



DR. KEVIN HAMLLEN

LOUIS A. BEECHERL, JR. DISTINGUISHED PROFESSOR  
COMPUTER SCIENCE DEPARTMENT  
CYBER SECURITY RESEARCH AND EDUCATION INSTITUTE  
THE UNIVERSITY OF TEXAS AT DALLAS

# SOFTWARE ATTACK SURFACE REDUCTION ON THE FLY

Dr. Kevin Hamlen

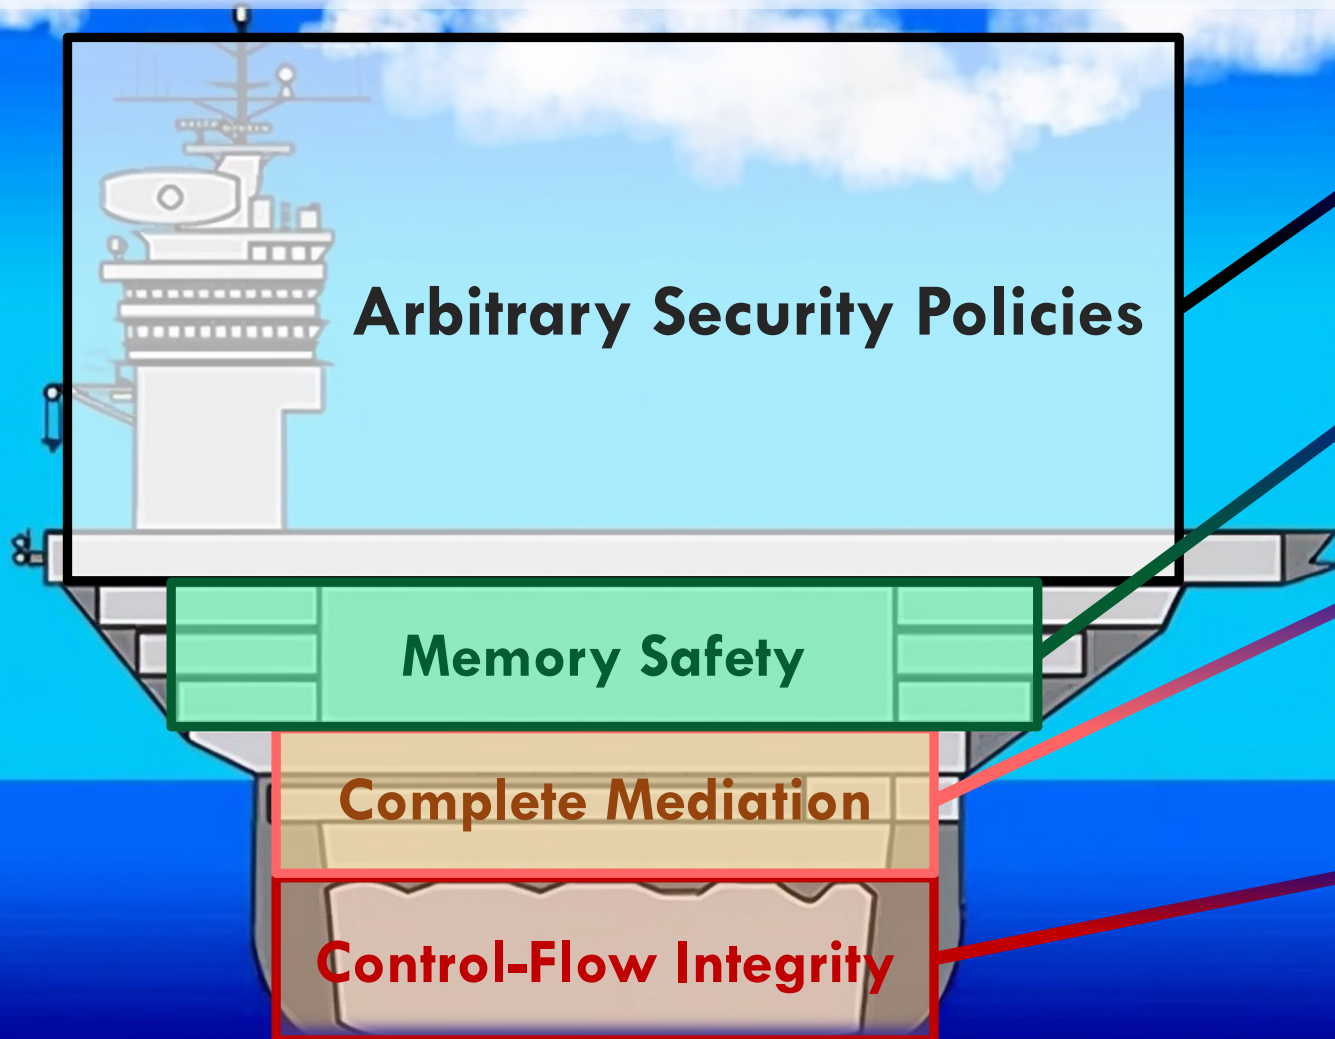
ONR Award N00014-21-1-2654



Any opinions, findings, conclusions, or recommendations expressed in this presentation are those of the author(s) and do not necessarily reflect the views of ONR, UTD or other supporters.

# Foundations of Software Security

2/31



- Safety + Liveness = Everything
- Safety: “bad events” prohibited
- Liveness: “good events” guaranteed

- Memory Safety
- Secure security state storage

- Complete Mediation
- Surround all security-relevant operations with uncircumventable guard code (security checks)

- Control-Flow Integrity (CFI)
- Constrains execution order to a graph of allowed flows
- Affords provably uncircumventable basic blocks of program instructions

# Foundations of Software Security

3/31

Static Code Instrumentation  
Dynamic Policy Enforcement

Arbitrary Security Policies

Memory Safety

Complete Mediation

Control-Flow Integrity

```
for (int i = 0; i < 32; ++i) {  
    if (safe(&buf[i]))  
        buf[i] = i;  
    else  
        abort("memory corruption!");  
    if (++counter > max)  
        abort("deadlock alert");  
    if (check(&obj->method)  
        && valid_args(buf, i)) {  
        obj->method(buf, i);  
    } else  
        abort("security alert!");  
}
```

# Binary Code Debloating Architecture


4



trace  
whitelist



- 1) Learn debloating policy from execution traces
  - ❑ consumer has no formal policy specification
  - ❑ consumer is not aware of all “undesired” program functionalities
- 2) Machine learning derives suitable policy from a whitelist of traces
- 3) Enforce learned policy with source-free, **context-sensitive CFI**
  - ❑ generalizes and subsumes non-contextual CFI and code byte erasure
- 4) Machine-validate binary hardening transforms for highest assurance
  - ❑ Picinæ: Platform In Coq for Instruction-level Analysis of Executables



Machine Learning




policy



Binary, Context-sensitive CFI



debloated  
binary




Formal Validation

# Binary Code Debloating Architecture

5



trace  
whitelist



Machine Learning

- 1) Learn debloating policy from execution traces
  - ❑ consumer has no formal policy specification
  - ❑ consumer is not aware of all “undesired” program functionalities
- 2) Machine learning derives suitable policy from a whitelist of traces
- 3) **Enforce learned policy with source-free, context-sensitive CFI**
  - ❑ generalizes and subsumes non-contextual CFI and code byte erasure
- 4) Machine-validate binary hardening transforms for highest assurance
  - ❑ Picinæ: Platform In Coq for Instruction-level Analysis of Executables



