

# PACStack

## an Authenticated Call Stack

*Hans Liljestrand<sup>†</sup>, Thomas Nyman<sup>‡</sup>, Lachlan J. Gunn<sup>‡</sup>, Jan-Erik Ekberg<sup>§‡</sup>, N. Asokan<sup>†‡</sup>*

# Control-flow attacks

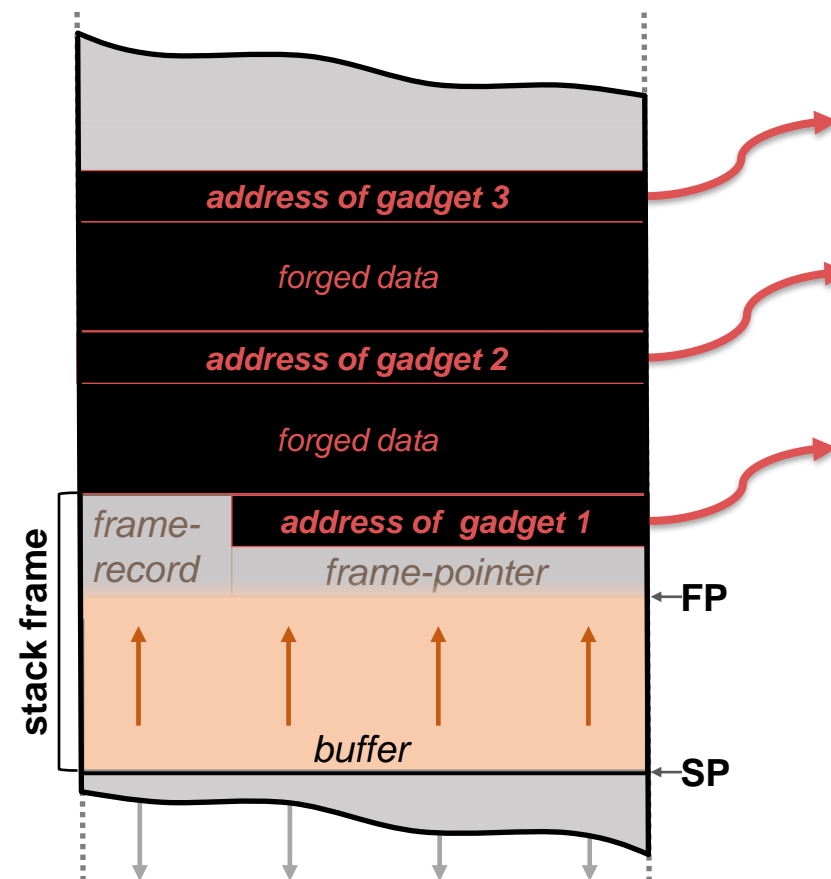
Exploit memory error (e.g. buffer overflow) to:

- Inject [shellcode](#) into writable memory (usually stack)
- Corrupts [return address](#) to redirect execution flow

Code injection is prevented by modern defenses

**Return address corruption** still prevalent!

- Code reuse attacks, e.g., ROP



Elias Levy (as *Aleph One*), [Smashing the stack for fun and profit](#), Phrack 7 (1996)

Cowan et al., [StackGuard: Automatic adaptive detection and prevention of buffer-overflow attacks](#), USENIX Security (1998)

Szekeres et al., [SoK: Eternal War in Memory](#), IEEE SP (2013)

# ARMv8.3-A Pointer Authentication

## General purpose hardware primitive **approximating pointer integrity**

- Ensure **pointers** in memory remain **unchanged**
- Uses message authentication codes embedded in pointers

## Introduced in ARMv8.3-A specification (2016)

- First compatible processors 2018 (Apple A12 / [iOS12](#))
- Userspace support in [Linux 4.21](#), in-kernel support in [5.7](#)
- Instrumentation support introduced in [GCC 7.0](#) and [Clang 8](#)

Apple Inc. "[Pointer Authentication – Clang documentation](#)", git (2019)

