

# Building Efficient Software Dynamic Translators

#### **Bruce Childers**

Jack Davidson

University of Pittsburgh Department of Computer Science Pittsburgh, PA 15260 childers@cs.pitt.edu University of Virginia Department of Computer Science Charlottesville, VA 22904 jwd@virginia.edu

1



#### Outline

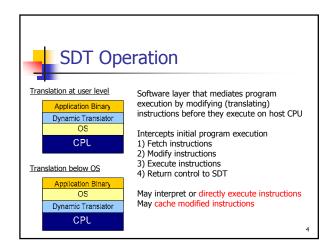
- Part I: Introduction to SDT
- Part II: The Strata SDT Framework
  - Translation virtual machine
  - Indirect branch handling
  - Performance
- Part III: Code security applications

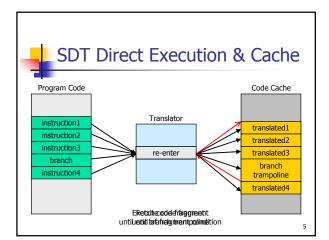
2



#### What is SDT?

- Any software that intercepts, controls, or modifies a program as it runs
- Subsumes:
  - Dynamic optimization / compilation
  - Dynamic binary translation
  - Dynamic instrumentation (e.g., profiling)
  - Host virtualization
  - Debugging





# Why Use SDT? Improve program performance Adapt program to its execution environment Overcome economic barriers Allow one architecture's binaries to run on another Application specific ISA improvements Code decompression, encryption Resource management Power, memory footprint, resource protection Software engineering and dependability Performance monitoring, fault isolation, security



#### Challenges

- Significant impediments to SDT uses
- 1. Tightly tied to host platform
- 2. It's another layer: Run-time overhead
- 3. User debugging translated programs



# Close Coupling to ISA & OS

- Steep learning curve
  - Low-level understanding of platform and translator
- Decoding & translating instructions
   E.g., delay slots? variable length ISA?
- Switching context
- E.g., how, what & where to save/restore?
- Signal handling
- E.g., ensure translator maintains control?
- Multi-threading
   E.g., sharing data structures? locking?



#### Run-time Overhead

- Translator overhead
  - Fetch, decode, & translate loop
    - Possibly, re-translate already seen code
  - Context save/restore
  - Platform interactions with translated code
- Translated code overhead
  - Invoking translator after execution
  - Handling control transfers
  - Program instrumentation needed by translator



# **Debugging Translated Code**

- Typically, forced to debug at instruction level
- Static debug mappings become inconsistent
  - Dynamically generated, optimized code
- Multiple and changing statement instances
- Translated vs. untranslated code
  - Insert/remove breakpoints
- Extra code (e.g., trampolines and instrumentation)
- AADEBUG '05
- Debugging the dynamic translator itself!

10



# **Example Applications & SDTs**

- Many existing and beneficial SDT uses
- Many systems typically, re-implement translator for a new use of the technology
  - Architecture study (Shade, Embra)
  - Host virtualization (SimOS, VMware, Flex86)
  - Binary translation (DAISY, FX!32, Transmeta CMS, Co-designed VMs, Walkabout)
  - Code optimization (Dynamo, LLVM, ADORE)
  - Dynamic instrumentation (Pin, INSOP)
  - Code security (Dynamo/RIO)

11



#### Shade, Embra

- Architecture study: Examine trade-offs
- Very fast & focused: Model only what is needed for some trade-off
- Translate application to simulate/model architecture features
  - Avoid interpretation for simulation
  - Translate and instrument to have simulated effect
  - A form of dynamically compiled simulation



# Shade, Embra

- Shade [Cmelik & Keppel, SIGMETRICS'93]
   User-written instrumentation

  - Translates, caches & chains at basic block level

  - Perform callbacks at instrumented points
     E.g., cache simulation instruments loads and stores
     Slowdown: 10-228x on SPEC2000, Sun Blade 100
- Embra [Witchel & Rosenblum, SIGMETRICS'96]
  - Same idea as Shade to model MIPS R4000

  - Accurate & fast memory hierarchy simulation
     Tightly stitch memory simulation code into application code
     Uses quick checks (lookups into large array) to avoid expensive searches for MMU and caches
  - Slowdown: 5-34x on SPEC2000 (Embra's techniques on SPARC)



#### DAISY: PowerPC to VLIW

- Binary translation [Ebcioglu & Altman, ISCA'97]
- Translation approach
  - Operates on page worth of PowerPC instructions
  - Translation done one instruction at a time & immediately scheduled into VLIWs
  - Maps untranslated pages to translated pages
  - On a missing translated page, re-invoke translator
- Performance: Within 20% of static VLIW compiler
  - 4315 instr/translated vs. 100,000 instr/translated



#### Dynamo

- Binary dynamic optimization [Bala et al., PLDI 2000]
  - Machine specific performance w/o recompiling application
- Optimization of PA-RISC binaries
  - Interpret instruction sequences, lightweight profiling
  - When hot, form *instruction traces* (NET algorithm)
  - Optimize instruction traces: primarily redundancy removal
  - Cache & directly execute the instruction traces
  - Flush code cache when increase in trace creation
- Performance: about -2% to 23%, avg. 9%, over -O2



