

# FRRACTAL

An Operating System Designed for Microarchitecture Reverse Engineering


Joseph Ravichandran & Mengjia Yan  
MIT CSAIL

# Research Question

Can we improve microarchitecture security research by creating a new operating system?

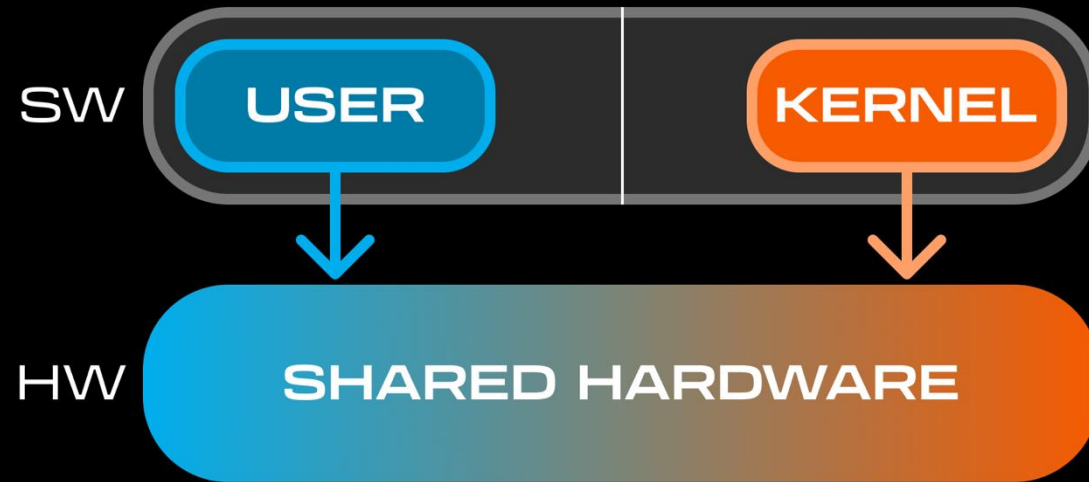


Linux/ macOS




Fractal


# ISA separates user & kernel code