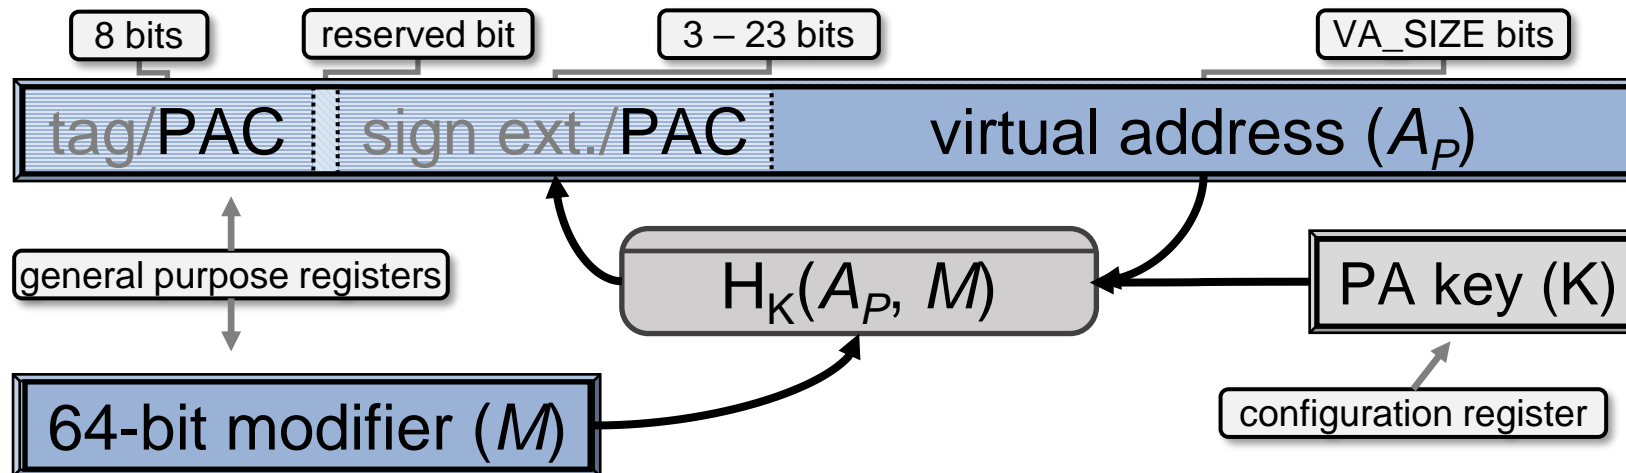
ARM. [Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile](#). Version E.a. July 2019

ARM. [Developments in the Arm A-Profile Architecture: Armv8.6-A](#). September 2019