# Dynamo/RIO

- Reconfigurable translation [Bruening et al., CGO'03]
  - x86 framework for implementing uses of SDT
  - Library & API: intercept & control translation
    - Callbacks on specific translator events and operations
- Features
  - Code caching & branch handling (chaining)
  - Multiple structure representations of instructions
  - Traces & trace creation (based on Dynamo's NET)
- Performance: 0.93-1.8x, avg. 1.25x, on SPEC2000

16



# Dynamo/RIO

- Program Shepherding [Kiriansky et al., USENIX Security '02]
  - Monitor control flow to enforce security policies
  - Transparent & efficient with dynamic code translation
- Check trustworthiness of code during translation
- During run-time, verify program addresses
- Sandboxing to avoid getting around checks
- Protected RIO data structures with page permissions
- Performance: 1.0-1.7x on Linux

1



#### Pin

- Customized program analysis using dynamic instrumentation [Luk et al., PLDI'05]
  - ATOM-like program analysis routines written by users
- Pintools
  - Instrumentation code written in C/C++ using Pin's API
  - Dynamic translation to generate and instrument code
  - Reduce instrumentation overhead by inlining, register reallocation, liveness analysis and instruction scheduling
- Slowdown: 2.5x on integer and 1.4x on floating point



#### **Outline**

- Part I: Introduction to SDT
- Part II: The Strata SDT Framework
  - Translation virtual machine
  - Handling indirect branches
  - Performance
- Part III: Code security applications

19



# Why do we need an SDT infrastructure?

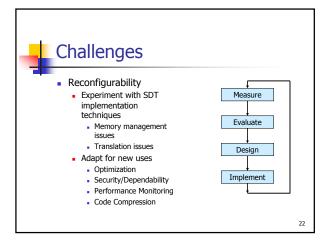
- Simple answer: make SDTs easier to develop
  - Allows experimentation with novel SDT systems
  - Accelerates research
  - Provides a model for other SDT systems
- How?
  - Factor out code needed across many SDTs
  - Allow for SDT reuse and composition
  - Provide support for multiple architectures to ease retargeting
  - Provide efficient translation mechanisms

20



#### Strata

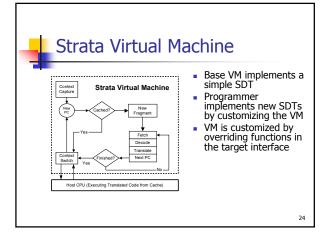
- Infrastructure designed for building SDTs
- Designed with extensibility in mind
  - Optimization
  - Code compression
  - Performance monitoring
  - Security
- Provides:
  - Platform independent common services
  - Target interface that abstracts target-specific support functions
  - Target-specific support functions

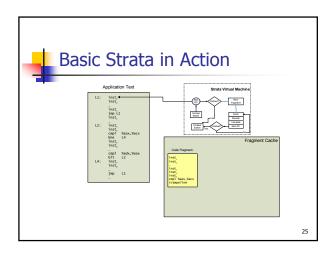


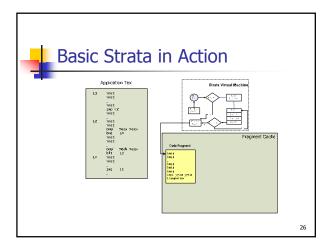


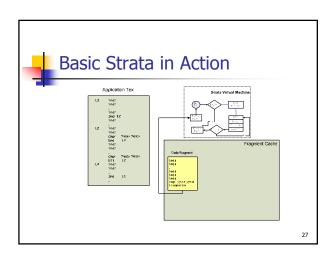
# Challenges

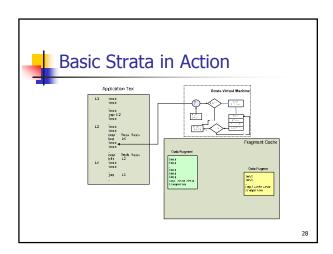
- Retargetability
  - Instruction-set architecture
    - Instruction decoding
    - Instruction semantics
    - Alignment
  - Operating system
    - Calling convention
    - Memory management and consistency
    - System calls
    - Signal handling
    - Thread semantics

