

# HDFI: Hardware-Assisted Data-flow Isolation

Presented by Ben Schreiber

**Chengyu Song**<sup>1</sup>, Hyungon Moon<sup>2</sup>, Monjur Alam<sup>1</sup>, Insu Yun<sup>1</sup>,  
Byoungyoung Lee<sup>1</sup>, Taesoo Kim<sup>1</sup>, Wenke Lee<sup>1</sup>, Yunheung Paek<sup>2</sup>

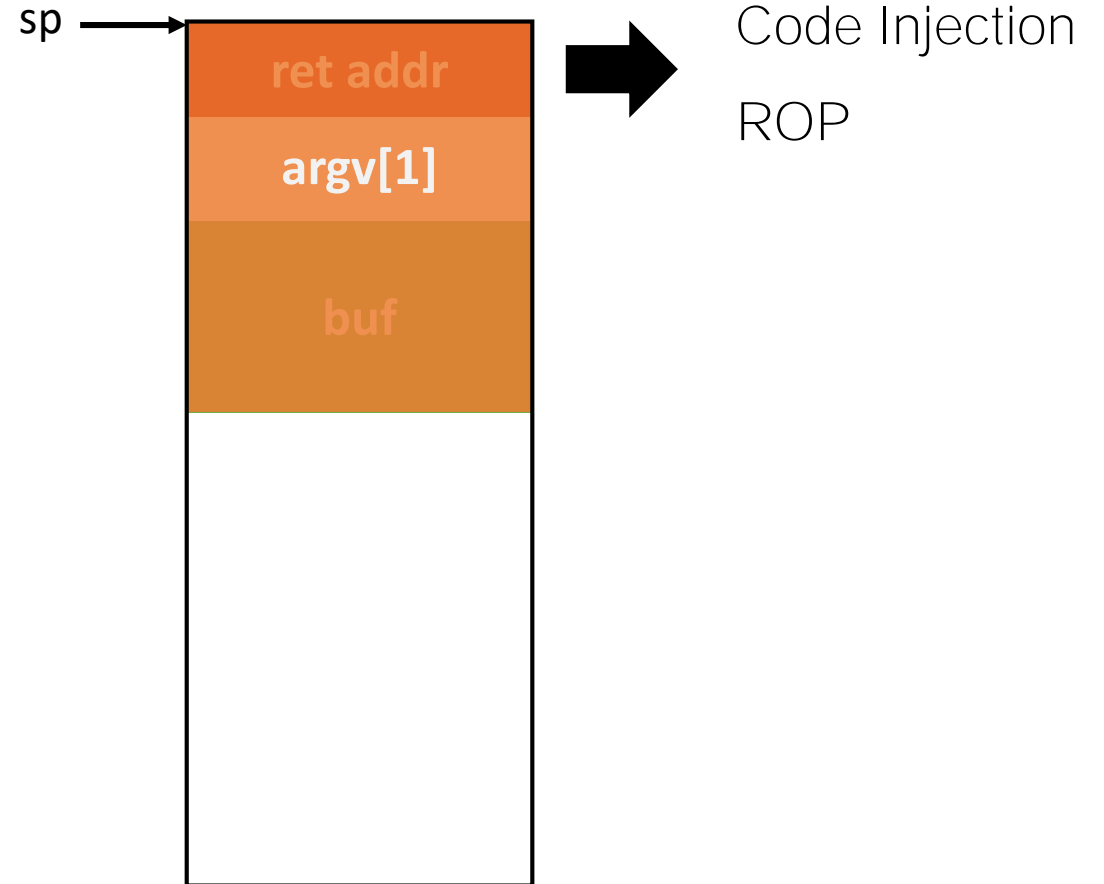
<sup>1</sup>**Georgia Institute of Technology**

<sup>2</sup>Seoul National University

# A simple stack overflow

```
int main(int argc, const char *argv[]) {  
    char buf[16];  
    strcpy(buf, argv[1]);  
    return 0;  
}
```

