not be directly converted to bitcode, but `uClibc` was adequate for this purpose.
- Simple forms of partial evaluation on LLVM have been explored:
  - Fujita uses cloning and LLVM's optimizer
  - Smowton and Hand inline file data



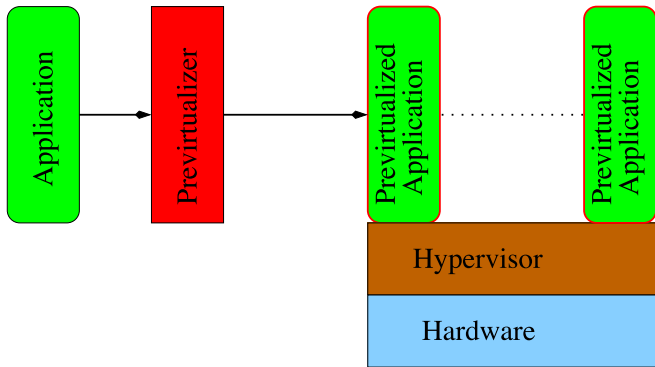
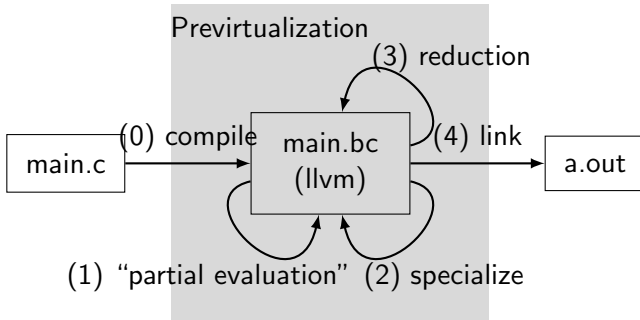


Figure 2: Previrtualizing Applications

- ➊ Reduce the “functionality” of a system
  - `nweb` is a simple webserver.
  - It doesn't need to be able to listen on arbitrary ports.
  - Make configuration options static.
- ➋ Overcome static analysis
  - Miniblog should never send email, so that functionality should not be in the system.
  - We need to cut it out, since mail is in the PHP standard library (compiled into the interpreter!).
- ➌ Monitor systems and enforce dynamic policies
  - Log function calls as the program runs.
  - Check security properties.




# Module Previrtualization Overview



# Partial Evaluation

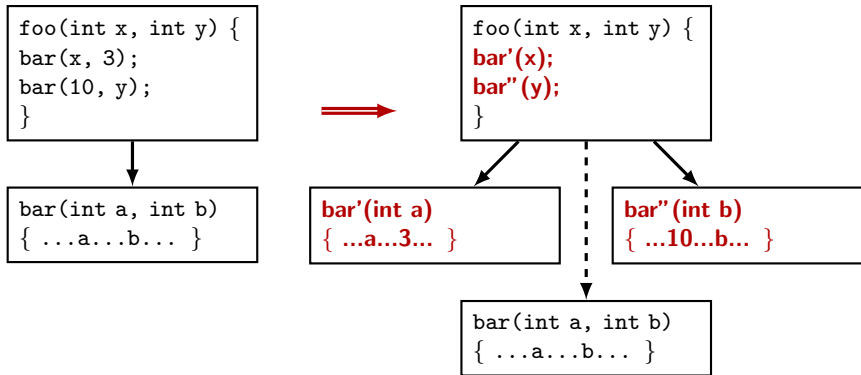
- Simplify the program as much as possible, want to expose constants.



- Use LLVM's `-O3`.

# Specialization

- Specialize functions when they take constant arguments.