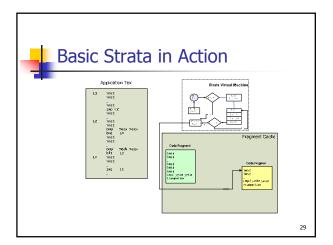


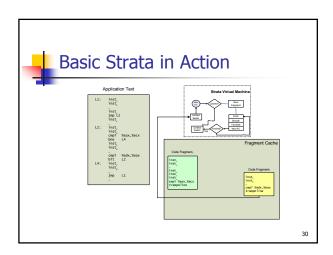


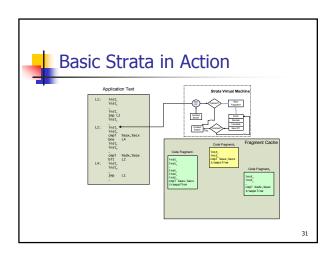


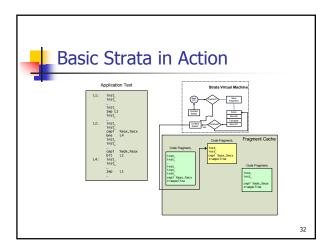


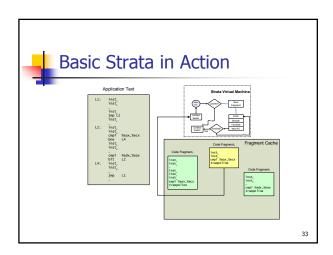


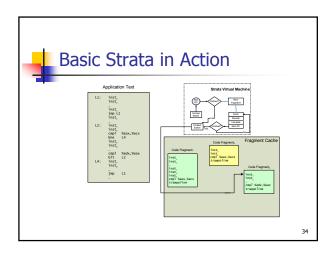


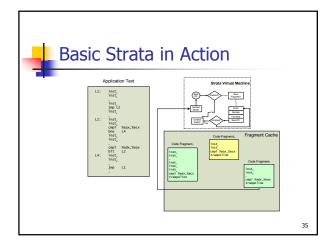


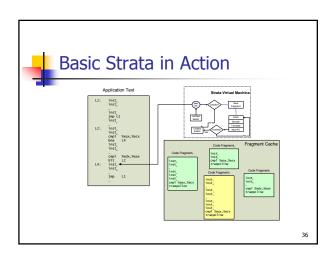


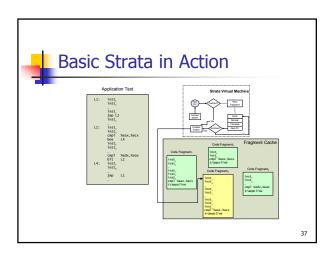


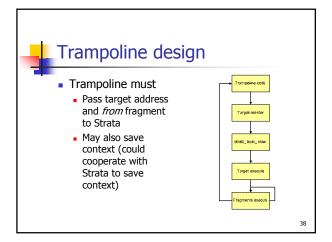


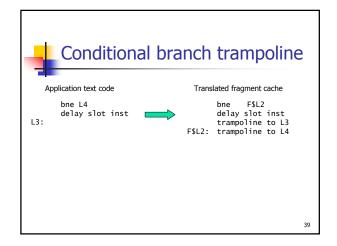


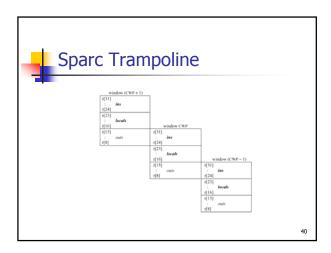


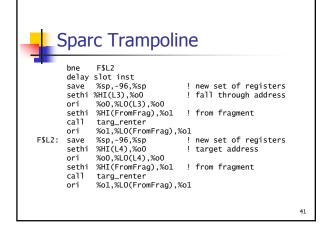


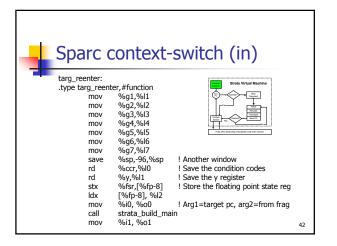


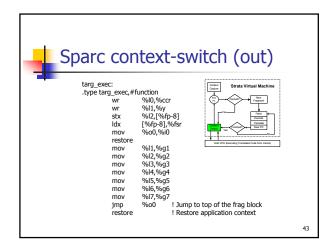


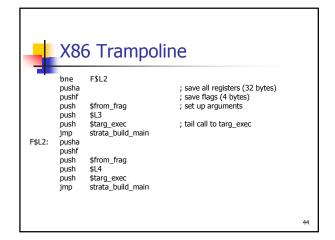


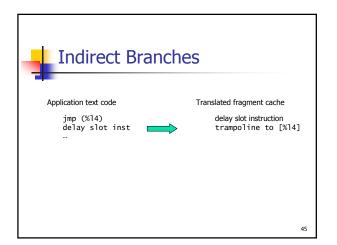


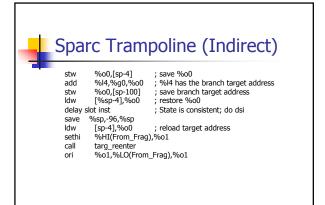










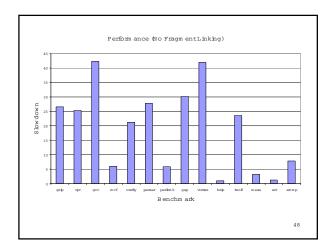


push targ\_address pusha pushf push \$from\_frag push push \$targ\_exec jmp strata\_build\_main ; must be pushed before sg ; save all registers (32 byte ; save flags (4 bytes) ; set up arguments ; tail call to targ\_exec ; tail call to targ\_exec

ine (Indirect)

; must be pushed before sp changes; save all registers (32 bytes); save flags (4 bytes); set up arguments