# ARMv8.3-A PA – PAC Generation

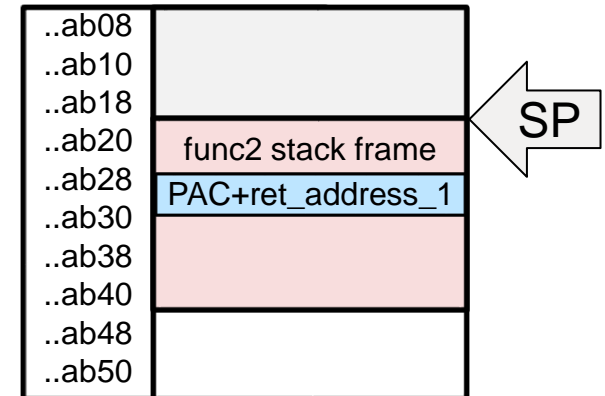
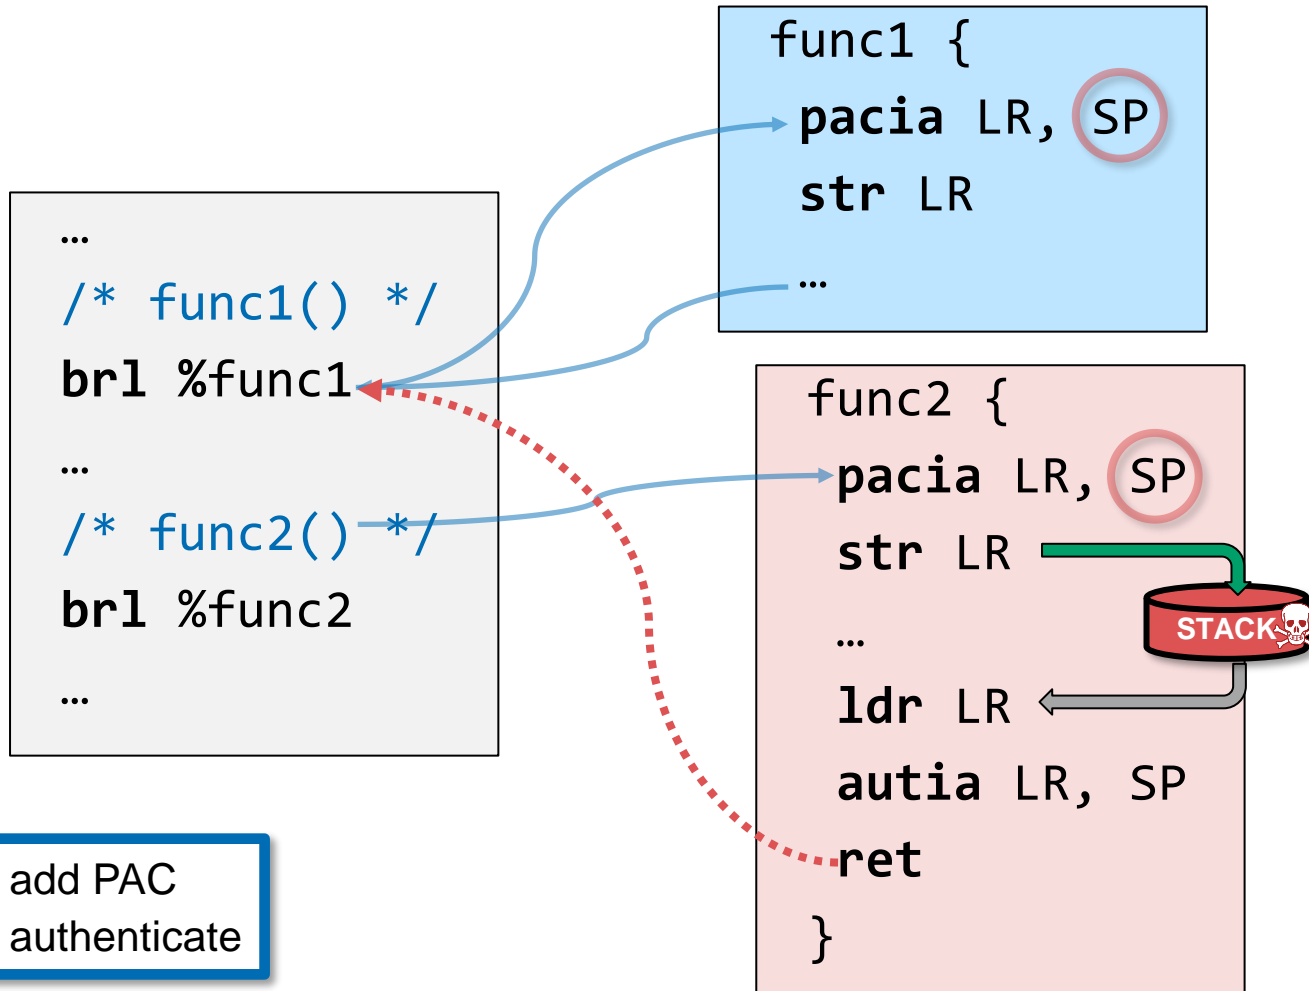
Adds Pointer Authentication Code (**PAC**) into unused bits of pointer

- Keyed, tweakable **MAC** from **pointer address** and 64-bit **modifier**
- PA keys protected by hardware, modifier is decided where pointer **created and used**



# Reuse attacks on PA


Adversary may **reuse** PACs



# PACStack: high-level idea

## Chained MAC of cryptographically bound return addresses

- Modifier bound to all previous return addresses on the call stack
- Statistically unique to control-flow path
  - prevents reuse
  - allows precise verification of returns