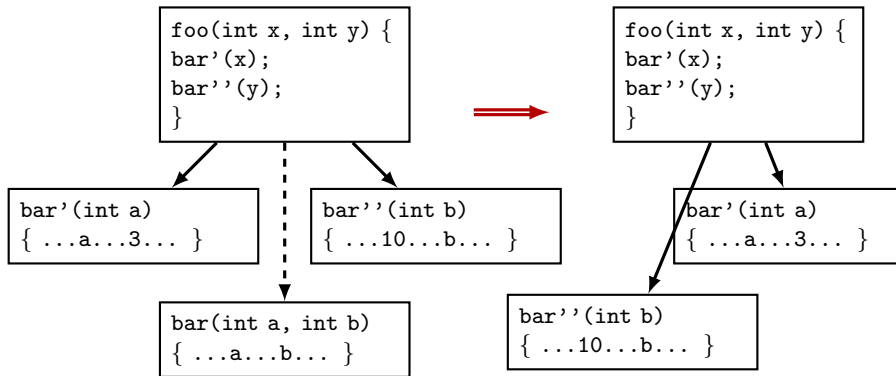


- Clone functions and inline constants in a custom LLVM pass.



# Reduction

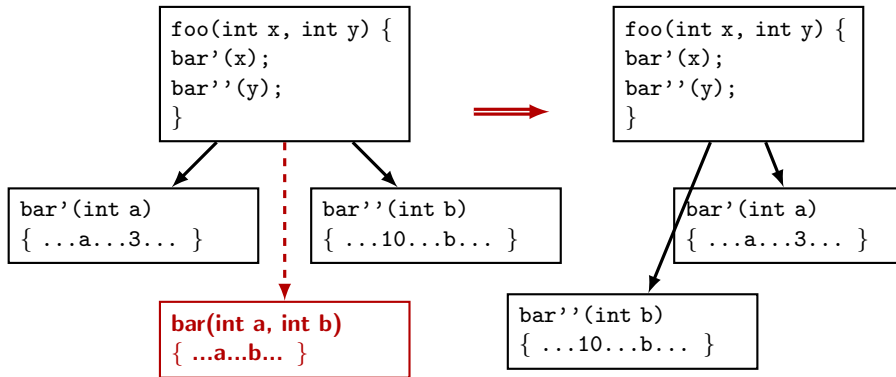
- Eliminate unused code.



- LLVM dead-code/global elimination pass.

# Reduction

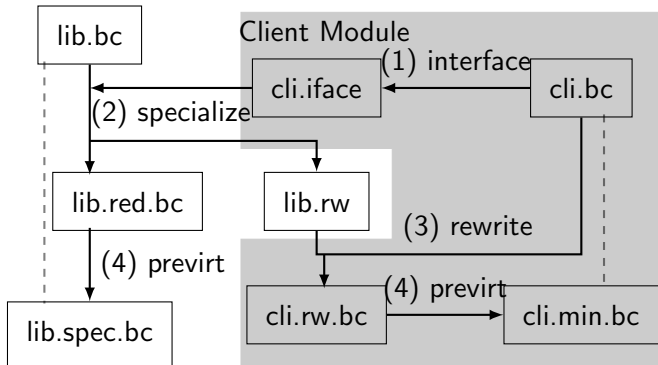
- Eliminate unused code.



- LLVM dead-code/global elimination pass.

- 1 Specialization of command line parameters
- 2 Partial evaluation by optimization
- 3 Aggressive specialization of dangerous functions
- 4 Dead-code elimination
- 5 Goto 2!