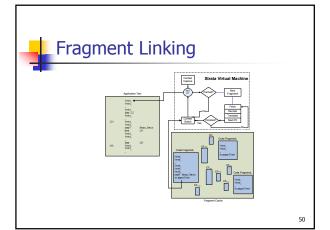
; tail call to targ\_exec

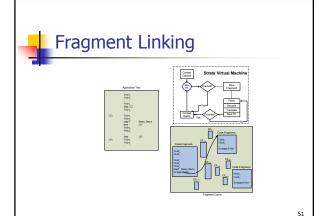


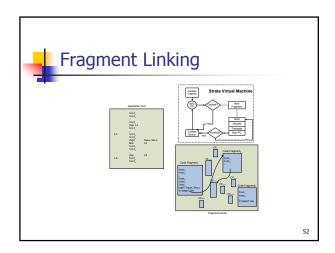


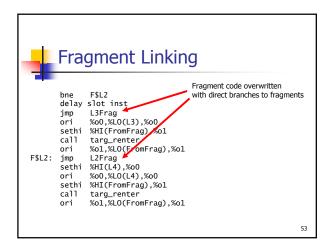
# Fragment Linking

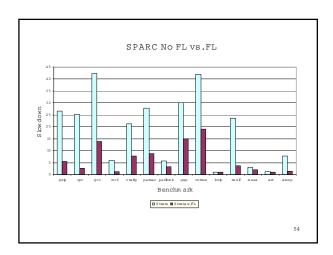
- If target of branch exists in fragment cache, do not context switch
- Patch branches in fragment cache to jump directly to the fragment













#### The Common Services

- Used by VM and target-specific functions
- Platform independent
- Provides:
  - Strata dynamic memory management (via arenas)
     Code cache management

  - Fragment linking
- Strata performance monitoring and reporting
- Easily extended to incorporate new services

<u>Application</u>	
Context Management Linker	
Memory Management  Cache Management  Target Interface	ine
Cache Management Strata Virtual CPO	Machine
Target Interface	2
Target Specific Functions	
Host CPU and OS	



# The Target Interface

- Interface used by VM to perform target specific functions
- Used to customize the VM for new software dynamic translators
- Each target (SPARC/Solaris, MIPS/IRIX, x86/Linux, etc.) must implement all of the target interface

Applicat	ion	
Context Management	Linker	l
Memory Management	Strata Virtual CPU	in a
Cache Management		Strata Virtu
Target Inte	erface_	# 41
Target Specific	Functions	
Host CPU a	and OS	



#### Strata

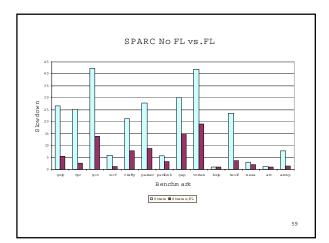
- 3.9K LOC (Platform/Machine independent)
- Platform/Machine Dependent
  - 2.3K LOC (Sparc Solaris)
  - 2.1K LOC (MIPS Irix)
  - 2.0K LOC (x86 Linux)



#### **Outline**

- Part I: Introduction to SDT
- Part II: The Strata SDT Framework
  - Translation virtual machine
  - Handling indirect branches
  - Performance
- Part III: Code security applications

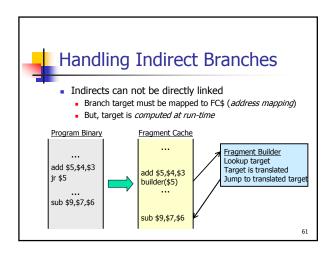
58

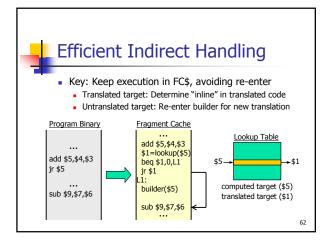


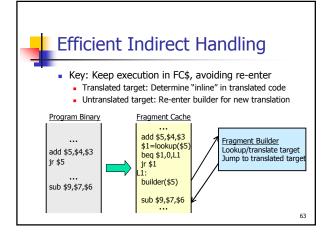


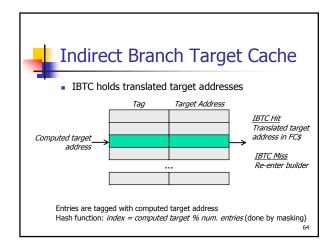
# Handling Indirect Branches

- Remaining majority of builder re-enters
- Sources of indirect branches
  - Switch statements
  - Returns
  - Function pointers
- Indirect branches treated in two parts
  - Actual indirects (e.g., jr \$r7)
  - Returns (indirects via the return address register)











# IBTC Design Trade-offs

- Size of IBTC: Fixed vs. adaptive policy
  - How many entries? Should table grow or shrink?
- Number of IBTCs: Shared vs. non-shared policy
- One IBTC for all indirects? An unique IBTC per indirect?
- Table lookup: *Inlined address mapping policy* 
  - Should some address mappings be done directly as a series of instructions, rather than a table lookup

6



#### Fixed IBTC Size

- Table size affects
  - Memory footprint
  - Hardware caching behaviorDistribution of addresses and conflict misses
- Small size: Good footprint, good D-cache behavior, possibly poor conflict miss behavior
- Large size: Poor footprint, possibly poor D-cache behavior, fewer conflict misses