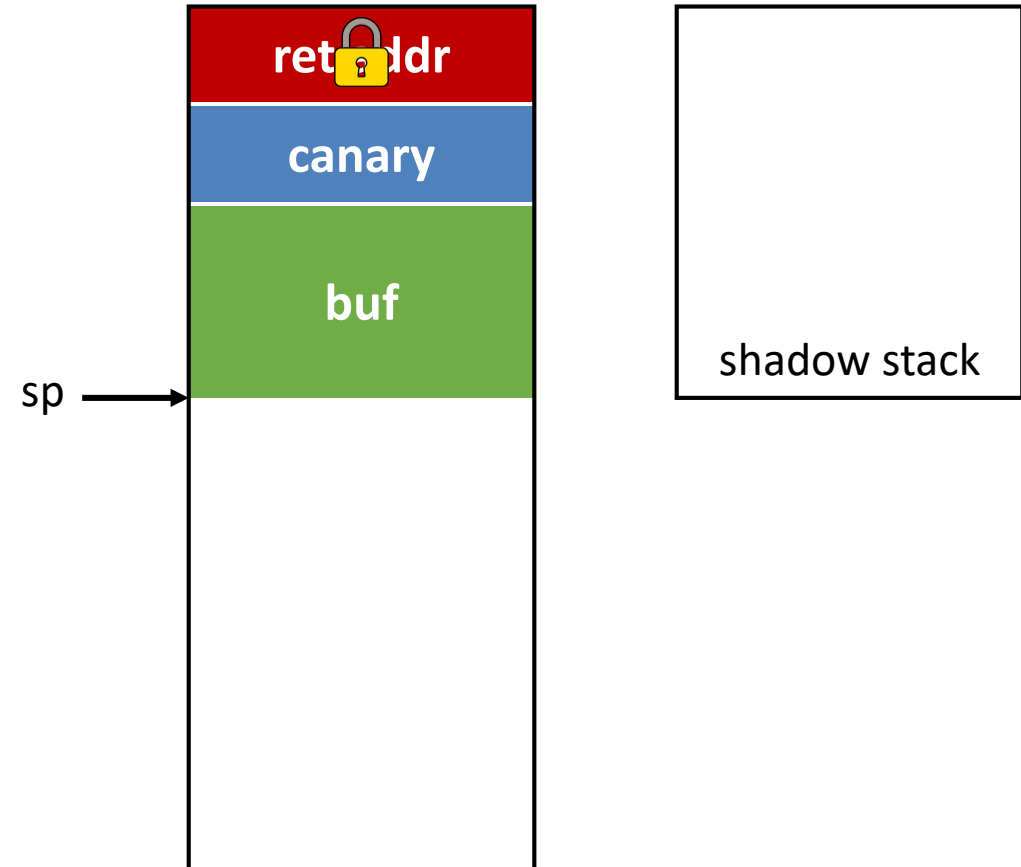
```
1 main:  
2 add    sp, sp, -32  
3 sd     ra, 24(sp)  
4 ld     a1, 8(a1)    ; argv[1]  
5 mv     a0, sp       ; char buff[16]  
6 call   strcpy      ; strcpy(buff, argv[1])  
7 li     a0, 0  
8 ld     ra, 24(sp)  
9 add    sp, sp, 32  
10 jr    ra           ; return
```



# Defense mechanisms

```
int main(int argc, const char *argv[]) {  
    char buf[16];  
    strcpy(buf, argv[1]);  
    return 0;  
}
```

```
1 main:  
2   add    sp, sp, -32  
3   sd    ra, 24(sp)  
4   ld    a1, 8(a1)      ; argv[1]  
5   mv    a0, sp         ; char buff[16]  
6   call  strcpy        ; strcpy(buff, argv[1])  
7   li    a0, 0  
8   ld    ra, 24(sp)  
9   add   sp, sp, 32  
10  jr    ra              ; return
```



# Prior Limitations

- Software: lacks good isolation mechanisms in 64-bit world
  - SFI and virtual address space: **secure** but **expensive**
  - Address randomization: **efficient** but **insecure**
- Hardware: lacks **flexibility**
  - Context saving/restoring (setjmp/longjmp), deep recursion, kernel stack, etc.
  - Other data: code pointers, non-control data
- Data shadowing: adds **overheads**
  - Breaks data locality, needs additional step to look up or reserved register(s)
  - Occupies additional memory