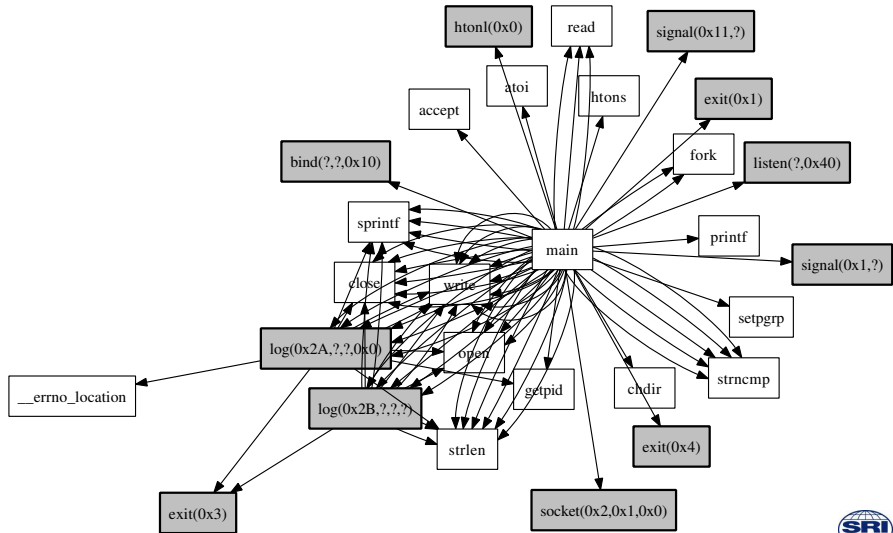
# Cross-module Previrtualization



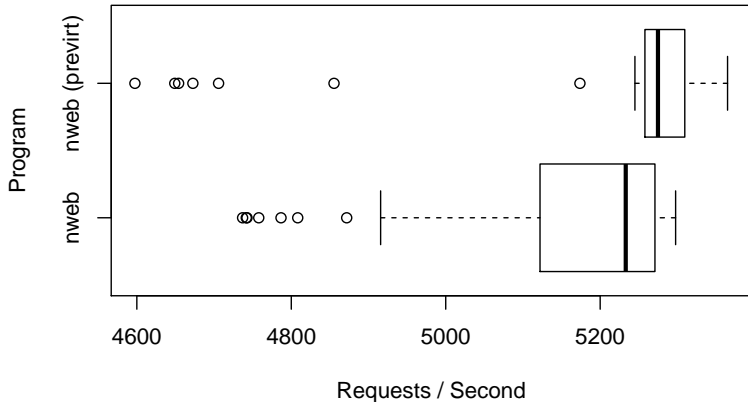




# Callgraphs: Before & After



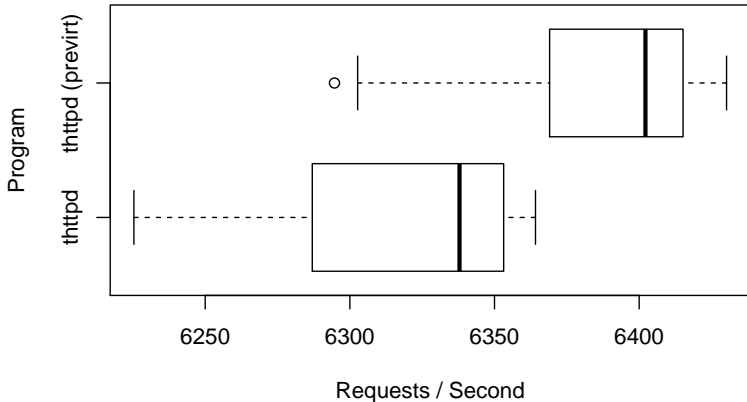
## Previrtualized nweb Performance



nweb performance before previrtualization.

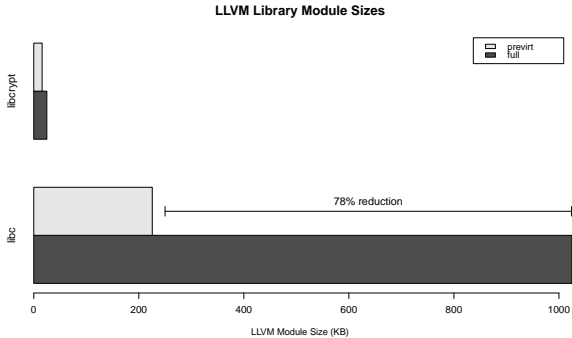
# Performance of Previrtualized Software (thttpd)