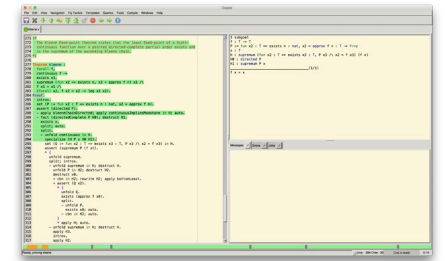
policy



Binary, Context-sensitive CFI

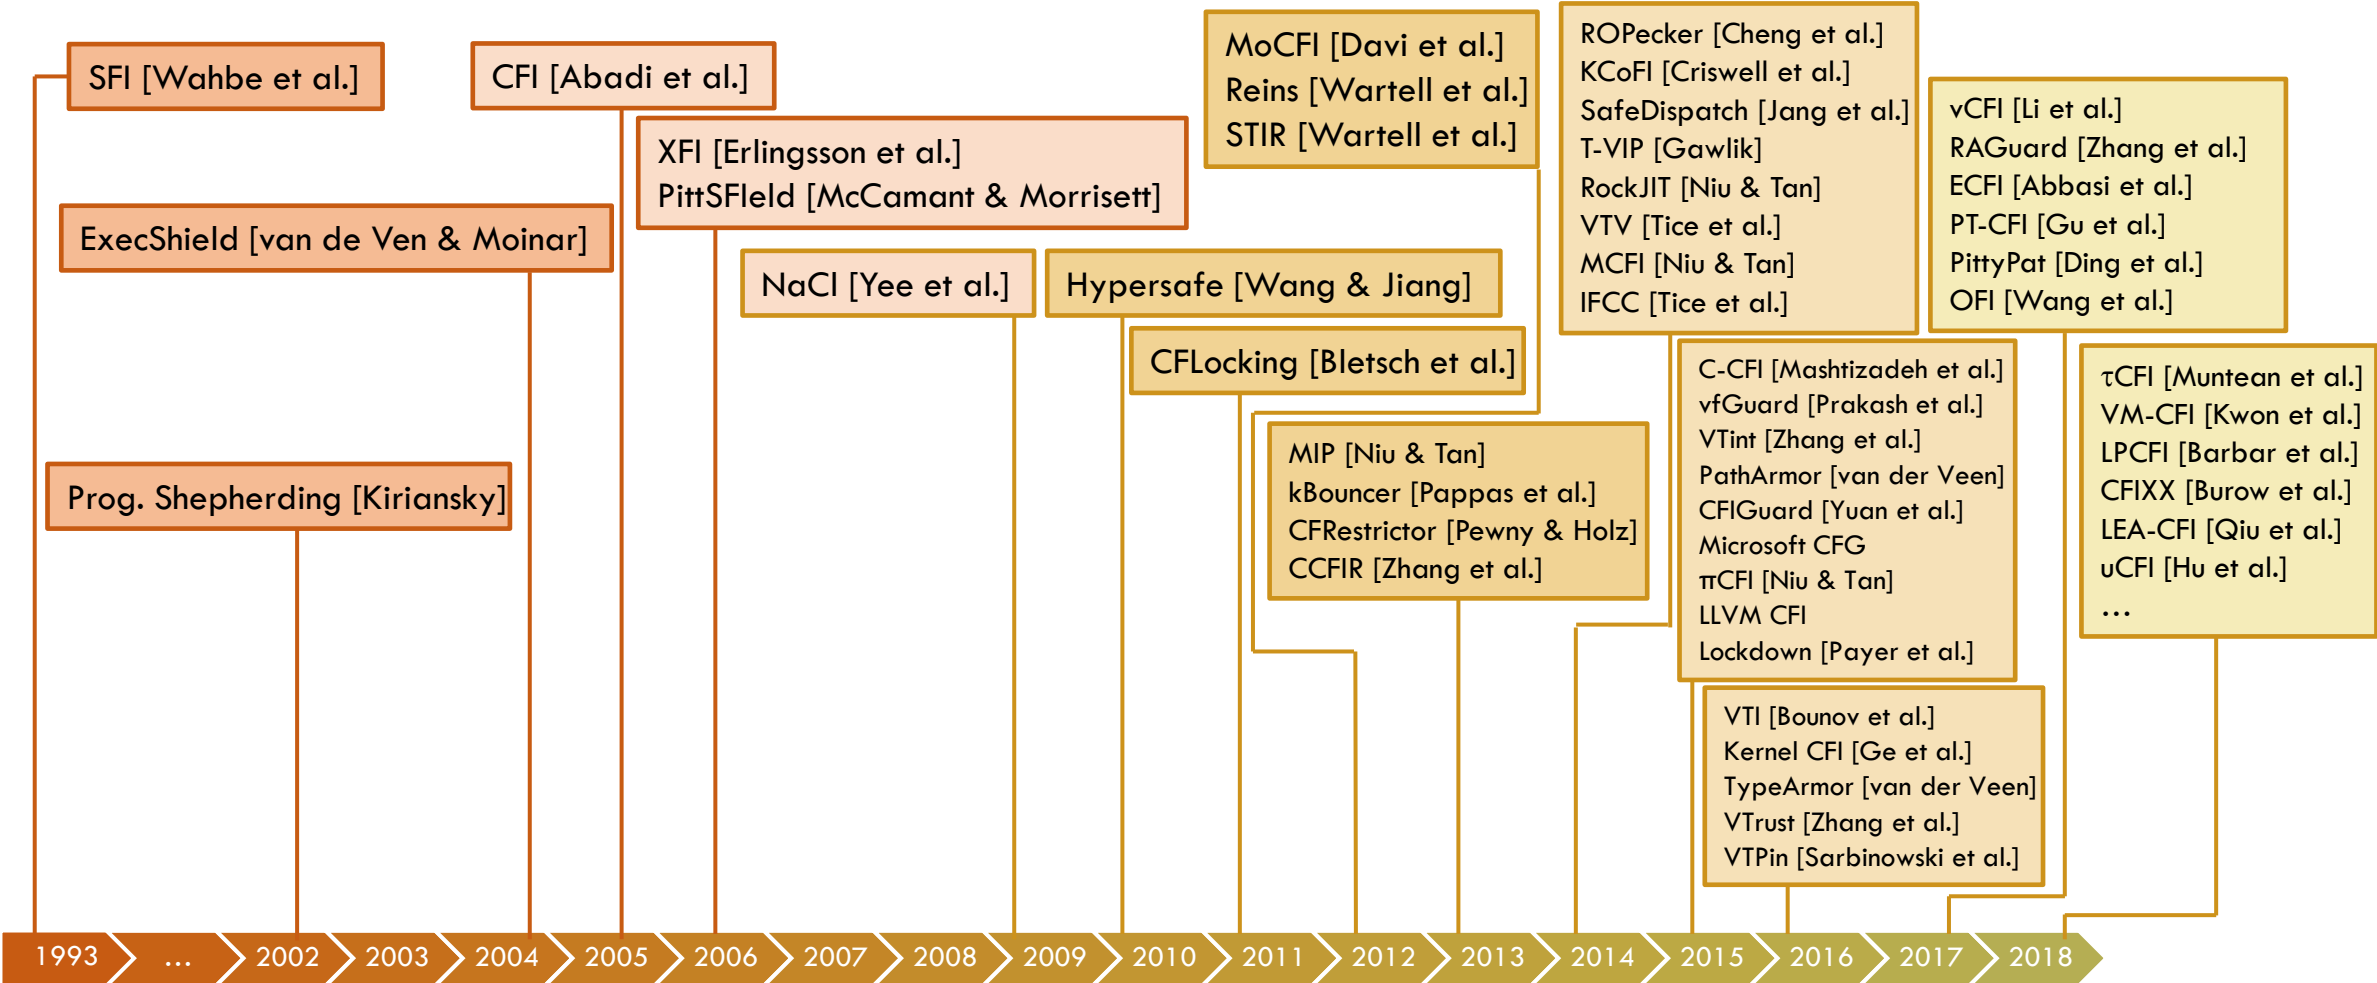


debloated  
binary

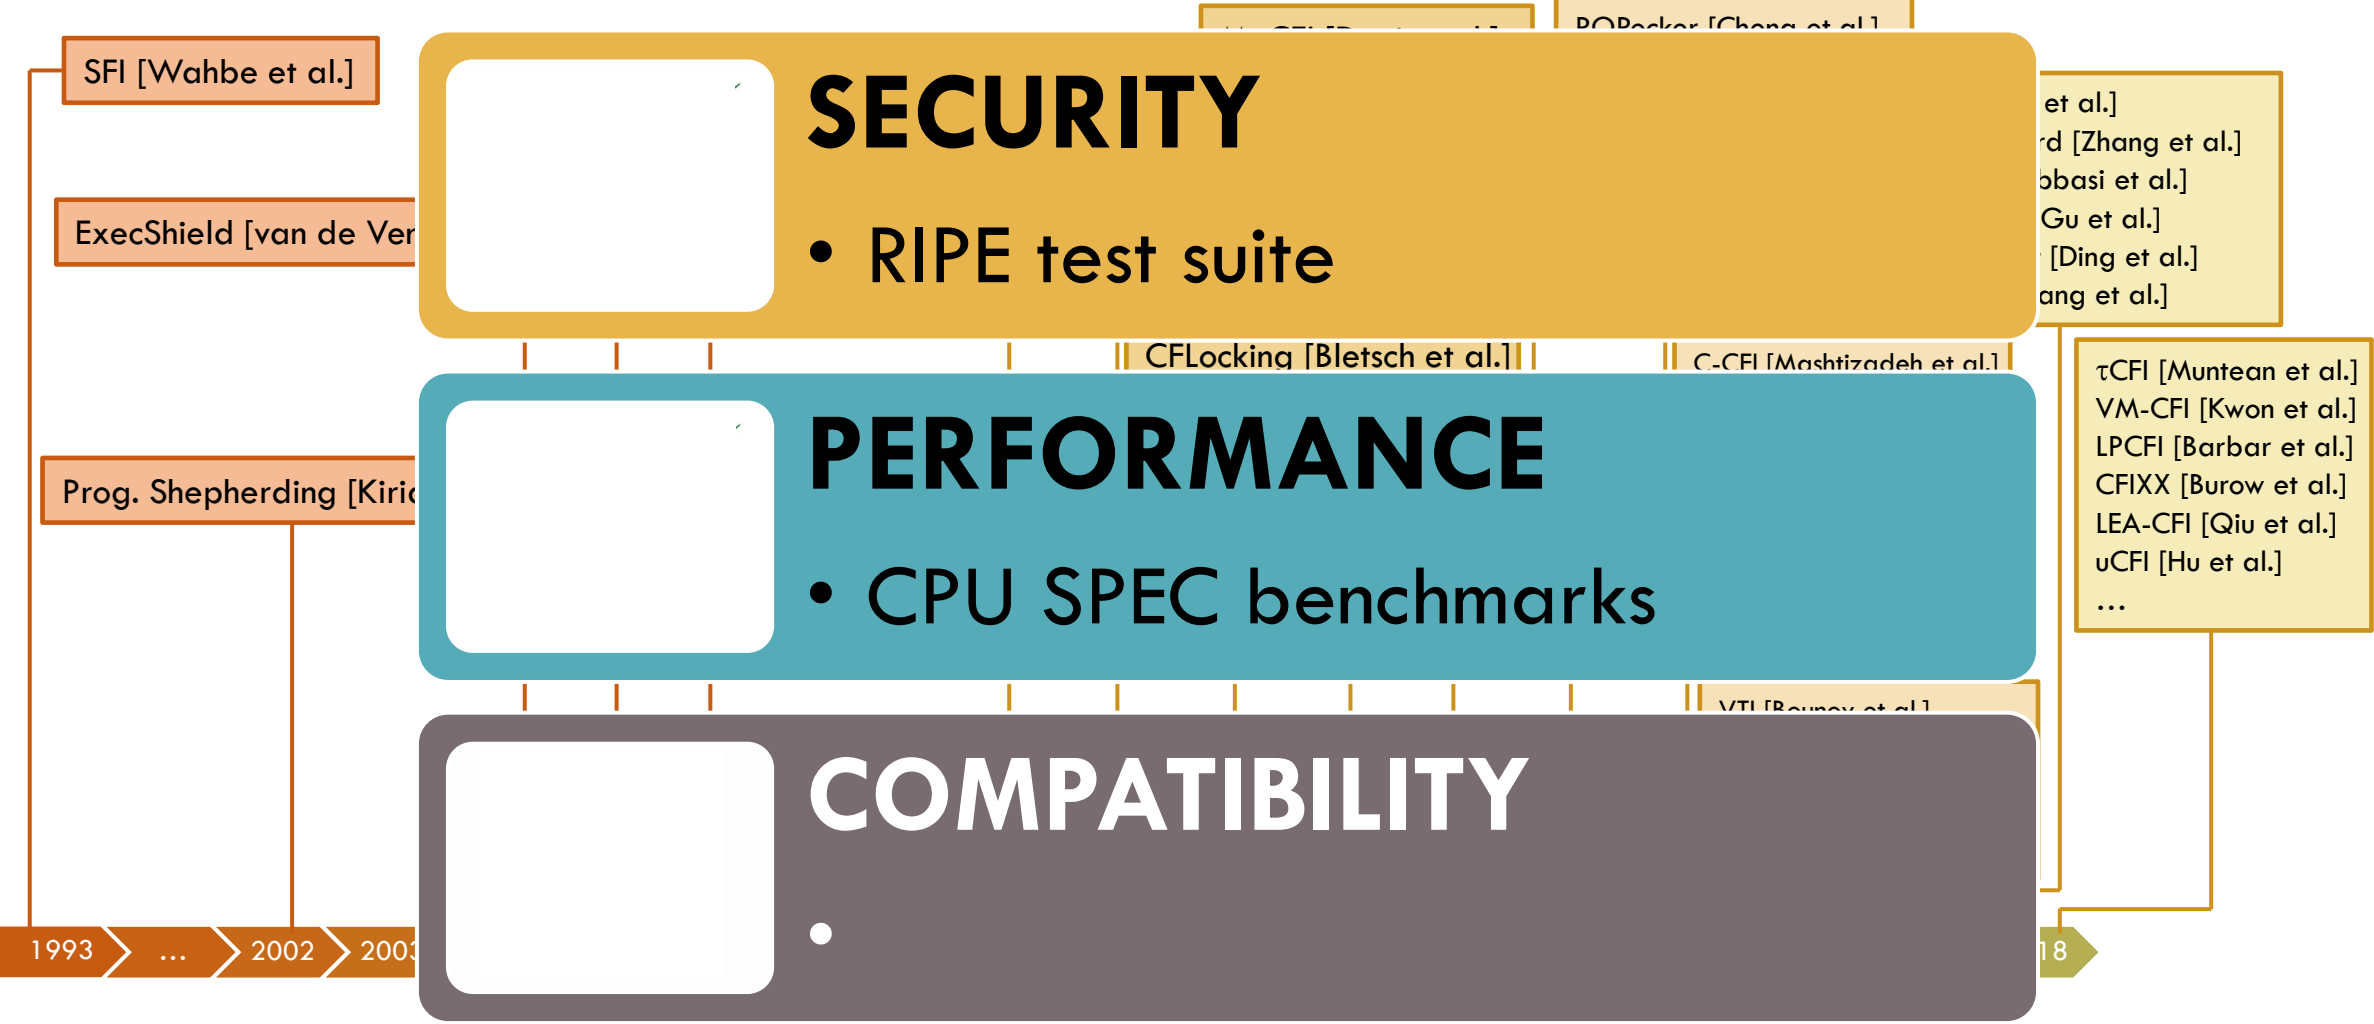


Formal Validation

# CFI Research Timeline






# CFI Research Timeline



# Scalability Gap

8/31




\*Papers containing at least one experiment where at least one **COMPLETE** non-benchmark application for the indicated OS was rewritten & secured



# Where the Wild Things Are


[RecordedFuture, CTA-2018-0327]

### Top 10 Vulnerabilities Used by Cybercriminals



DATA RANGE: JAN. 1, 2017-DEC. 31, 2017

### Desktop OS Market Share



# Soft(ware) Targets

10/31

- ❑ Windows/MacOS in mission-critical environments
  - ❑ “About 75% of control systems are on **Windows XP** or other **nonsupported OSES.**”  
–Daryl Haegley, Office of Assistant Secretary of Defense for Energy, Installations and Environment
  - ❑ More than 25% of all government computers currently run an **outdated Windows** or **MacOS** operating system. [BitSight, 6/1/17]
  - ❑ DHS, Coast Guard, and Secret Service currently store top secret information on **outdated Windows 2003** servers. [OIG-18-56, 3/1/18]
  - ❑ Hundreds of *satellites* run **Windows 95** and/or are controlled by **Windows Mobile** devices.



# 20 Widespread Classes of CFI Compatibility Problems

11/31

<b>Compatibility Metric</b>	<b>Real-world Software Examples</b>
<b>Function Pointers</b>	7-Zip, Adobe Reader, Apache, Calculator, Chrome, Dropbox, Firefox, JVM, ...
<b>Callbacks</b>	7-Zip, Adobe Reader, Apache, Calculator, Chrome, Dropbox, Firefox, JVM, ...