# Adaptive IBTC Size

- Adaptively increase IBTC size on a conflict miss
  - Good memory footprint of small IBTC
  - Fewer conflict misses of large IBTC, when needed
  - Tailored to behavior of particular indirect branches
- Adaptive IBTC requires
  - Amount to grow table by (e.g., double size on a miss)
  - Re-allocating memory for different IBTC sizes
  - Re-hashing IBTC entries when growing size
  - Re-writing IBTC hash function in fragment cache



#### Shared vs. Non-shared

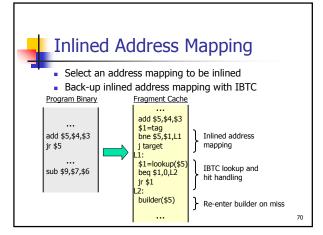
- Shared: One table for <u>all</u> indirect branches
  - Smallest memory footprint
  - Too small size: Many conflicts possible
  - Can have lower total cold start misses
- Non-shared: One table *per* indirect branch
  - Large memory footprint (depends on num. unique indirects)
     Cache contents "tailored" to each indirect branch

  - Reduce conflict misses, but higher total cold start misses
  - Use with adaptive IBTC size



# **Inlined Address Mapping**

- IBTC lookup
  - 2 memory accesses: read tag, read address
  - May cause eviction of application data from D-cache
  - Address mapping may have locality!
- Inline address mapping
  - Generate code in FC\$ to do mapping w/o table lookup
  - Inline address mapping hit faster than IBTC hit
- Policies of what to inline
  - First address mapped after translation
  - After every miss
  - After the n<sup>th</sup> miss





# **Example Configurations**

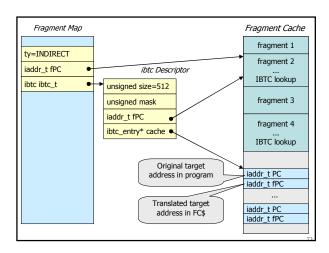
- Fixed, Shared
- Simple, small footprint
- Fixed, Non-shared
- Simple, medium to large footprint, can spread out conflicts
- Adaptive, Non-shared
  - Good memory footprint, re-hashing distributes conflicts
- Adaptive, Non-shared, Inlined address mapping
  - Inlines first address mapping to avoid IBTC access
- Adaptive, shared: Doesn't get benefit of tailoring IBTC to particular indirect branches.

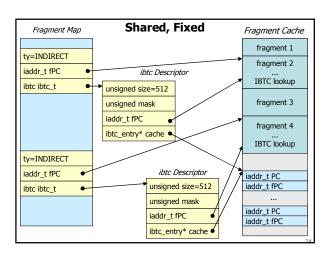
7

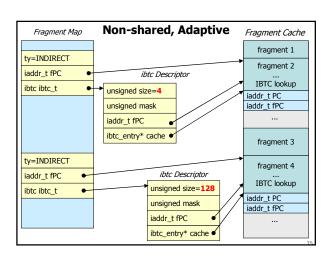


# Strata Implementation

- IBTC implementation
  - Data structures and hashing into IBTC
  - Translating indirect branches
  - IBTC lookup and hit handling
  - IBTC miss handling
- Primary example: non-shared, adaptive









# Handling Indirects

IBTC is affected at five different points:

- 1. When Strata gets control to initialize
- 2. Translating the indirect branch (translation)
- 3. IBTC cold start miss first time fragment is executed (execution)
- 4. A regular IBTC miss (execution)
- 5. A regular IBTC hit (execution)

76



# **IBTC Initialization**

- strata\_init\_ibtc()
  - Initializes & allocates IBTC data structures
  - Single structure allocated for shared policy
- strata\_reset\_ibtc()
  - Resets data structures on FC\$ flush
  - IBTC allocated in FC\$
  - Hence, a flush frees the IBTC memory

77



# Translating the Indirect

- Builder is re-entered on a new target
  - suppose target has an indirect branch
- strata\_build\_main(to\_PC,from\_frag)
  - *to\_PC* is program target with indirect branch
  - from\_frag is originating fragment in FC\$
  - uses targ\_ind\_branch(frag,next\_PC,insn)
  - terminates the fragment after indirect
  - executes the fragment



# Translating the Indirect

- targ\_ind\_branch(frag,next\_PC,insn)
  - Target dependent interface for translating the indirect
  - Does code emission for indirect & IBTC
  - frag is current fragment being translated
  - next\_PC is address of next instruction
  - *insn* is the indirect branch



# Translating the Indirect

%00,[%sp-4] %14,%g0,%o0 %00,[%sp-100] [%sp-4],%o0 stw add ldw delay slot inst # context save save %sp,-96,%sp ldw [%sp-4],%o0 # table lookup now comes next, # then the builder re-enter code

# extract target address

# do table lookup HI(ibtc\_p->cache),%o2 LO(ibtc\_p->cache),%o2 %o1,throw\_away,%o0 %o1,ibtc\_p->mask,%o1 sethi ori andi %01,3,%01 %01,%02,%02 [%02],%01 add ldw cmp bne %00,%01 LO # had a match in the IBTC ldw [%o2+4],%o3 jmpli %o3,0,%g0 restore