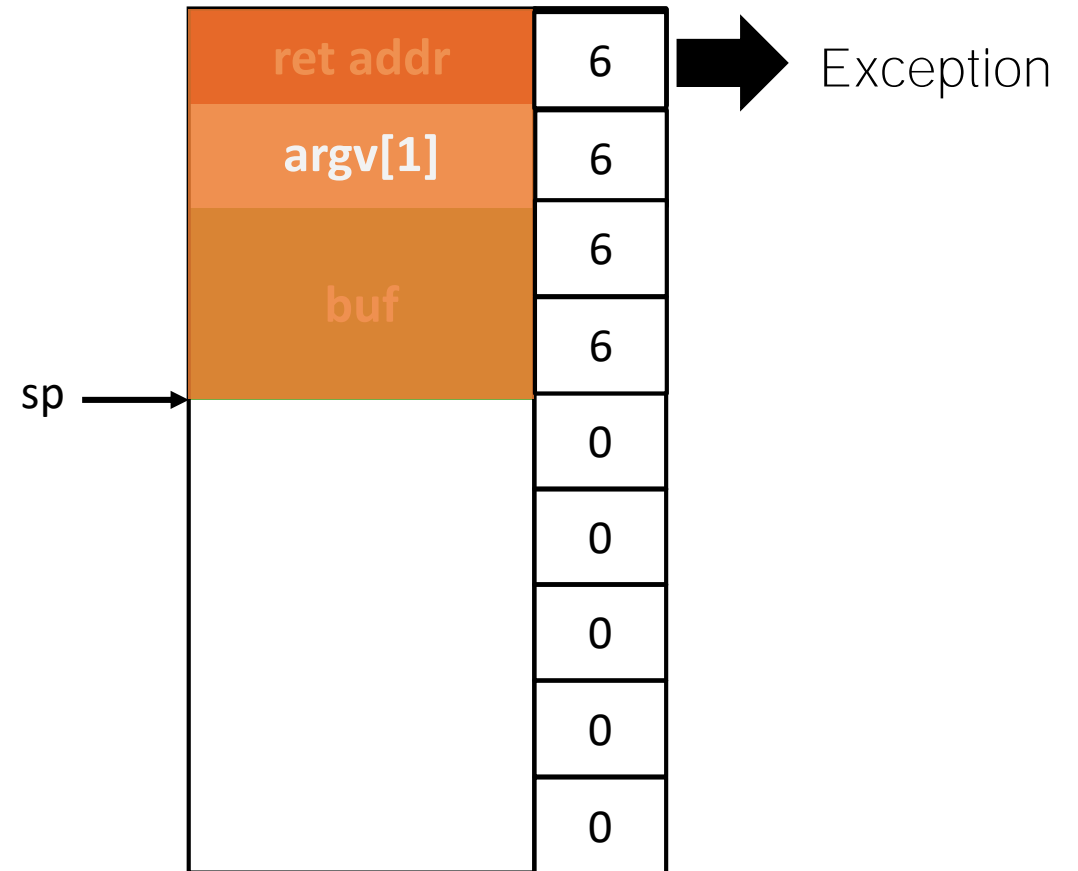
# Hardware-assisted data-flow isolation

- **Secure** and **efficient**
  - Low performance overhead and strong security guarantees
- **Flexible**
  - Capable of supporting different security model/mechanisms
- **Fine-grained**
  - No more data-shadowing
- **Practical**
  - Minimized hardware changes

# Data-flow Integrity [OSDI'06]

Runtime data-flow should not deviate from static data-flow graph

```
1 main:
2   add    sp, sp, -32
3   sd     ra, 24(sp)
4   ld     a1, 8(a1)    ; argv[1]
5   mv     a0, sp      ; char buff[16]
6   call   strcpy      ; strcpy(buff, argv[1])
7   li     a0, 0
8   ld     ra, 24(sp)
9   add    sp, sp, 32
10  jr     ra          ; return
```