<b>Dynamic Linking</b>	7-Zip, Adobe Reader, Apache, Calculator, Chrome, Dropbox, Firefox, JVM, ...
<b>Delay-Loading</b>	Adobe Reader, Calculator, Chrome, Firefox, JVM, MS Paint, MS Powerpoint, ...
<b>Exporting/Importing Data Symbols</b>	7-Zip, Apache, Calculator, Chrome, Dropbox, Firefox, MS Paint, MS Powerpoint, ...
<b>Virtual Functions</b>	7-Zip, Adobe Reader, Calculator, Chrome, Dropbox, Firefox, JVM, Notepad, ...
<b>Writable Vtables</b>	programs with UI's based on GTK+ (Linux) or COM (Windows)
<b>Tail Calls</b>	programs compiled with tail-call optimization (e.g., -O2 or /O2)
<b>Switch-Case Statements</b>	7-Zip, Adobe Reader, Apache, Calculator, Chrome, Dropbox, Firefox, JVM, ...
<b>Returns</b>	almost every benign program
<b>Unmatched Call/Return Pairs</b>	Adobe Reader, Apache, Chrome, Firefox, JVM, MS PowerPoint, Visual Studio, ...
<b>Exceptions</b>	7-Zip, Adobe Reader, Apache, Calculator, Chrome, Dropbox, Firefox, JVM, ...
<b>Calling Conventions</b>	almost every program has functions
<b>Multithreading</b>	7-Zip, Adobe Reader, Apache, Calculator, Chrome, Dropbox, Firefox, JVM, ...
<b>TLS Callbacks</b>	Adobe Reader, Chrome, Firefox, MS Paint, TeXstudio, UPX
<b>Position-Independent Code</b>	7-Zip, Adobe Reader, Apache, Calculator, Chrome, Dropbox, Firefox, JVM, ...
<b>Memory Management</b>	7-Zip, Adobe Reader, Apache, Chrome, Dropbox, Firefox, MS PowerPoint, ...
<b>JIT Code</b>	Adobe Flash, Chrome, Dropbox, Firefox, JVM, MS PowerPoint, PotPlayer, ...
<b>Self-Unpacking</b>	programs decompressed by self-extractors (e.g., UPX, NSIS)
<b>Runtime API Hooking</b>	Microsoft Office, including MS Excel, MS PowerPoint, etc.