# had a mismatch in the IBTC

L0: re-enter the builder



#### **Cold Start Miss**

- IBTC is not loaded during indirect translation
  - For shared: May already contain valid entries
  - For non-shared: Initialized to be empty
- Hence, when fragment first executes, a cold start miss can happen
- Can't preload because the target address may actually be computed in the fragment



# Handling an IBTC miss

- 1. strata\_build\_main(to\_PC, from\_frag)
  - Re-enter when IBTC lookup didn't find a match
- 2. check for target fragment in FC\$
  - If missing: Then translate & execute fragment
  - Otherwise, handle the miss (step 3)
- 3. strata\_indirect\_branch\_miss(from, to)
- 4. execute target fragment

82



# Handling an IBTC Miss

- Miss handling depends on policies
  - Shared: Load missing mapping
  - Non-shared: Load missing mapping
  - Adaptive: Increase capacity
  - Inlining: Generate mapping for first miss
- strata\_indirect\_branch\_miss()
  - target independent implementation of the different policies

83



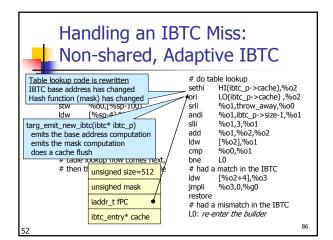
# Handling an IBTC Miss: Shared, Fixed IBTC

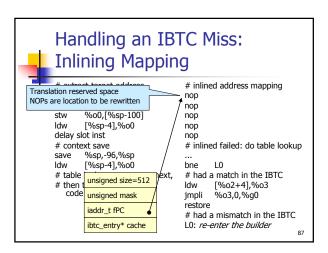
- strata\_indirect\_branch\_miss(from, to)
  - from is fragment with the indirect that missed
  - $\, \bullet \,$   $\,$  to is fragment that is the translated target of indirect
  - if to fragment wasn't translated when re-entered, it is fetched and translated to get the correct target for the indirect
- Update the IBTC
  - hash on to->PC and update with to->fPC
  - IBTC[hash(to->PC)].PC = to->PC
  - IBTC[hash(to->PC)].fPC = to->fPC

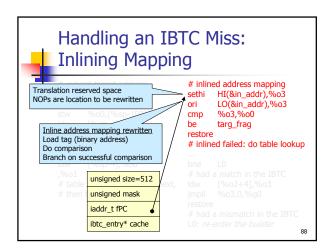


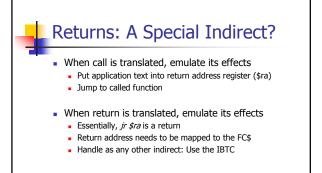
# Handling an IBTC Miss: Non-shared, Adaptive IBTC

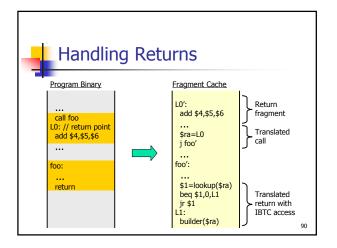
- strata\_indirect\_branch\_miss(from, to)
  - On a cold start miss, simply load the IBTC
  - On a conflict miss, possibly increase size
- Conflict miss
  - If IBTC size > MAX\_SIZE, handle like a cold start miss
  - Otherwise, allocate a new IBTC with double capacity
  - Using original IBTC, load new IBTC
  - Rewrite hash lookup in FC\$ for from fragment
  - Load mapping (to->PC,to->fPC) into new IBTC









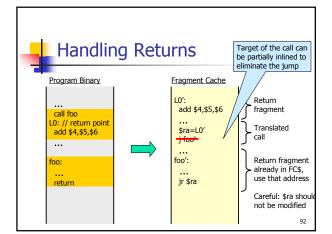




# Fast Return Handling

- Every return requires an IBTC lookup
- During call translation
  - Return fragment may already be in FC\$
  - Thus, return fragment address may be known!
  - At that point (during translation), map the return address
- Return fragment may not be translated
  - During call translation, prefetch & translate return fragment
  - Return fragment address is then known

91





# Implementing Fast Returns

- Similar to IBTC
  - Implementation is target dependent
    - $\, \bullet \,$  E.g., On MIPS, can't always identify the call
- Translating the call instruction
  - targ\_call(frag, PC, insn)
- Translating the return instruction
  - targ\_return(frag, PC, insn)



# **Translating Call**

- targ\_call(frag, PC, insn)
  - frag is the fragment with the call being translated
  - *PC* is the next program address (i.e., the return address)
  - insn is the actual call
- Steps for translation
  - Fetch and translate the delay slot instruction (SPARC/MIPS)
  - Lookup return address (PC+8 on the SPARC)
  - If return fragment untranslated, fetch & translate it
  - Emit code to load return address
  - Emit the delay slot instruction