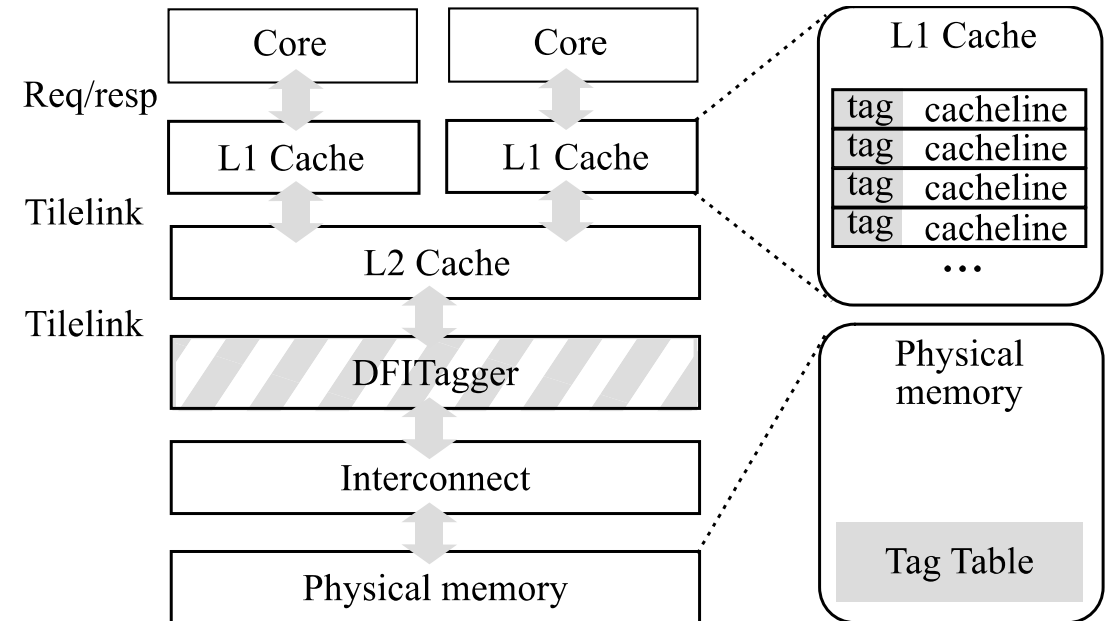


# ISA extension

- Tagged memory
  - Machine word granularity
  - Fixed tag size → currently only *1 bit* (sensitive or not)
- Three new *atomic* instructions to enable **DFI-style** checks
  - `sdset1`, `ldchk0`, `ldchk1`
- New semantic of old instructions (backward compatible)
  - `sd` : `sdset0`
  - `ld` : does not check tags

# Hardware extension

- Cache extension
  - Extra bits in the cache line for storing the tag (reusing existing cache coherence interconnect)
- Memory Tagger
  - Emulating tagged memory without physically extending the main memory
  - Tags are stored in separate table in protected main memory





# Optimizations

- Memory Tagger introduces additional performance overhead
  - Naive implementation: 2x memory accesses, 1 for data, 1 for tag
- Three optimization techniques
  - Tag cache
  - Tag valid bits (TVB)
  - Meta tag table (MTT)

# Tag Cache

- A 64 byte cache line needs only 8 tag bits
- The memory interface is a fixed 64 bytes
  - Many lines of tags (enough for 4 KB block) can be fetched at the same time
- Tag cache located within DFITagger block
  - 64 byte entries corresponding to 4 KB of memory

# Tag Valid Bits (TVB)

- Observation: a majority of memory accesses do not check tags
- If the processor does not ask for tags, do not lookup/supply tags
  - TVB indicates whether tags are present with data in cache
- Loads will ask DFITagger for tags if they are needed and not present
- Writes always set TVB

# Meta Tag Table

- Observation: most memory is tagged with 0
- Meta Tag Table is an in memory data structure where each bit indicates whether a group of entries in the Tag Table are all zero
- Meta Tag Directory is a register that functions similarly on the Meta Tag Table
- These structures can eliminate some memory reads

# Return address protection

- Policy: return address should always have tag 1
- Benefits: secure and supports context saving/restoring, deep recursion, modified return address, kernel stack

```
1 main:
2   add    sp,sp,-32
3   *sdset1 ra,24(sp)
4   ld     a1,8(a1)      ; argv[1]
5   mv     a0,sp         ; char buff[16]
6   call  strcpy        ; strcpy(buff, argv[1])
7   li     a0,0
8   *ldchk1 ra,24(sp)
9   add    sp,sp,32
10  jr     ra            ; return
```