$auth_i$ ,  $i \in [0, n - 1]$  bound to corresponding return addresses,  $ret_i$ ,  $i \in [0, n]$ , and  $auth_n$

# Collisions still happen!

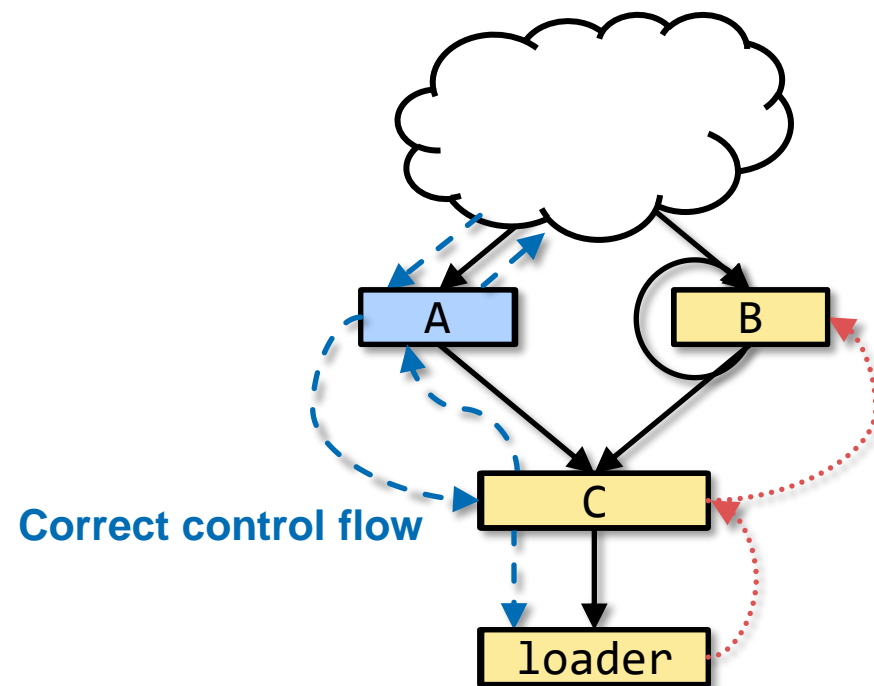
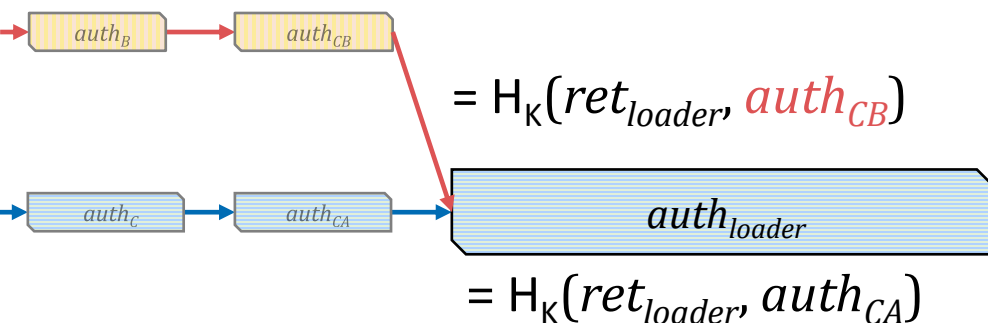
PAC modifiers meant to **prevent** pointer reuse

But **collisions allow reuse**

- Registers cannot be directly corrupted
- But values loaded from the stack are vulnerable

Collisions allow **reuse of PAC chain**

- Harvest and reuse repeatedly for **arbitrary paths**



# Mitigation of hash-collisions: PAC masking

**Challenge:** PAC collisions occur on average after  $1.253 * 2^{b/2}$  return addresses

- For  $b = 16$ ,  $n = 321$  addresses

**Solution:** Prevent **recognizing** collisions by masking each *auth*

- pseudo-random mask generated using `pacia(0x0, authi-1)`
- exploitation changes from **deterministic** to **probabilistic**

Attack	w/o Masking	w/ Masking
Reuse previous auth collision	1	$2^{-b}$
Guess auth to existing call-site	$2^{-b}$	$2^{-b}$
Guess auth to arbitrary address	$2^{-2b}$	$2^{-2b}$