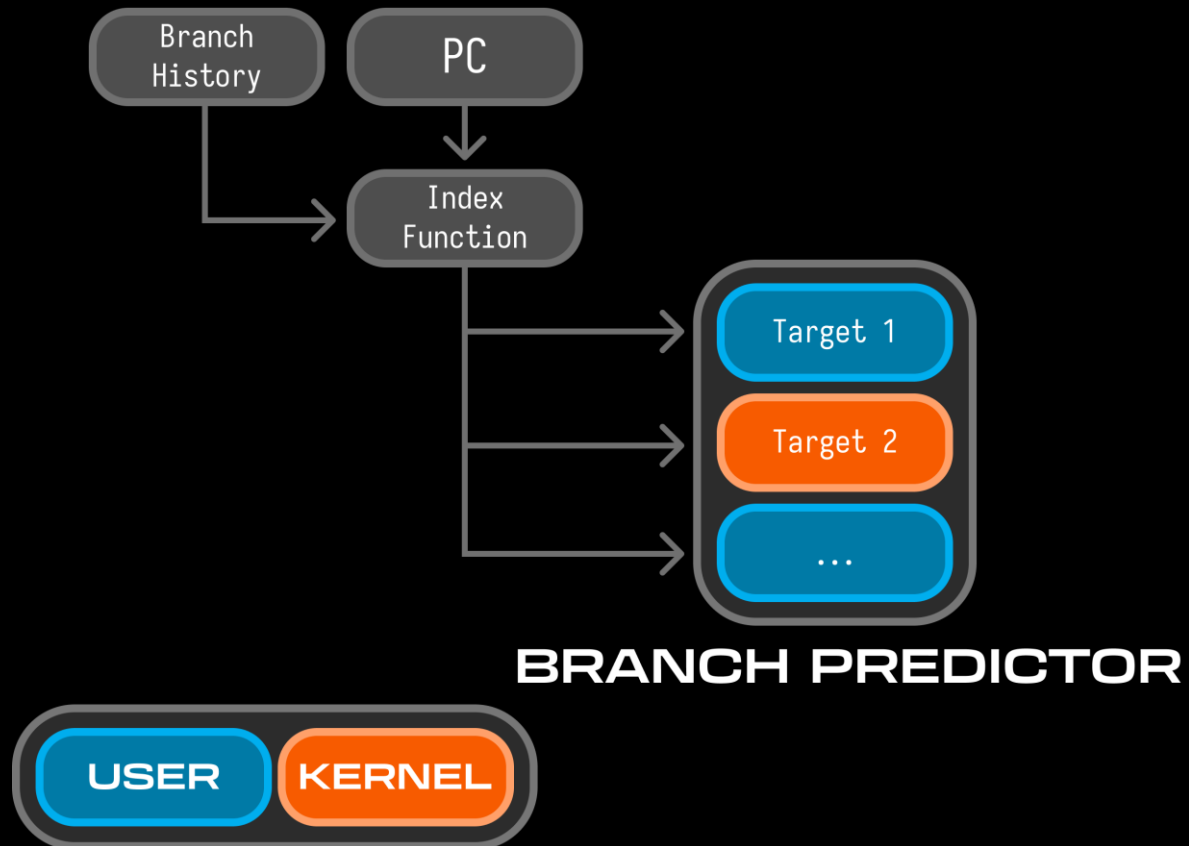
# Recent defenses separate hardware



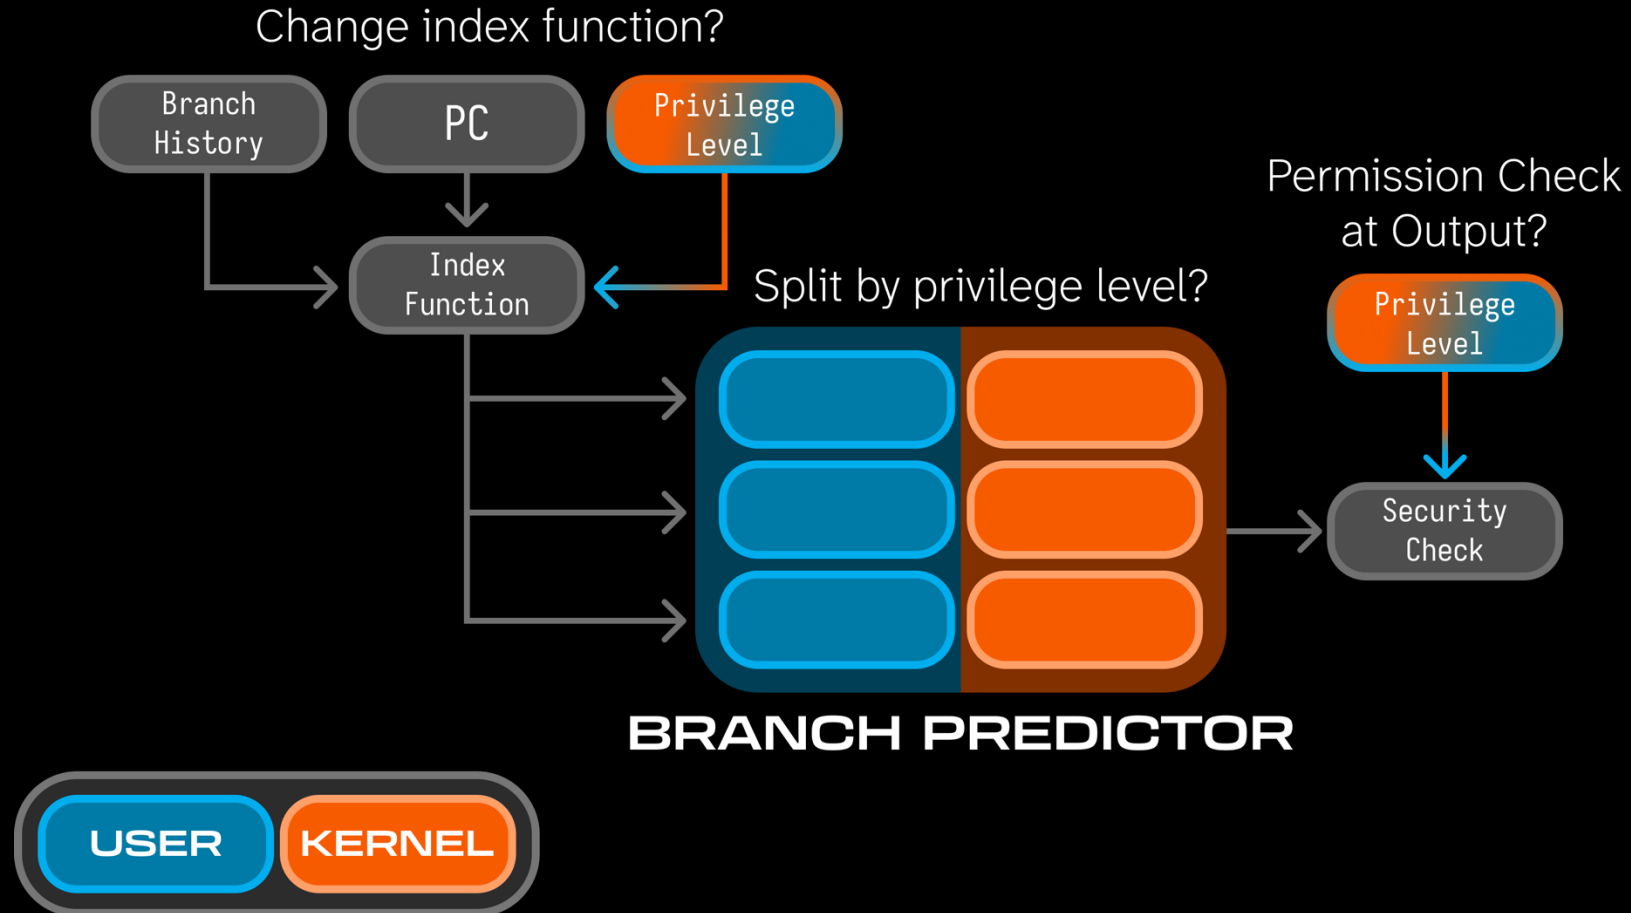
# Fractal blends software modes to better understand hardware