# Standard Library Protections

- Heap Metadata
  - Protect pointers within ptmalloc
- Global Offset Table
  - GOT is a data structure used for dynamic linking
  - Add tags to table at same time as ASLR is applied
- Exit handler
  - Currently pointer is (weakly) encrypted
  - Use tag on pointer to ensure security

# Various applications

	Application	Security Policy (invariants)
Integrity Protection	Shadow Stack	return address and register spills should have tag 1 (push / pop)
	vptr Protection	vptr should have tag 1 (constructor / virtual function call)
	Code Pointer Separation	code pointer should have tag 1 (CPI [OSDI'14])
	C Library Enhancement	important data/pointers should have tag 1 (manual modification)
	Kernel Protection	sensitive kernel data should have tag 1 (Kenali [NDSS'16])
Leak Detection	Heartbleed Prevention	crypto keys should have tag 1
		output buffer should have tag 0

# Implementations

- Hardware
  - RISC-V RocketCore generator: 2198 LoC
  - Instantiated on Xilinx Zynq ZC706 FPGA board
- Software (RISC-V toolchain)
  - Assembler gas: 16 LoC
  - Kernel modifications: 60 LoC
  - Security applications: 170 LoC



# Effectiveness of optimizations

- Memory bandwidth and latency

Benchmark	Tag Cache	+TVB	+MTT	+TVB+MTT
L1 hit	0%	0%	0%	0%
L1 miss	14.47%	5.26%	14.47%	5.26%
Copy	13.14%	4.44%	11.84%	4.26%
Scale	10.62%	4.79%	9.45%	4.67%
Add	4.37%	1.26%	4.13%	1.2%
Triad	9.66%	1.96%	8.8%	1.83%

- SPEC CINT2000

Benchmark	Tag Cache	+TVB	+MTT	+TVB+MTT
164.gzip	16.09%	2.18%	6.85%	1.87%
175.vpr	29.51%	3.26%	7.71%	1.43%
181.mcf	36.89%	3.08%	13.66%	-0.11%
197.parser	16.11%	2.27%	7.61%	1.53%
254.gap	12.19%	1.04%	6.53%	0.71%
256.bzip2	14.52%	2.65%	3.63%	0.84%
300.twolf	26.71%	2.97%	7.37%	0.36%

- Negative value is artifact from minor differences between runs

# Security experiments

- With synthesized attacks

<b>Mechanism</b>	<b>Attacks</b>	<b>Result</b>
Shadow stack	RIPE	✓
Heap metadata protection	Heap exploit	✓
VTable protection	VTable hijacking	✓
Code pointer separation (CPS)	RIPE	✓
Code pointer separation (CPS)	Format string exploit	✓
Kernel protection	Privilege escalation	✓
Private key leak prevention	Heartbleed	✓

# Impacts on security solutions

- Security
  - Hardware-enforced isolation
- Simplicity
  - No data shadowing
- Usability
  - Implementation/port is very easy

Application	Language	LoC
Shadow Stack	C++ (LLVM 3.3)	4
VTable Protection	C++ (LLVM 3.3)	40
CPS	C++ (LLVM 3.3)	41
Kernel Protection	C (Linux 3.14.41)	70
Library Protection	C (glibc 2.22)	10
Heartbleed Prevention	C (OpenSSL 1.0.1a)	2

# Impacts on security solutions (cont.)

- Efficiency
  - GCC (-O2)
  - Clang (-O0)

Benchmark	Shadow stack (GCC)	SS+CPS (Clang)
164.gzip	1.12%	2.42%
181.mcf	1.76%	3.54%
254.gap	3.34%	13.23%
256.bzip2	3.05%	4.61%

# Security analysis

- Attack surface
  - Inaccuracy of data-flow analysis
  - Deputy attacks
- Best practices
  - CFI is necessary (e.g., CPS + shadow stack)
  - Recursive protection of pointers
  - Guarantee the trustworthiness of the written value
  - Use runtime memory safety technique to compensate inaccuracy of static analysis

# Limitations and Improvements

- Direct memory access (DMA)
  - Attacker could directly alter tag table in memory
- Further optimizations
  - Add tag prefetch, improved cache, etc. to complement an OoO core
- Dynamic Code Generation
  - Not currently supported, but tags could be used to secure JIT'ed code

**Q & A**

**Thank you!**