

# An Analysis of Call-site Patching Without Strong Hardware Support for Self-Modifying-Code

Tim Hartley, <u>Foivos Zakkak</u>, Christos Kotselidis, Mikel Lujan

first.last@manchester.ac.uk

MPLR'19 2019-10-22



### Direct branching

### Indirect branching





## **Call-Site Patching**

- Tiered compilation
- De-optimization
- Etc.



# **JIT compilation and Caches**

The University of Manchester



Code-stream vs Data-stream

- 1. Code gets fetched to I-Cache
- 2. Data get fetched to D-Cache
- 3. CPU executes code from I-Cache
- 4. CPU writes data to D-Cache
- 5. D-Cache writes-back to memory
- 6. D-Cache fetches code to be edited
- 7. CPU writes code to D-Cache
- 8. D-Cache writes-back code



# Low-power architectures and call-site patching

The University of Manchester

### Fixed size instructions

- Limit the range of direct branches/calls
  - +- 128MiB on AArch64
  - +- 1MiB on RISC-V

### Require multiple instructions to perform long-range calls

#### AArch64





- Weak memory models and self-modifying-code (SMC) support
  - SW explicitly issues memory barriers
  - Code-stream handled separately from data-stream (need to sync them)
- Not all instructions are safe to patch
  - ARM (armv7 and armv8) and IBM (Power) limit the instructions that are safe to be patched while executing
    - Even if using atomic writes



# Patchable call-site implementations in AArch64

Direct Branching (short-range only)	Relative-Load Indirect Branching		
<b>B</b> TARGET	CALLEE_1 : .quad 0 x0123456789ABCDEF  CALLEE_N : .quad 0 x01234ABCDEF56789 START : LDR X16, CALLEE_1 BLR X16		
Absolute-Load Indirect Branching	Trampolines (OpenJDK approach)		
MOVZ X16, #0xABCD ; Craft the address MOVK X16, #0xEF89, lsl #16 ; holding MOVK X16, #0x7654, lsl #32 ; the MOVK X16, #0x0213, lsl #48 ; target LDR X16, [X16] BLR X16	L: LDR X16, CALLEE BR X16; Don 't link CALLEE: .quad 0 x0123456789ABCDEF START: BL SHORT_TARGET; or L		



# **Comparison of call-site implementation approaches**

		Absolute-load	<b>Relative-load</b>		
	Direct	Indirect	Indirect	Trampolines	Any   Safe-point
Characterisation	Code	Data	Hybrid	Hybrid	Any
<b>Requires SMC support</b>	Yes	No	No	Yes	Maybe
Size (in 32-bit instructions)	1	6	2 +2	1 +4 to 3 +2	Any $(\geq 1)$
Patching Complexity	Medium	Low	Medium	High	Any
Supported Call Range	Limited	Any	Any	Any	Any
Stop-the-world pause	No	No	No	No	Yes



- Odroid-C2
  - Quad-core Cortex-A53 @ 1.54GHz (pinned)
    - 8-stage pipelined processor with 2-way superscalar, in-order pipeline
  - 2 GB DDR3 RAM
  - Ubuntu 18.04.02 LTS
  - Kernel: Odroid 3.16..68-41
  - GCC 8.3.0
  - MaxineVM 2.8.0
  - OpenJDK 8 u212





## Microbenchmark

START:



- Generates inline call-sites
- Callers are ret-only methods
- To patch we call a patcher method instead of a ret-only
- Patcher always patches the next call-site (allows us to control number of patches
- Patcher performs the necessary barriers as it would in a real system



## **Microbenchmark results**

The University of Manchester



2019-10-22

MPLR'19 @foivoszakkak



# Dacapo and MaxineVM

- We take the best two performing approaches (Direct and Relative-Load Indirect) and evaluate them with DaCapo using MaxineVM
- We had to tweak Relative-Load Indirect to make it work with MaxineVM
  - Due to its metacircular nature, MaxineVM can only operate with offsets (relative branches), since at boot image creation the absolute targets are not known yet

Indirect-Maxine				
OFFSET: CALL:	ADR X17, CALL LDR X16, OFFSET ADD X16, X16, X17 B #8 .int CALL - CALLEE BLR X16	; Get address of BLR ; Load offset ; Add them ; Jump over inline offset _1		







### **DaCapo Results**





 OpenJDK's method seems the best for AArch64 since it penalizes only long-range branches and avoids explicit instruction cache invalidations on callers.

- If you have a higher  $\frac{\#long-range\ calls}{\#short-range\ calls}$  ratio then maybe Relative-Load is better

The most promising approach in theory would be combining the following gadgets
Indirect (long-rang)

Direct (short-range only) B TARGET ADRP X16, CALLEE ADD X16, X16, :lo12:CALLEE BLR X16

 On AArch64 this is not possible though since ADRP and ADD cannot be safely overwritten if they are being executed concurrently with the modifications.