# Shared entries can lead to security issues!



# Multiple Possible Designs: We want to reverse engineer



## Approach: Build attacks to study defenses

1. Create a simple single-domain attack
2. Vary the privilege level
3. Learn how hardware prevents our attack across domains



# Attacking the indirect branch predictor

Train branch predictor to jump to target X

Call kernel extension

Run aliased branch

Measure whether trained branch influenced aliased branch

Userspace

Kernelspace



# Attacking the indirect branch predictor

Train branch predictor to jump to target X

Call kernel extension

Measure whether trained branch influenced aliased branch

**Different address spaces:  
Changing privilege level  
changes VA layout**

Run aliased branch

**System noise from  
syscall overhead**

Userspace

Kernelspace

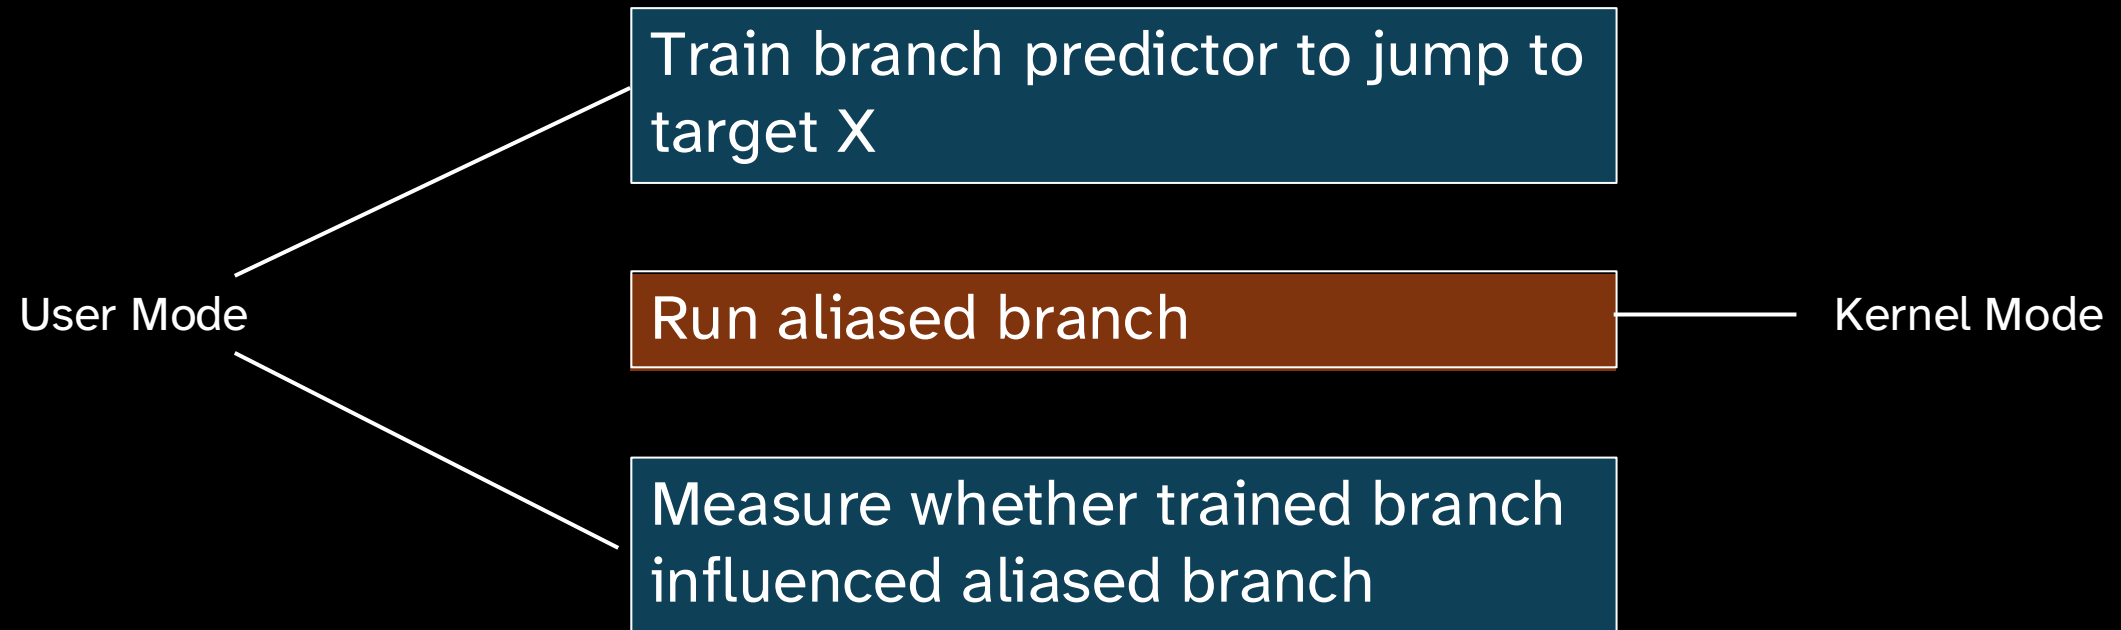


## Solution


Make user and kernel share an  
address space & code



# Fractal lets you change the privilege level of specific parts of your program



# How do we change a running program's privilege level?



Duplicate the page table with different permissions



# The Full Picture



# Vary Just the Privilege Level

User → Kernel



Kernel → User



User → User



Kernel → Kernel



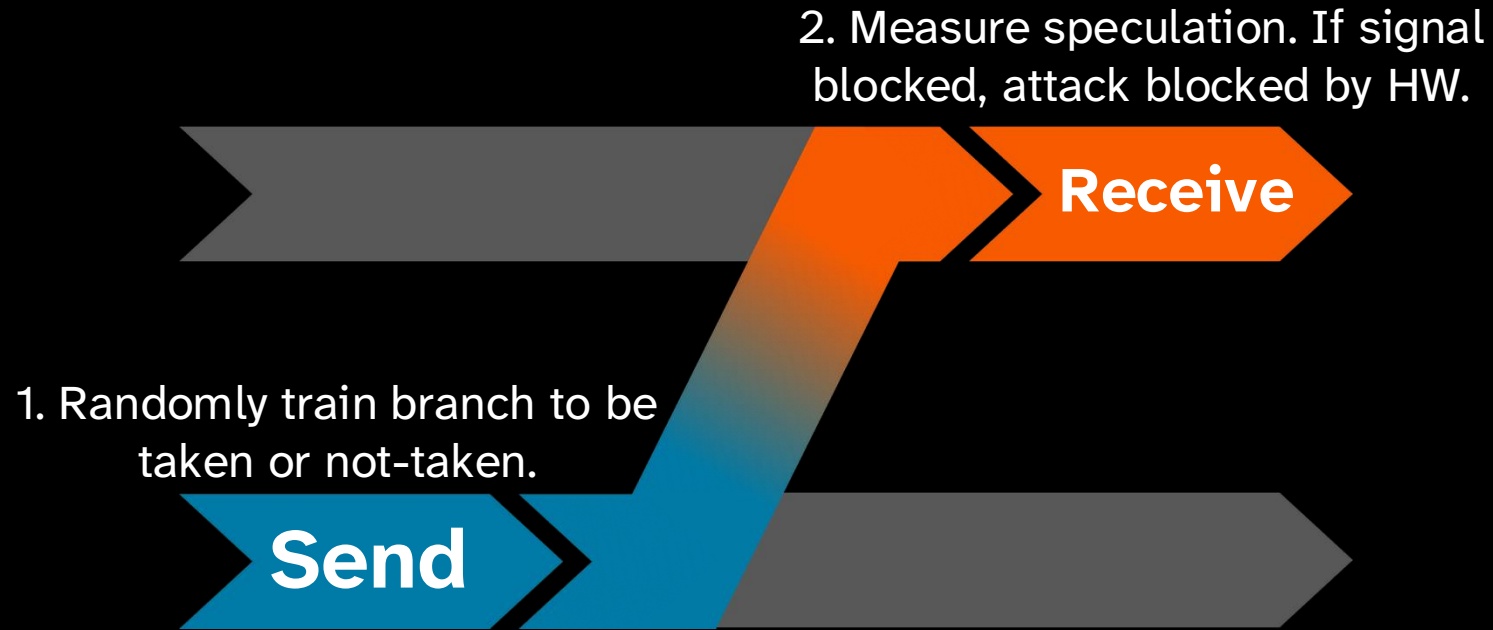
USER

KERNEL

PAGETABLE  
SWITCH



# Sender & Receiver



USER

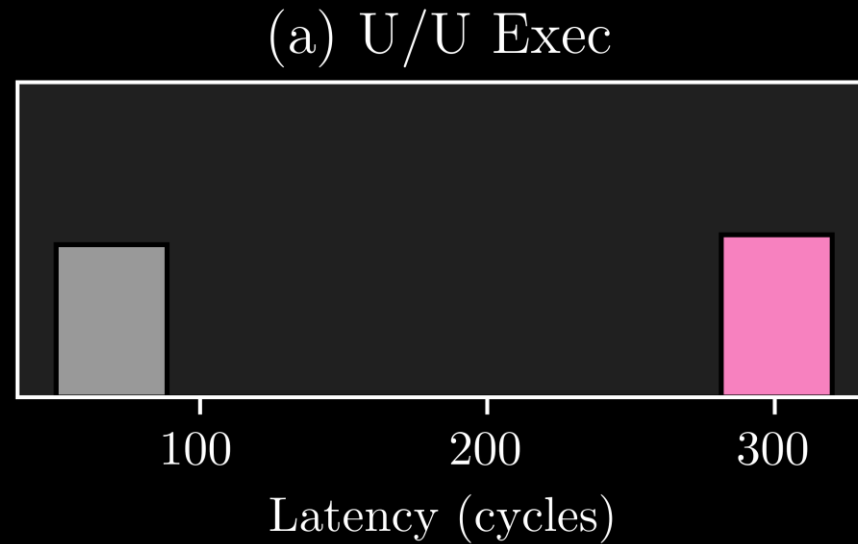
KERNEL

PAGETABLE  
SWITCH




# Evaluation on Apple M1

When running same-privilege receiver, hardware uses branch targets trained by sender