Maximum probability of success for different attacks



# Provable security

PA-based schemes can be analyzed as crypto protocols

We prove security reduction for PACStack

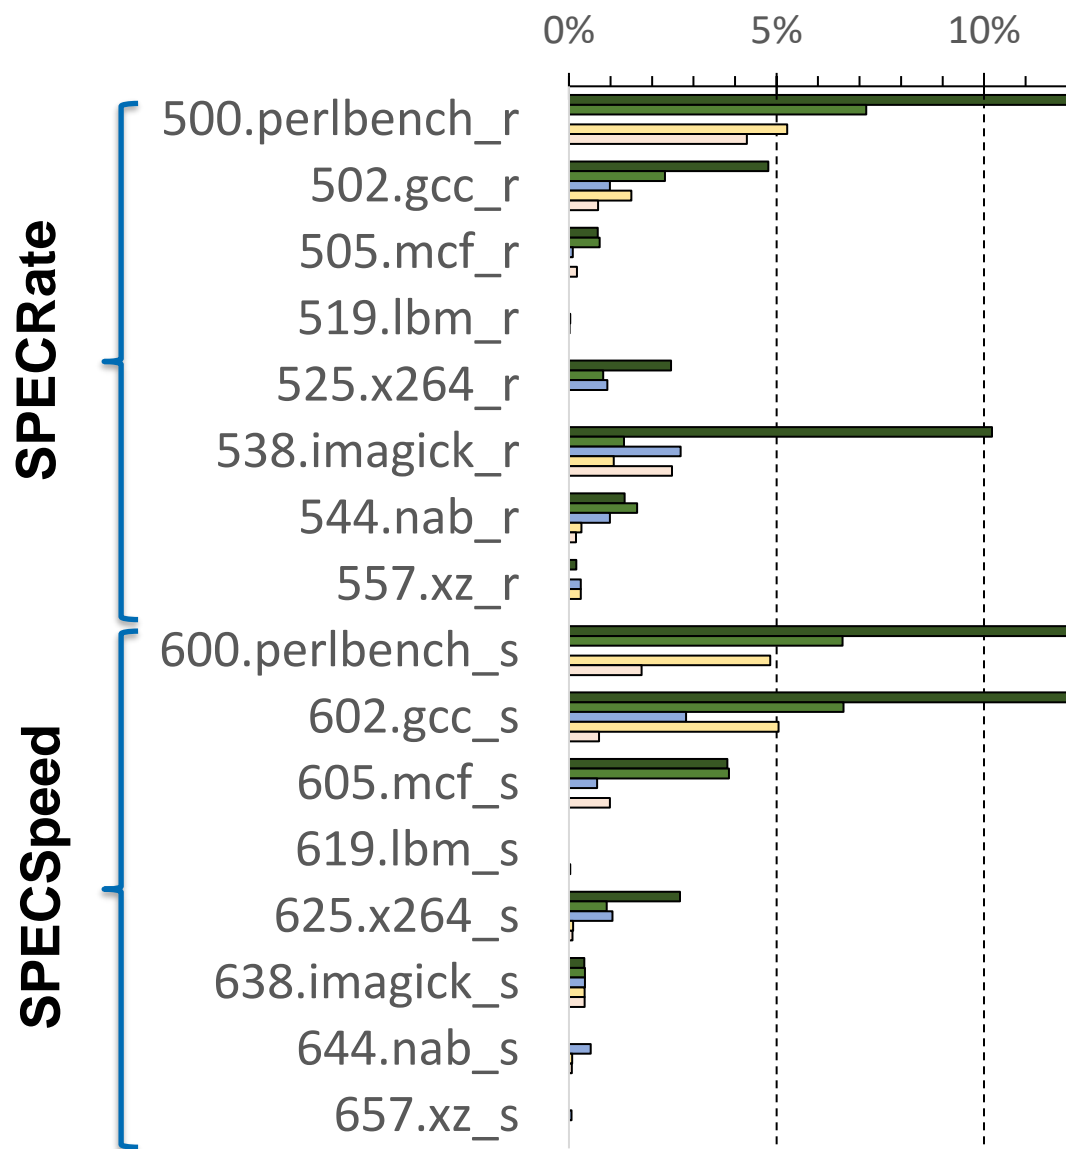
- to collision-resistance of PAC primitive