# ConFIRM CFI Compatibility Benchmark Suite

12

**ConFirm CFI Compatibility Benchmark Suite**

<https://github.com/SoftwareLanguagesSecurityLab/Confirm>





# Major Findings

14

- Multithreading + Unmatched call/return pairs = Trouble
  - ▣ unmatched call/returns arise from: exceptions, tail-call optimization
  - ▣ cross-thread stack smashing beats all CFI defenses we tested
  - ▣ seems hard to fix without huge performance overheads
- Runtime Code Generation
  - ▣ more prevalent than generally expected
    - rise of JIT-compiled languages, runtime hooking, self-extracting components
  - ▣ most RCG is beyond the reach of all CFI algorithms
- Questionable whether SPEC CPU adequately tests CFI performance
  - ▣ SPEC CPU benchmarks chosen/designed to test CPU speeds
  - ▣ Operation profiles prioritize opcodes that bottleneck non-CFI software
  - ▣ Mostly simple control-flow graphs

# CFI Performance Measurement Problems

15

SPEC CPU Benchmark	CFI Solution									Benchmark Correlation
	MCFG	Reins	GCC-VTV	LLVM-CFI	MCFI	$\pi$ CFI	$\pi$ CFI (nto)	PathArmor	Lockdown	
perlbench				2.4	5.0	5.0	5.3	15.0	150.0	0.09
bzip2	-0.3	9.2		-0.7	1.0	1.0	0.8	0.0	8.0	-0.12
gcc					4.5	4.5	10.5	9.0	50.0	0.02
mcf	0.5	9.1		3.6	4.5	4.5	1.8	1.0	2.0	-0.39
gobmk	-0.2			0.2	7.0	7.5	11.8	0.0	43.0	-0.09
hmmmer	0.7			0.1	0.0	0.0	-0.1	1.0	3.0	0.33
sjeng	3.4			1.6	5.0	5.0	11.9	0.0	80.0	-0.03
h264ref	5.4			5.3	6.0	6.0	8.3	1.0	43.0	-0.09
libquantum				-6.9	0.0	-0.3	-1.0	3.0	5.0	0.51
omnetpp	3.8		5.8		5.0	5.0	18.8			-0.52
astar	0.1		3.6	0.9	3.5	4.0	2.9		17.0	0.92
xalancbmk	5.5		24.0	7.2	7.0	7.0	17.6		118.0	0.94
milc	2.0			0.2	2.0	2.0	1.4	4.0	8.0	0.40
namd	0.1		-0.1	0.1	-0.5	-0.5	-0.5	3.0		0.98
dealII	-0.1		0.7	7.9	4.5	4.5	4.4			-0.36
soplex	2.3		0.5	-0.3	-4.0	-4.0	0.9	12.0		0.89
povray	10.8		-0.6	8.9	10.0	10.5	17.4		90.0	0.88
lbm	4.2			-0.2	1.0	1.0	-0.5	0.0	2.0	-0.22
sphinx3	-0.1			-0.8	1.5	1.5	2.4	3.0	8.0	0.31
CONFIRM median	9.51	4.59	33.56	5.19	30.83	-11.10	-11.60	648.01	140.82	0.36

# CFI vs. Runtime Code Generation

16/31

- CFI Fundamental Assumptions (Abadi et al., 2005)
  - Non-Writable Code (NWC)
  - Non-Executable Data (NXD)
- Most Modern Software Violates Both
  - Rise of Just-In-Time (JIT) Languages since 2005
    - Lua, JavaScript, Python, Java, Ruby, PHP, Erlang, Wasm, Lisp, Ethereum, ...
    - Everything on .NET, all Microsoft COM software, ...
  - All self-unpacking components (e.g., cloud), installers (e.g., UPX), ...
  - Many forms of dynamic loading & hooking (example: Microsoft Office)



# Existing Solutions

17/31

- Manually customize the code generator (RockJIT [Niu & Tan, CCS'14])
  - ▣ Extremely high maintenance burden
  - ▣ Only works for certain very specific code generation patterns (old JITs)
  - ▣ Incompatible with all modern software (**~9 years of 100% incompatibility**)
- Turn off all dynamic code generation
  - ▣ Massive performance hit (**e.g., 1600% overhead on JS Octane**)
  - ▣ Impossible for many products (.NET)
  - ▣ Introduces new compatibility & security problems