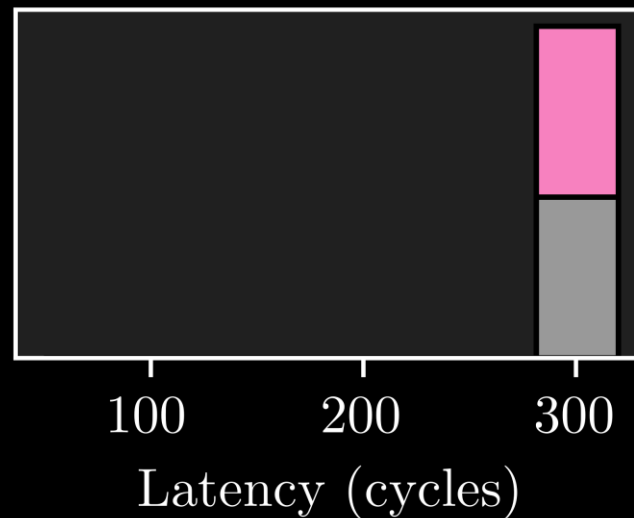


# Evaluation on Apple M1


(a) U/U Exec



(b) U/K Exec

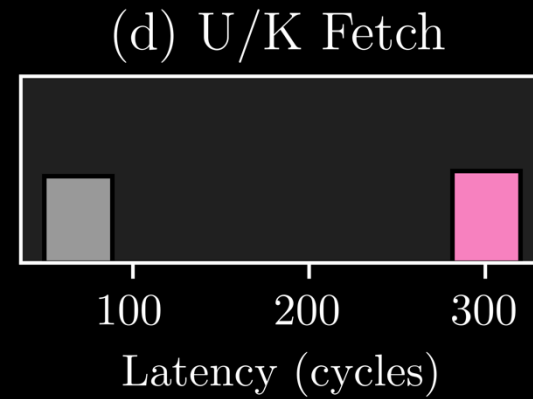


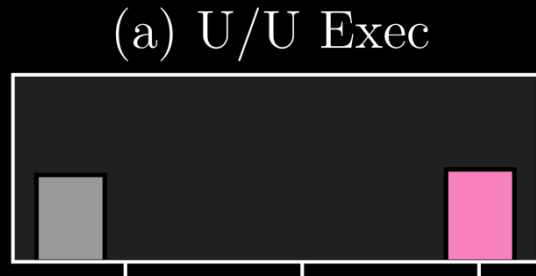


When receiver is moved to kernel mode, hardware stops using target for predictions!



# Evaluation on Apple M1

Surprisingly, hardware still **fetches** trained targets no matter the privilege mode!



Surprising!



# Evaluation on Apple M1

- CBP has no protection
- Cross-ASID acts like cross-privilege on IBP

	Same Domain Fetch	Same Domain Exec	X-Priv Fetch	X-Priv Exec	X-ASID Fetch	X-ASID Exec
CBP	✓	✓	✓	✓	✓	✓
IBP	✓	✓	✓	✗	✓	✗

✓ = Attack Succeeds  
✗ = Attack Blocked by Hardware





# Evaluation on Apple M1

- Phantom speculation: replace receiver branch with NOP, check whether it gets speculatively decoded as a branch
- Phantom fetches exist, but Phantom speculative execution blocked

	Same Domain Fetch	Same Domain Exec	X-Priv Fetch	X-Priv Exec	X-ASID Fetch	X-ASID Exec
CBP	✓	✗	✓	✗	✓	✗
IBP	✓	✗	✓	✗	✓	✗

✓ = Attack Succeeds  
✗ = Attack Blocked by Hardware





<https://fractal-os.com>

By creating new system software,  
microarchitecture reverse engineering  
workflows can be improved.

More in the paper:

CPU stack aliasing, syscall returns, shared C library,  
shadow memory maps, memory management  
approaches, and more!



This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. (2141064). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.