Formal analysis identified vulnerabilities in early design

- Identified collision attack against non-masked PACStack
- Proved masked PACStack secure

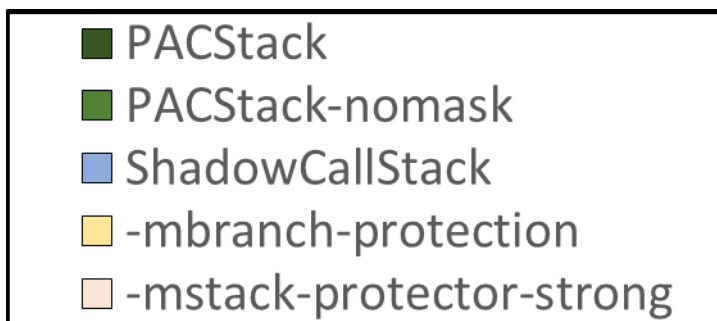
```
 $G_{ACS}^{\mathcal{A}}(1^\lambda, H, C, q)$ 
1:  $K \xleftarrow{\$} \{0,1\}^\lambda$ 
2: / Give  $\mathcal{A}$   $q$  tokens from call-graph traversals.
3: for  $i \in \{1, \dots, q\}$  do
4:    $p_{1..m+1} \leftarrow \mathcal{A}_{\text{oracle-request}}()$ 
5:   / Is the request for a real path through the call-graph?
6:   if  $\exists j : p_j \rightarrow p_{j+1} \notin \text{edges}(C)$  then
7:     return 0
8:   endif
9:    $auth_m \leftarrow T_K(p_m, T_K(p_{m-1}, \dots) \parallel p_{m-1}) \parallel p_m$ 
10:   $\mathcal{A}_{\text{oracle-response}}(auth_m)$ 
11: endifor
12:  $ptr_{\text{jumper}}, ptr_{\text{correct}}, auth_{\text{correct}}, t_{\text{correct}},$ 
13:    $ptr_{\text{adv}}, auth_{\text{adv}}, t_{\text{adv}} \leftarrow \mathcal{A}_{\text{ACS-Violation}}()$ 
14: / The substituted masked authenticated return address must be different.
15: if  $ptr_{\text{correct}} = ptr_{\text{adv}} \wedge auth_{\text{correct}} = auth_{\text{adv}}$  then
16:   return 0
17: endif
18: / Does the return pointer authenticate correctly with the adversary's
19: / new masked authenticated return address as the modifier?
20: if  $H_K(ptr_{\text{jumper}}, auth_{\text{correct}} \parallel ptr_{\text{correct}})$ 
21:    $\neq H_K(ptr_{\text{jumper}}, auth_{\text{adv}} \parallel ptr_{\text{adv}})$  then
22:   return 0
23: endif
24: / Did the adversary provide a valid masked authenticated return address?
25: if  $auth_{\text{adv}} = H_K(ptr_{\text{adv}}, t_{\text{adv}})$ 
26:   return 1
27: endif
28: return 0
```

# SPEC CPU 2017 benchmarks



**Performance overhead measured based on 4-cycle PA-analogue**

- **with masking ~ 3% (geo.mean)**
- **without masking ~ 1,2 % (geo.mean)**



# Thank you

**PACStack provides security comparable to shadow stacks**

- **but, probabilistic vs. deterministic**
- **Small overhead without additional hardware requirements**

**PA is a general-purpose mechanism**

- **Can support novel non-obvious use cases**
- **New emergent properties when combined with other HW?**



 [hans@liljestrand.dev](mailto:hans@liljestrand.dev)

 [@HansLiljestrand](https://twitter.com/HansLiljestrand)

 [pacstack.github.io](https://pacstack.github.io)

 [ssg.aalto.fi/research/projects/harp](https://ssg.aalto.fi/research/projects/harp)