# Real-world Example: Edge Browser

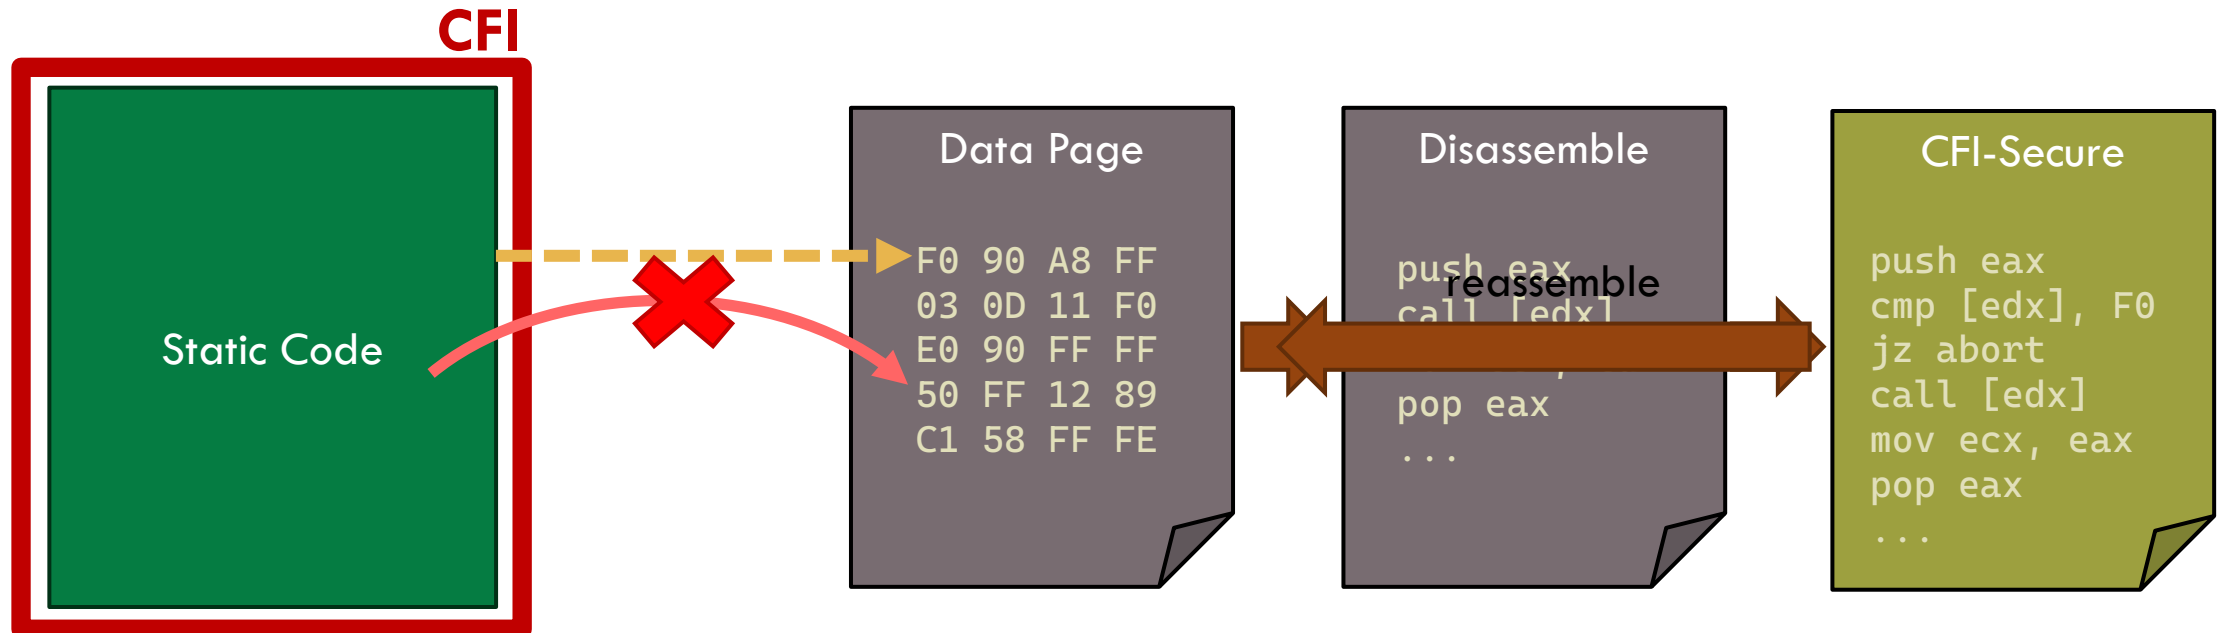
18/31

- Browser devs begin to realize that JIT is a huge security risk:
  - Over half of in-the-wild Chrome exploits from 2018-2021 abuse JIT vulnerabilities. [Mozilla Research, 2021]
  - More than 70% of the top programming languages are JIT-compiled, including JavaScript in almost all browsers.
- Microsoft turns off the JIT completely to enforce CFI in “secure mode”
  - *“As of Microsoft Edge 98, Control-flow Enforcement Technology (CET) and Arbitrary Code Guard (ACG) will be enabled in the renderer process when a site is in enhanced security mode. These additional mitigations **prevent dynamic code generation** in the renderer processes ...”* [Microsoft Browser Vulnerability Research Lab, 2022]
- But JIT compilers are dynamic code generators (critical for performance)...

# RENEW: Rewriting Newly Executable pages after Writes

19/31

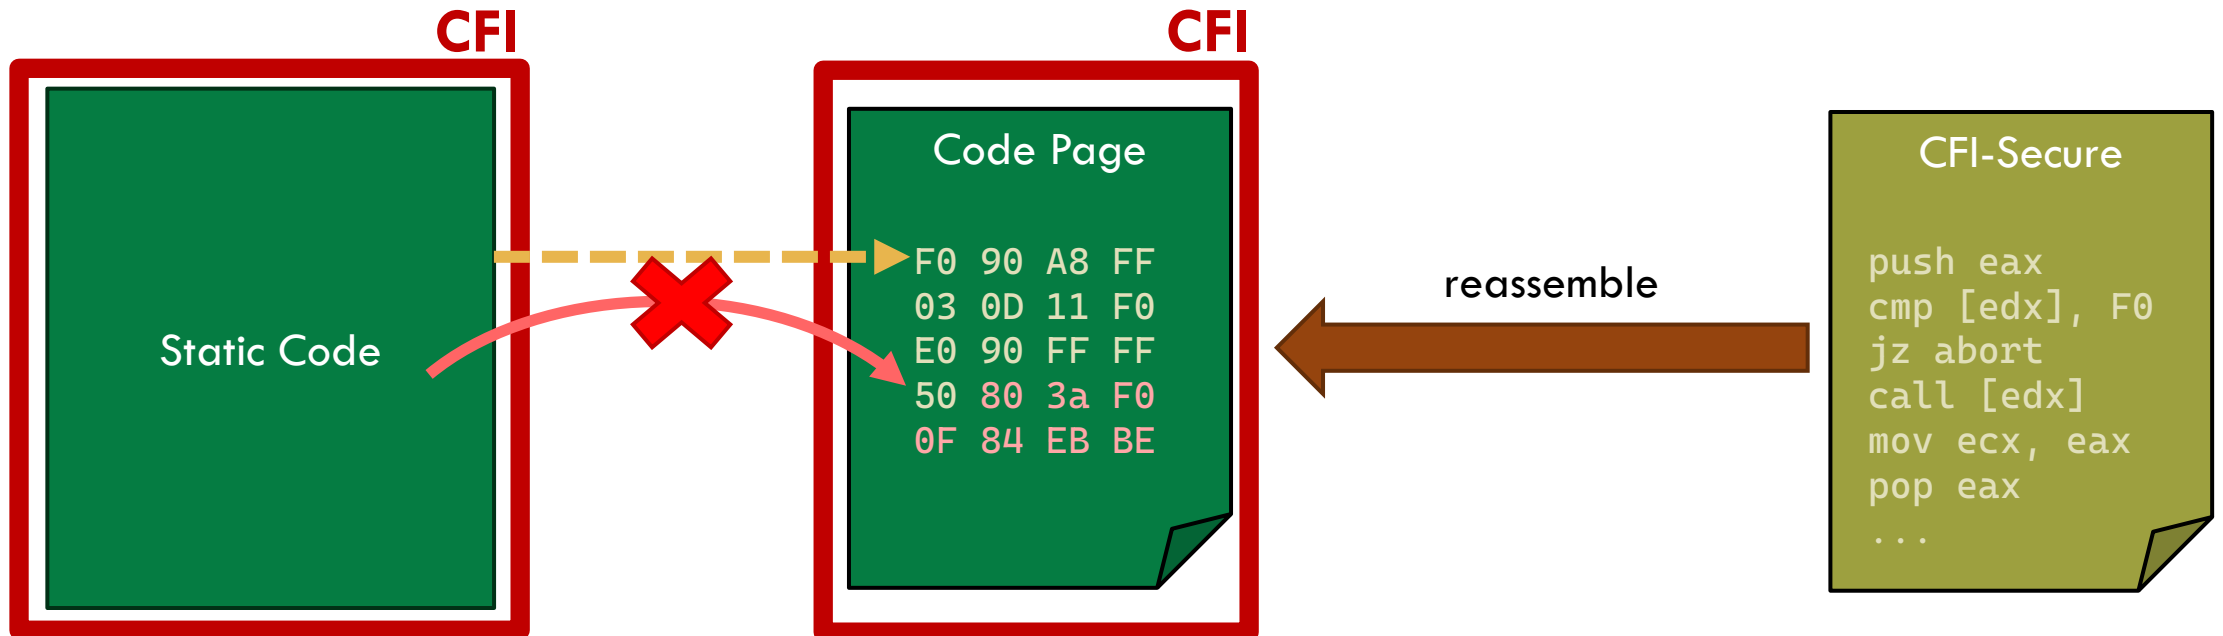
E. Bauman, J. Duan, K.W. Hamlen, and Z. Lin, “**Renewable Just-In-Time Control-Flow Integrity**,” In *Proc. 26<sup>th</sup> Int. Sym. on Research in Attacks, Intrusions and Defenses (RAID)*, October 2023.



# RENEW: Rewriting Newly Executable pages after Writes

20/31

## Static+Dynamic Code Instrumentation Dynamic Policy Enforcement



# Challenges

21/31

- Disassembly + Reassembly must be *fast* and *secure*
  - ▣ Disassembly alone is provably undecidable in general!
- Must support *recursive (generational) dynamic code generation*
  - ▣ dynamic code may write new dynamic code
  - ▣ static code may edit dynamic code pages mid-execution
- Must support calls from dynamic code to static code
  - ▣ Static code pointer passed to dynamic code
  - ▣ Pointer might not target a CFI-sanctioned entry point!
- Real-world apps sometimes *read generated code as data* (ugh!)

# Disassembly Undecidability

22/31

```
FF E0 5B 5D C3 0F
88 52 0F 84 EC 8B
```

□ Disassemble this hex sequence

■ CISC disassembly is undecidable! [Fred Cohen, '86]

Valid Disassembly	
FF E0	jmp eax
5B	pop ebx
5D	pop ebp
C3	retn
0F 88 52 0F 84 EC	jcc
8B ...	mov

Valid Disassembly	
FF E0	jmp eax
5B	pop ebx
5D	pop ebp
C3	retn
0F	db (1)
88 52 0F 84 EC	mov
8B ...	mov

Valid Disassembly	
FF E0	jmp eax
5B	pop ebx
5D	pop ebp
C3	retn
0F 88	db (2)
52	push edx
0F 84 EC 8B ...	jcc

# Innovation: Superset Disassembly

Byte Sequence: FF E0 5B 5D C3 0F 88 B0 50 FF FF 8B

23/31

● Disassembled

✘ Invalid

	Hex
●	FF
●	E0
●	5B
●	5D
●	C3
●	0F
●	88
✘	B0
	50
✘	FF
	FF
●	8B

Included Disassembly
jmp eax
pop
L1: pop
retn
jcc
L2: mov
loopne
jmp L1
mov
jmp L2

# Machine learning-based Disassembly Pruning

[Wartell, Zhou, Hamlen, Kantarcioglu, PAKDD'14]

24/31

- Insight: Distinguishing real code bytes from data bytes is a “noisy word segmentation problem”.
  - Word segmentation: Given a stream of symbols, partition them into words that are contextually sensible. [Teahan, 2000]
  - Noisy word segmentation: Some symbols are noise (data).
- Machine Learning based disassembler
  - based on  $k$ th-order Markov model
  - Estimate the probability of the sequence  $B$ :

$$p(B|M_\alpha) = -\log \prod_{i=1}^{|B|} p(b_i | b_{i-k}^{i-1}, M_\alpha)$$

Wartell, Zhou, Hamlen, Kantarcioglu. “Shingled Graph Disassembly: Finding the Undecidable Path.” PAKDD 2014.