94



# Translating Call and Return

#### Regular Call

# PC is program return address sethi HI(PC),%o7 ori %o7,LO(PC),%o7

# partially inlined block

# fragment with return

IBTC lookup to handle return

#### Call with Fast Return

# ret is the return fragment sethi HI(ret->fPC-8),%o7 ori %o7,LO(ret->fPC-8),%o7

# partially inlined block

# fragment with return

... jmpl %o7,%g0



# Translating Call and Return

#### Regular Call

# PC is program return address sethi HI(PC),%o7

%07,LO(PC),%07

# Return address has been mapped to a fragment address! # fragment with return

...
IBTC lookup to handle return

IBTC access replaced by an actual return through %07

#### Call with Fast Return

# ret is the return fragment sethi HI(ret->fPC-8),%o7 ori %o7,LO(ret->fPC-8),%o7

# partially inlined block

# fragment with return

\_jmpl %o7,%g0

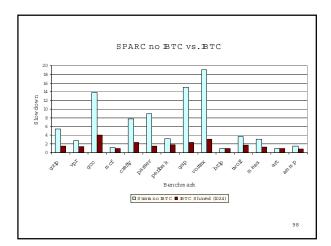
32

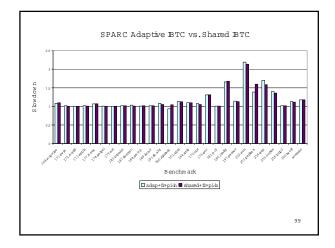


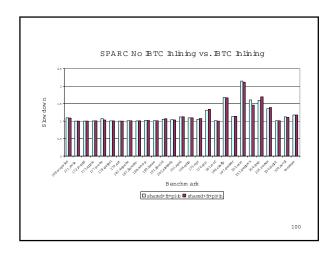
# Outline

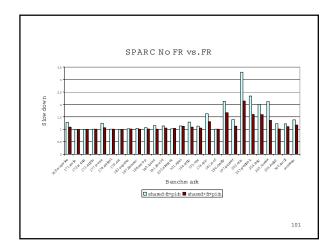
- Part I: Introduction to SDT
- Part II: The Strata SDT Framework
   Translation virtual machine

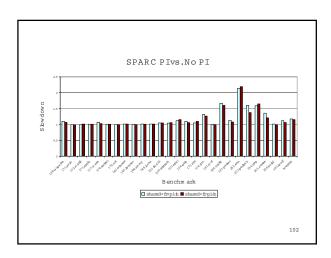
  - Handling indirect branches
     Performance
- Part III: Code security applications

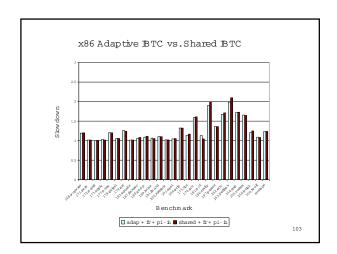


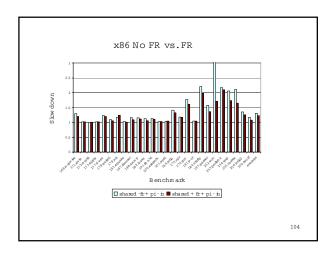


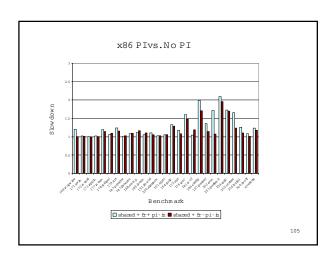


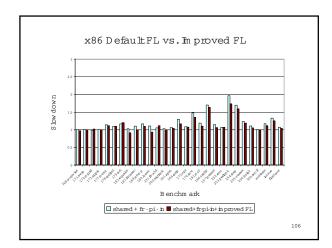














#### Outline

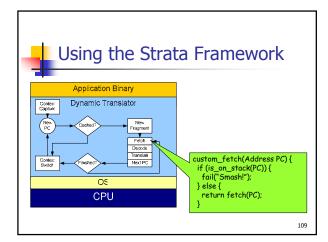
- Part I: Introduction to SDT
- Part II: The Strata SDT Framework
  - Translation virtual machine
  - Handling indirect branches
  - Performance
- Part III: Code security applications

107



# Software Security with SDT

- SDT provides support for a variety of approaches to making software resistant to attack
  - Enforcement of security policies
    - Examine and modify code before execution
    - Control what code is executed
    - Control use of resources





#### Related Work

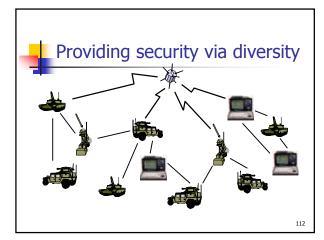
- Jones, M. B., Interposition agents: Transparently interposing user code at the system interface. In *Symposium on Operating Systems Principles* (1993), pp. 80–93.

  Scott, J. K., and Davidson, J. W., Strata: A software dynamic translation infrastructure. In *IEEE Workshop on Binary Translation* (2001).
- Kiriansky, V., Bruening, D., Amarasinghe, S., Secure Execution via Program Shepherding, Proceedings of the 11th USENIX Security Symposium, August 2002, pp. 191-206.
- Scott, J. K. and Davidson, J. W., Safe Virtual Execution Using Software Dynamic Translation. In *Proceedings of the 18th Annual Computer Security Applications Conference*, Las Vegas, NV, December 2002, pp. 209–218.