## Previrtualized thttpd Performance

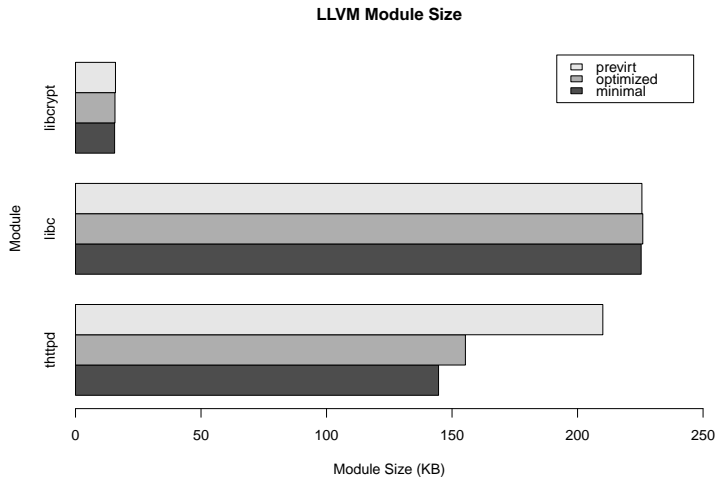


thttpd performance after previrtualization.

# Size of Previrtualized Software



# Size of Previrtualized Software



# Limits of Compile-time Analysis

- We can't automatically determine reachable code in all cases
  - Cross-language or cross-binary calls
  - Indirect function calls & static approximations (PHP problem)
  - Function pointers and binary compatibility
- **PHP built-in functions are implemented as a large function table.**
  - Static analysis has to say that all of these are reachable.
  - Lets the bad in with the good.

## PHP Snippet

```
const zend_function_entry basic_functions [] =
{ ... PHP_FE(system, arginfo_system), ... };
PHP_FUNCTION(system)
{ php_exec_ex (INTERNAL_FUNCTION_PARAM_PASSTHRU, 1); }
```



- 1 “Statically analyze” the PHP code and determine the functions that it will call.
  - For relatively static applications this can be done with a `grep`-like static analysis.
  - Miniblog requires about 46 PHP functions out of the 1028 functions that a minimal PHP install would have.
- 2 Implement a transformation that will replace these unused functions with a simple `exit(1)`.
  - Previrtualize the result to remove all the unnecessary code.

- We can specify subs the same way that we refer to specializations.

## Remove `system` Function

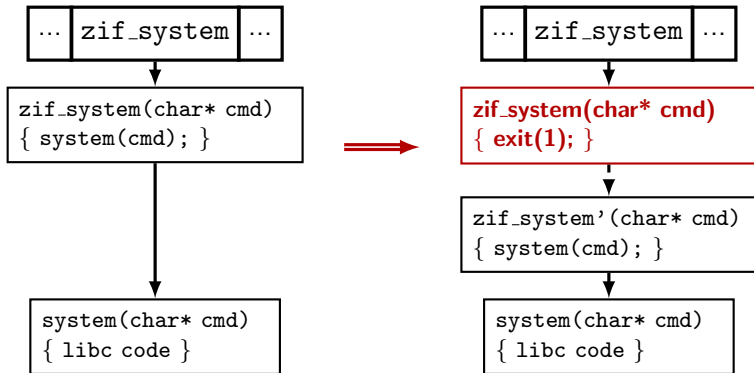
```
zif_system (?) => fail
```

- `fail` is a keyword meaning call `exit(1)`.
- Question marks specify wildcard arguments; here we stub *all* calls to `zif_system`.
  - Also support integer constants, so we can reject some calls but not others.



# Rewriting Code

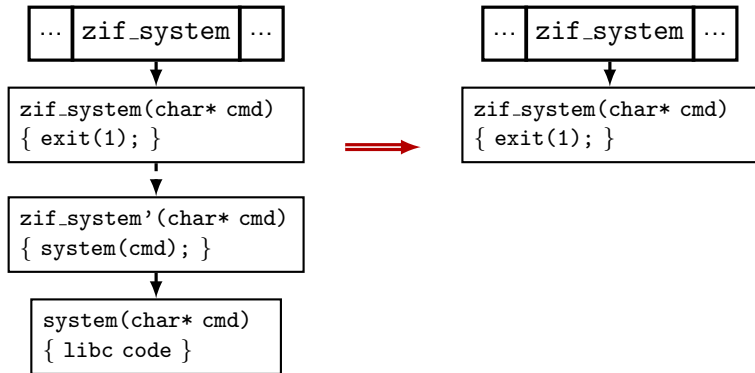
- Small transformation pass to replace function bodies.