Wartell, Zhou, Hamlen, Kantarcioglu, and Thuraisingham. “Differentiating code from data in x86 binaries.” ECML/PKDD 2011.



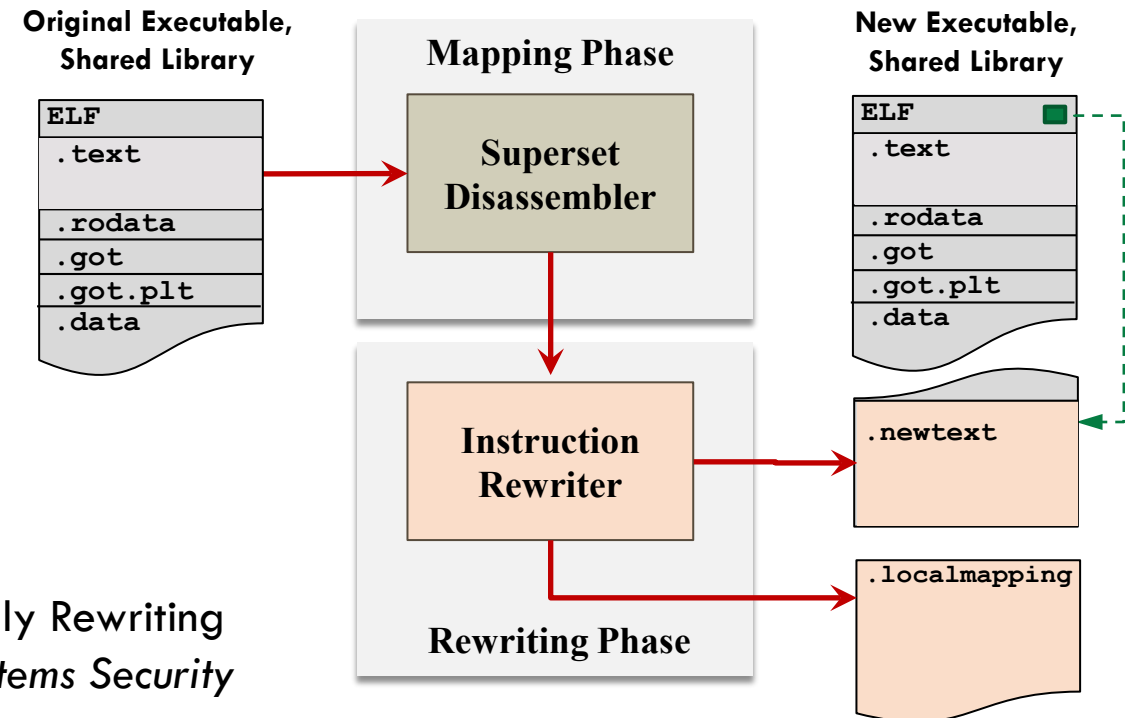
# Multiverse “Superset” Disassembler

[Bauman, Lin & Hamlen, NDSS’18]

25/31

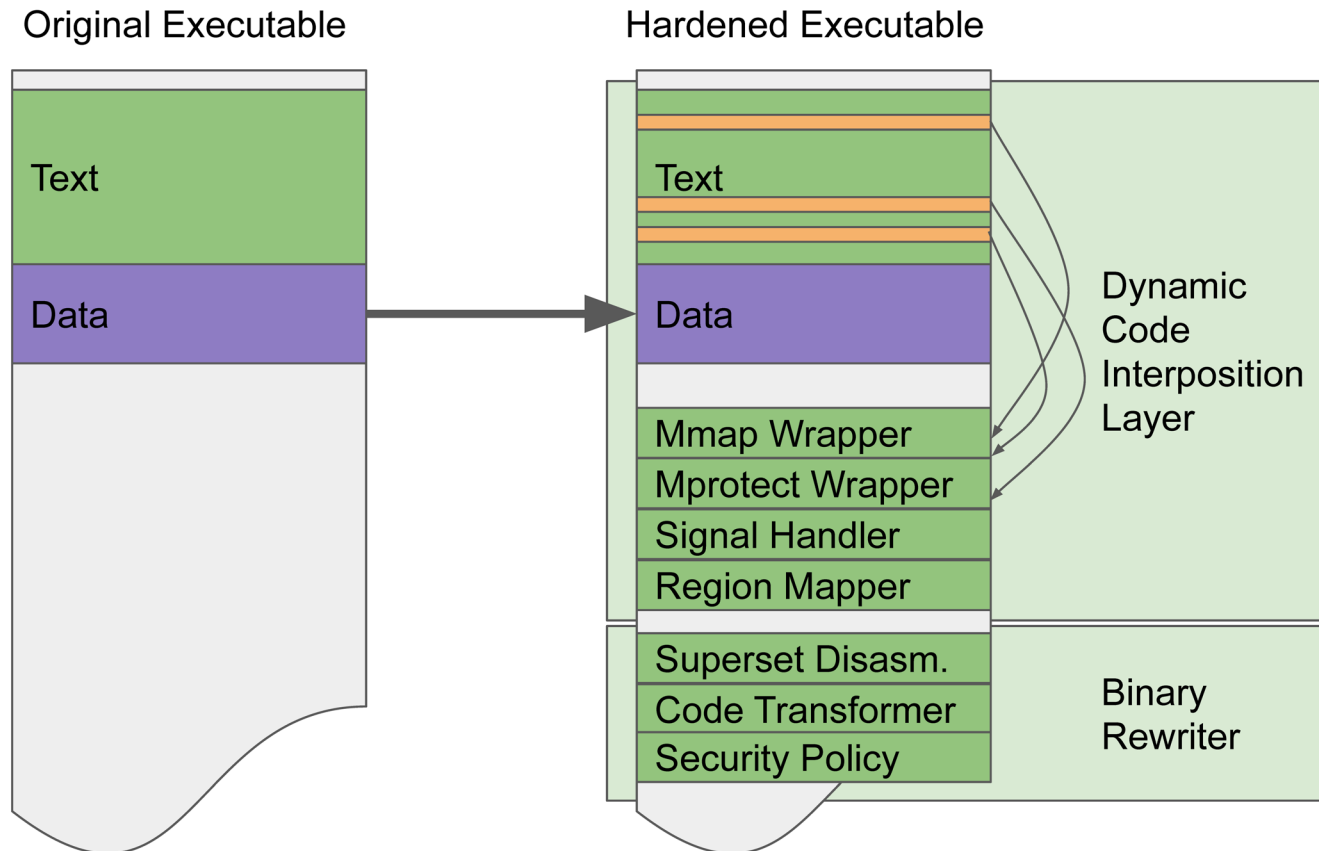
- ❑ Conservatively include every possible disassembly.
- ❑ Include and secure all of them.
- ❑ tested on 126 apps + 77 libs (all source-free)
- ❑ all application functionalities preserved
- ❑ 4-5x size increase of code segments (much smaller impact on overall file size)

E. Bauman, Z. Lin, and K.W. Hamlen. “Superset Disassembly: Statically Rewriting x86 Binaries Without Heuristics”. In *Proc. Network & Distributed Systems Security (NDSS)*, 2018.