# Providing Security via Diversity

- Biological systems are resilient because they have diversity
- Computer systems are vulnerable because they have no diversity
  - Software monoculture





# Providing Security via Diversity

- Introduce artificial diversity using SDT
  - ASR: address space randomization
  - ISR: instruction set randomization
  - Calling convention diversity
- Provides resiliency against attacks against distributed systems

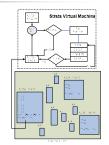


# Strata diversity transformations

- Default diversity
  - Address space randomization
    - Code is relocated in F\$
  - Run-time stack is modified
     Control-flow randomization
    - Control-flow randomization

      Basic block structure is modified (no unconditional branches, direct function calls can be eliminated)

      Indirect jumps and calls transformed





# Strata Diversity is Effective

- Simple Demonstration

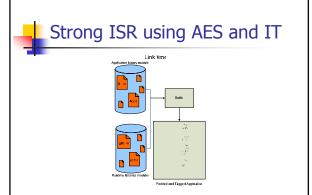
  - nple Demonstration
    Compile code normally and
    execute and attack is successful
    (you get a shell) (shell-normal)
    Compile code and run under
    control of Strata and attack fails
    (program terminates normally)
    (shell-strata)
    Adapt attack and run under
    Strata and attack works (shell-strata-adapt)

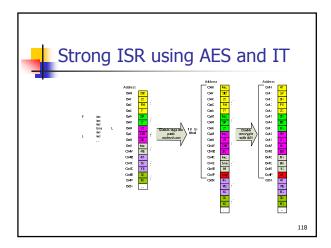
int main (int argc, char \*argv[]) {
 naive();
 printf("Nothing bad happened!\n");

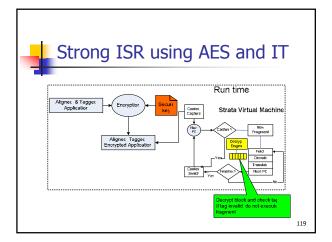


#### **Instruction Set Randomization**

- Encrypt application code prior to execution
- Decrypt code before it is executed
- Malicious code that is injected through some software vulnerability will be decrypted but because it was not encrypted, the resulting code, will not execute properly
- See
  - Randomized Instruction Set Emulation to Disrupt Binary Code Inject Attacks, Barrantes, Ackley, Forrest, et. al, CCS 03.
  - Countering Code-Injection Attacks with Instruction-set Randomization, Kc, Keromytis, Prevelakis, CCS 03.



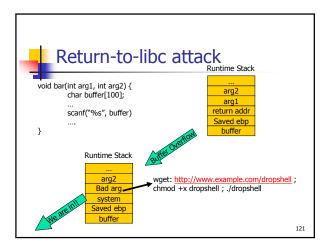


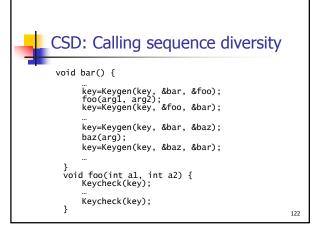




# CSD: Calling sequence diversity

- Compile-time/runtime technique to create a software population with many different calling sequences
- Effective defense against "return-to-libc" attacks (also known as arc injection, Pincus and Baker, IEEE Security and Privacy, 2(4), pp. 20-27)
  - Return-to-libc does not require injecting code into the application
  - ISR is not an effective defense against return-tolibc type attacks

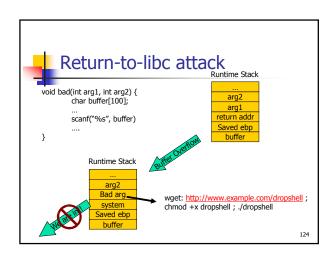


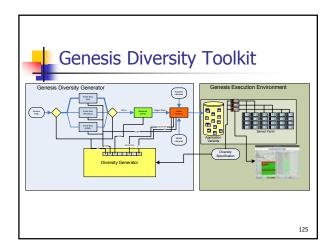


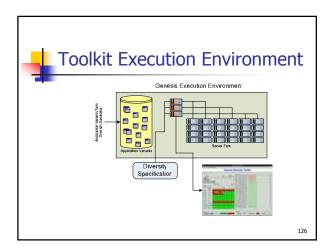


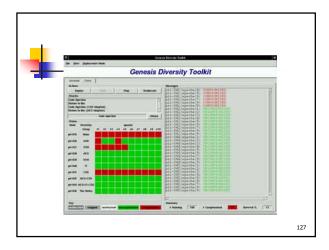
#### CSD: Calling sequence diversity

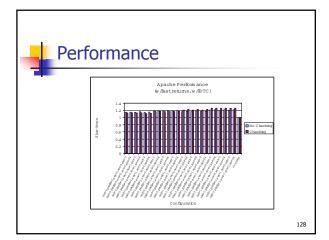
- Calls to Keygen and Keycheck routines are inserted by the compiler front end (lcc, edg, Phoenix)
- At runtime:
  - Strata generates a key for each function (stored in protected region)
  - Replaces calls with inline code to generate proper key or check that the key has the proper value













#### **Future Research**

- Architectural support for SDT
- Compile-time support for SDT
- Use SDT to detect and block nextgeneration viruses and worms (polymorphic, metamorphic, timebombs, and logic bombs, etc.)
- Use SDT to recover from attacks and automatically generate a patch