- Implemented as a custom LLVM transformation pass.

# Reusing the Previrtualization Hammer

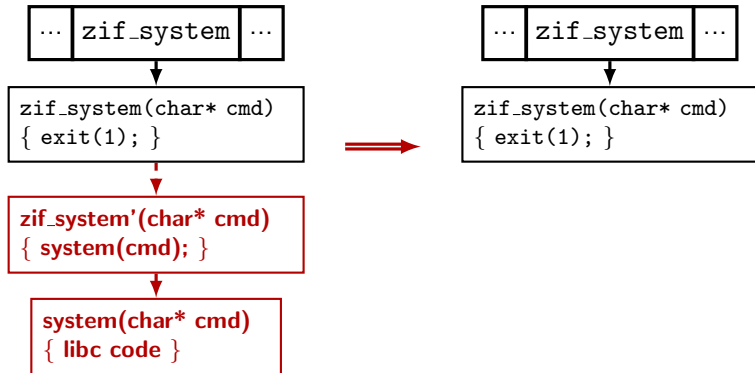
- Remove dead code using previrtualization.



- Reduce to an already solved problem!

# Reusing the Previrtualization Hammer

- Remove dead code using previrtualization.



- Reduce to an already solved problem!

- Remove dangerous PHP functions:
  - system
  - mail
  - etc.
- Previrtualization removes unused dependencies.

- Extend the enforcement mechanism.
- An implementation of *Aspect-Oriented Programming* for LLVM.
  - Monitor when execution enters/exits a function.
  - Support access to function arguments and return values.
  - Support conditional monitoring.
    - Allowing exit (1) on certain parameters.
  - Monitored binaries can be run *without* monitors.

# Does Previrtualization Increase Security?

- Many attacks exploit buffer overflow and format string vulnerabilities to inject code or invoke existing functionality.
- Protections like *StackGuard* have a runtime cost.
- Attacks like `return-to-libc` and return-oriented programming (ROP) can be defeated by address space layout randomization (ASLR) on 64-bit machines.
- Even through previrtualization introduces more potential attack sequences for ROP, this is dwarfed by the entropy introduced by 64 bits.
- The bigger gain is that many potential vulnerabilities are pruned by previrtualization.
- Specifically, vulnerabilities in start-up code are unavailable in the previrtualized application.



- Massalin's *Synthesis* kernel used partial evaluation for a form of run-time code generation for efficiency.
- McNamee, *et al.* optimized frequently used system calls with PE.
- Fujita exploited the LLVM optimizer for intra-module partial evaluation
- Snowton and Hand used this technique for inlining file data.
- Turnkey Linux distributes coarsely pruned appliances for several applications, including a JeOS (Just enough Operating System) stack.

- Improve the previrtualization toolchain
- Deeper partial evaluation
- Kernel previrtualization: Compiled FreeBSD with clang
- Inter-application previrtualization
- Adding security checks and monitoring during previrtualization
- VM previrtualization
- Other platforms: Android
- Previrtualization as a service.



- **Occam** is a tool for *previrtualization*
  - Program specialization to reduce *functionality*.
  - Partial evaluation through optimization.
  - Works well for generic platforms, e.g., languages with large libraries like PHP.
- Monitoring program execution
  - Monitors can be placed around functions and modify both inputs and outputs.
  - Monitors can be arbitrary C++ code and can maintain state between calls.
  - Aspect-oriented programming
- These techniques can be used for shrinking and wrapping the software stack for each deployment