# RENEW Overview

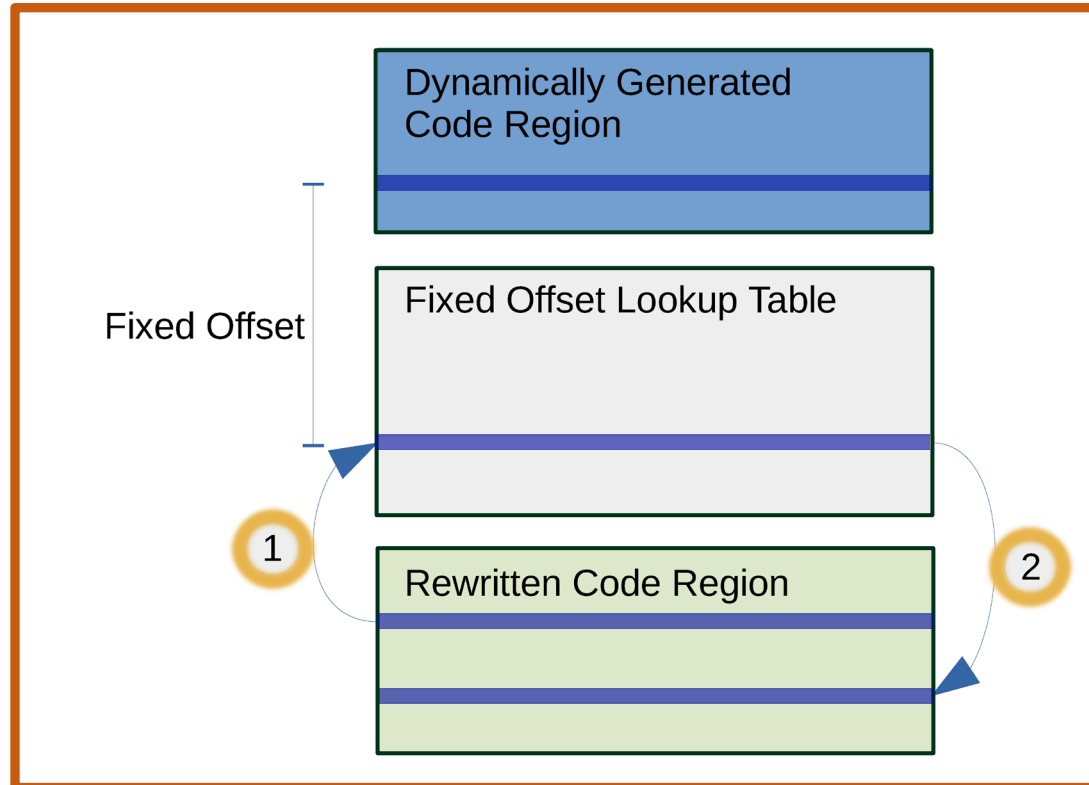
26/31



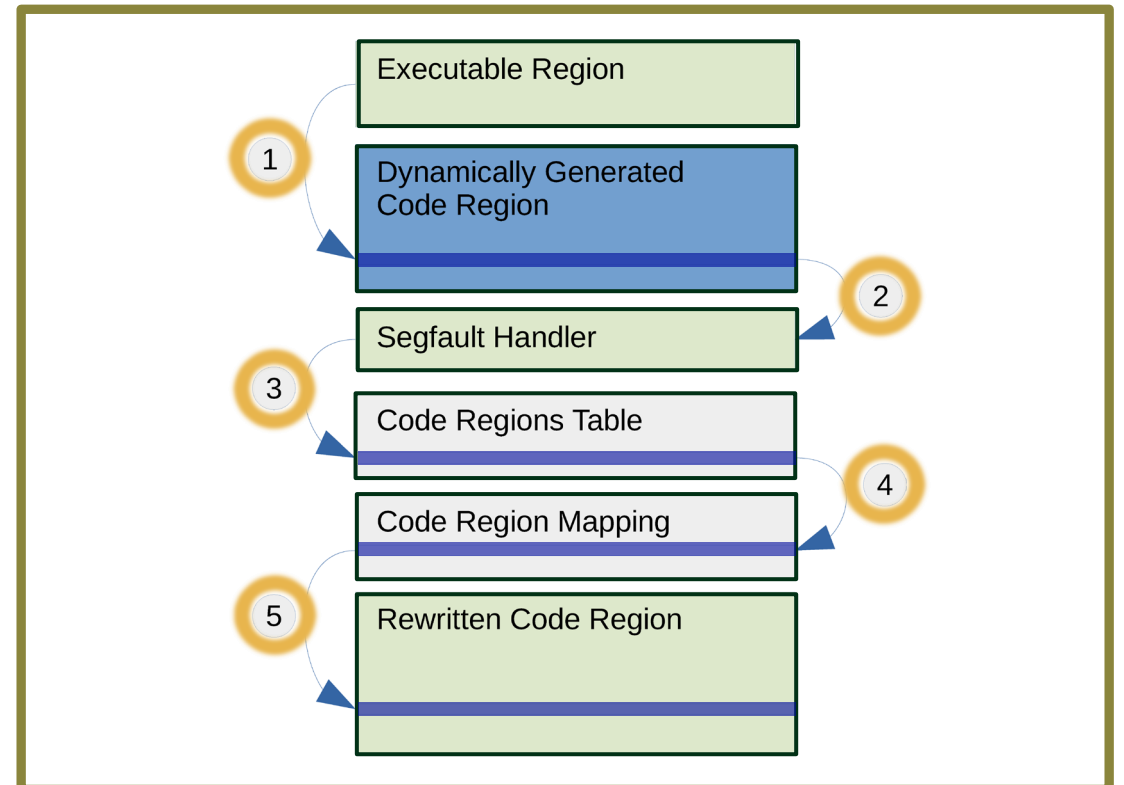
# Optimization Strategy: Fast path + Slow path

27/31

## Fast Path: Statically predictable flows



## Slow path: Fall back to signal handlers



# Proof-of-Concept Implementation

28/31

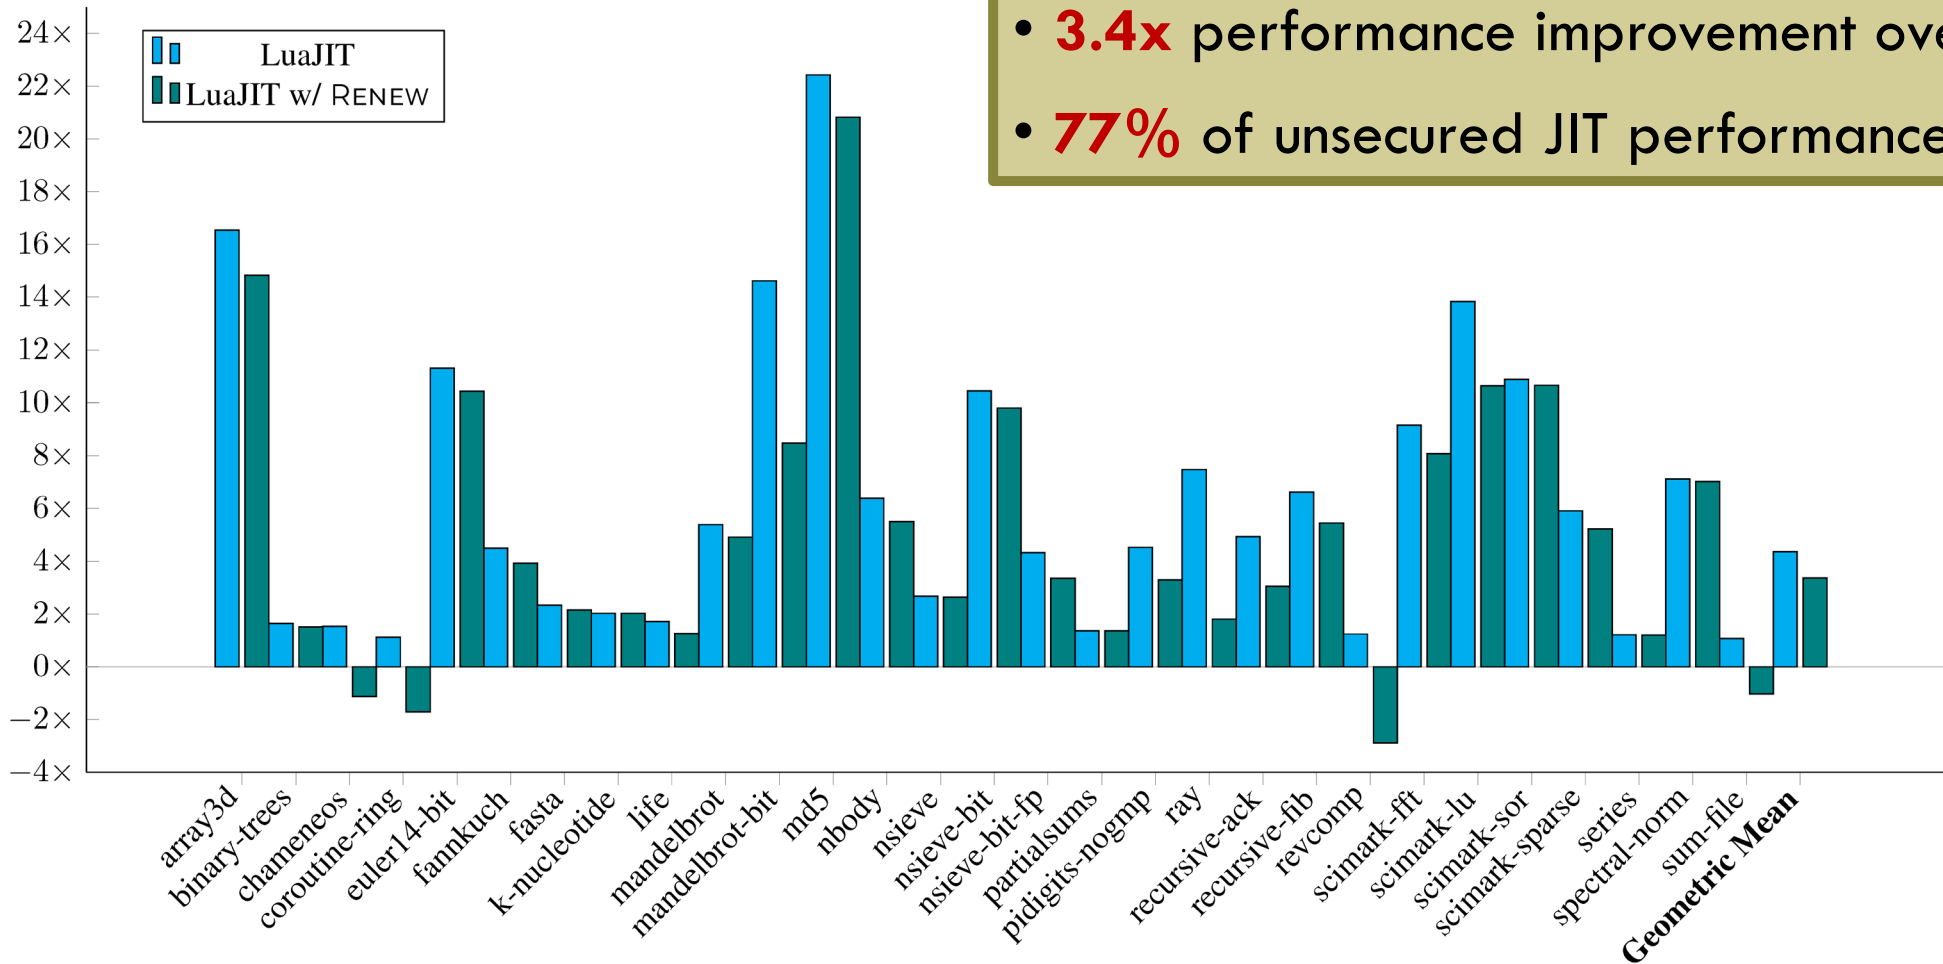
- Compact design and implementation
  - ~2000 lines of C code, plus in-lined assembly
  - includes entire disassembler, rewriter, interposition layer, etc.
  - included into target applications, so must be small
- Injected into target applications during compilation
  - shared library (`-Wl,-wrap=mmap -Wl,-wrap=mprotect`)
  - one-line change to application main function to call Renew initializer
- CFI static instrumentation of main app assumed
  - static CFI policy must permit Renew's static flows
  - Renew handles the dynamic flows
- Dramatically easier process than manual JIT redesign!

# Evaluation

29/31


- Three main case-studies:
  - ▣ Lua (JIT compiler)
  - ▣ Firefox JavaScript (Spidermonkey JIT compiler)
  - ▣ UPX (installer / code unpacker)
- Completely different rewriting strategies
  - ▣ No common code generation patterns
  - ▣ Completely different control-flow patterns
  - ▣ Highly optimized, highly complex, high churn (most popular JITs and unpacker today, latest versions at time of implementation)
- No existing CFI solution works correctly on any of these target applications.

# LuaJIT Evaluation Results



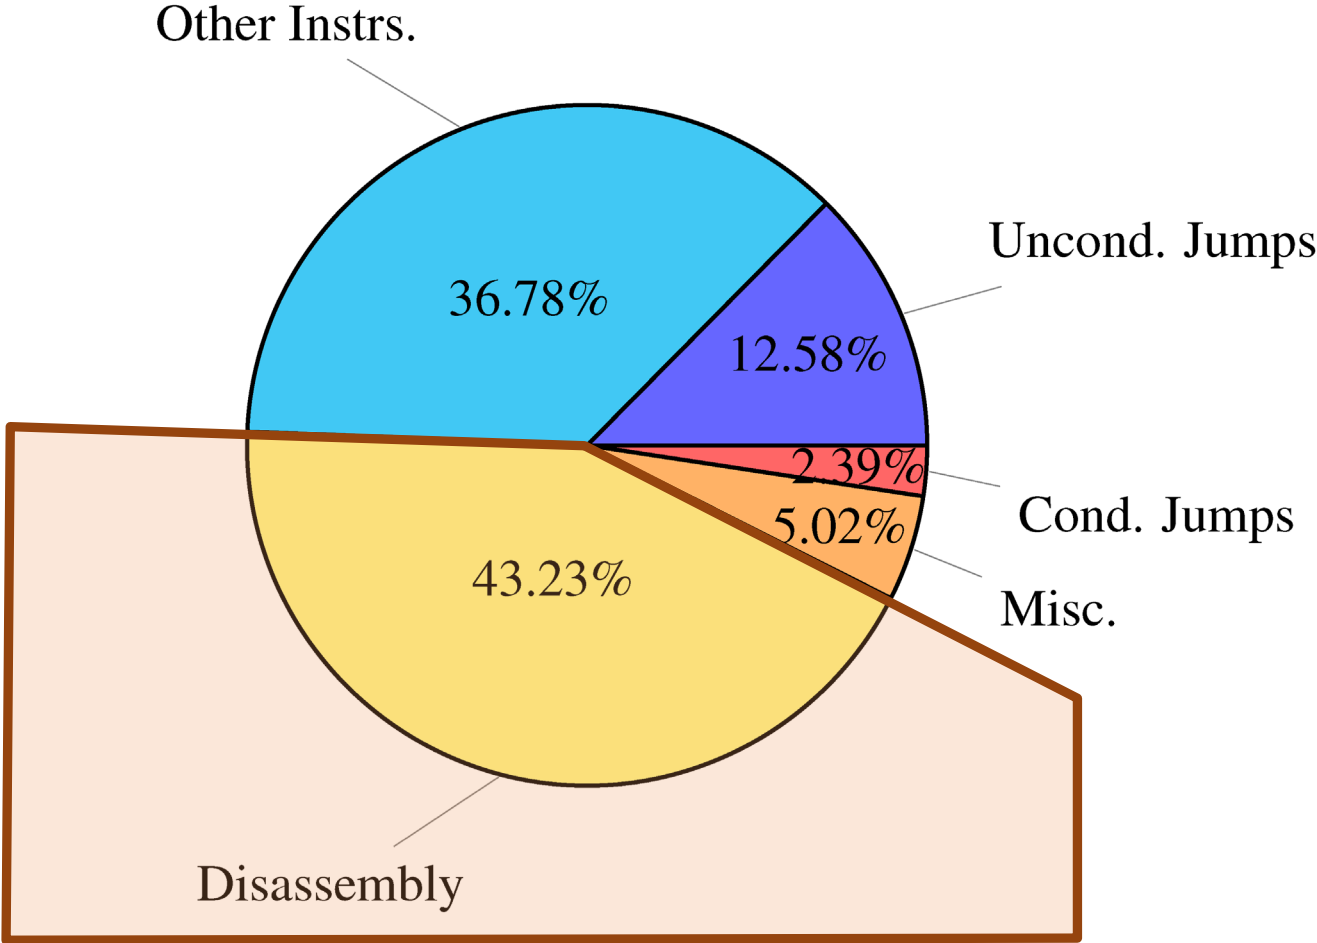
- **3.4x** performance improvement over JIT-off
- **77%** of unsecured JIT performance preserved

# Firefox JS Evaluation



- **3.9x** performance improvement over JIT-off
- **23%** of unsecured JIT performance preserved

# Overhead Breakdown





# UPX Evaluation

33/31

- SPEC CPU benchmarks + GNU binutils apps packed using UPX
- Extremely difficult compatibility challenge
  - ▣ UPX uses a custom binary header format and IAT to save space
  - ▣ completely arbitrary code generation behavior (depends on packed code)
  - ▣ two highly compressed layers of unpacking
  - ▣ first layer defies conventional disassembly
- Results
  - ▣ **all tests worked out-of-the-box** (no changes to Renew required)
  - ▣ predictably high overhead (3.6x slowdown)
    - not really a fair performance test; we mainly wanted to test compatibility

# Related Work Comparison

34/31

System	Year	SFI	CFI	Code-reuse			Source-agnostic
				Immunity	JIT	Packers	
NaCl-JIT	2011	✓	✗	✓	✓	✗	✗
Librando	2013	✗	✗	⚠	✓	✗	✓
RockJIT	2014	✓	✓	✓	✓	✗	✗
SDCG	2015	✗	✗	✗	✓	✗	✗
JITScope	2015	✓	✓	✓	✓	✗	✗
JITGuard	2017	✗	✗	⚠	✓	✗	✗
RENEW	2022	✓	✓	✓	✓	✓	✓

⚠ = Defense is diversity-based, so can be compromised by information disclosure.



THANK YOU!



DR. KEVIN HAMLIN

LOUIS A. BEECHERL, JR. DISTINGUISHED PROFESSOR  
COMPUTER SCIENCE DEPARTMENT  
CYBER SECURITY RESEARCH AND EDUCATION INSTITUTE  
THE UNIVERSITY OF TEXAS AT DALLAS

# SOFTWARE ATTACK SURFACE REDUCTION ON THE FLY

ONR Award N00014-21-1-2654

Any opinions, findings, conclusions, or recommendations expressed in this presentation are those of the author(s) and do not necessarily reflect the views of ONR, UTD or other supporters.