

Monitoring program execution (and more!) on ARM processors



Toulouse
Hacking
Convention_

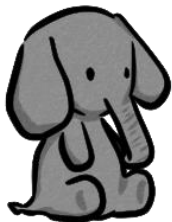
Pascal Cotret
pascal.cotret@gmail.com / @Pascal_r2

N7, March 9, 2018



Hello!

- Embedded software security engineer
 - Researcher in my spare-time
- (also former associate professor)



HardBlare project (3 labs, 2 PhDs...)

Threat model

Buffer overflow example with strcpy()

www.hackingtutorials.org

```
void main()
{
    char source[] = "username12"; // username12 to source[]
    char destination[7]; // Destination is 8 bytes
    strcpy(destination, source); // Copy source to destination

    return 0;
}
```

Buffer (8 bytes)

Overflow

U	S	E	R	N	A	M	E	1	2
0	1	2	3	4	5	6	7	8	9

```
Billys-N90AP:/var/mobile root# printf "AAAABBBBCCCCDD
DDEEEE\x30\xbe\x00\x00\xff\xff\xff\xff\x70\xbe\x00\x0
0" | ./roplevel1
Welcome to ROPLevel1 for ARM! Created by Billy Ellis
(@bellis1000)
warning: this program uses gets(), which is unsafe.
Everything seems normal.
string changed.
executing string...
Applications          app          roplevel1.c
Containers            exploit.sh  roplevel1.zip
Developer             heap        taptapskip
Documents             heap.c     vuln
Library               hello      vuln.c
Media                 hello.c
MobileSoftwareUpdate roplevel1
Billys-N90AP:/var/mobile root#
```

Playing with such attacks on ARM:

<https://billy-ellis.github.io> (@bellis1000)

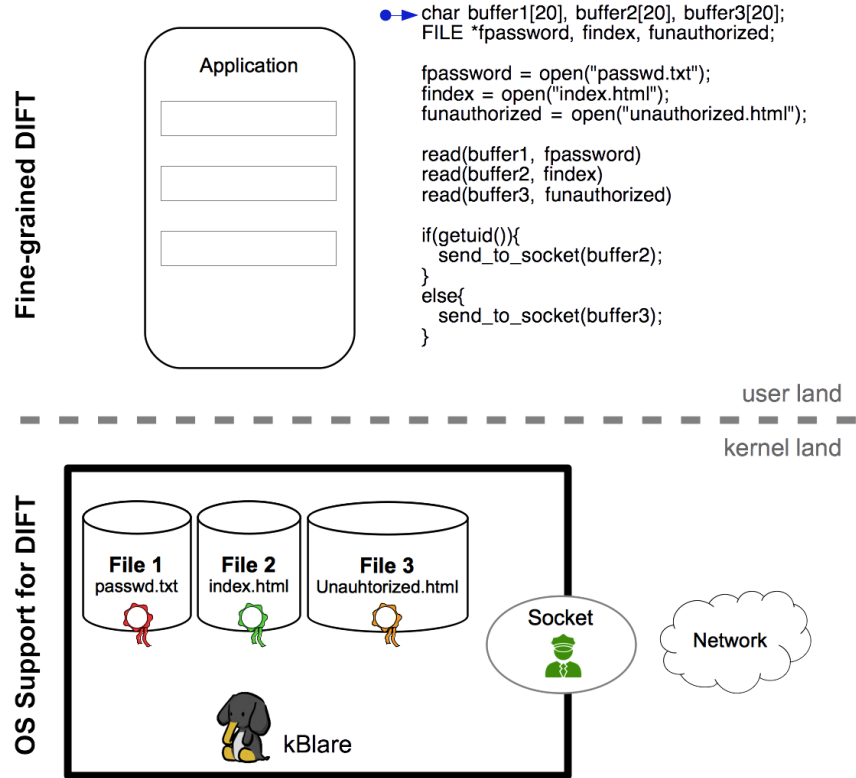
<https://www.root-me.org/?page=recherche&lang=en&recherche=ARM>

<https://azeria-labs.com/> (@Fox0x01)

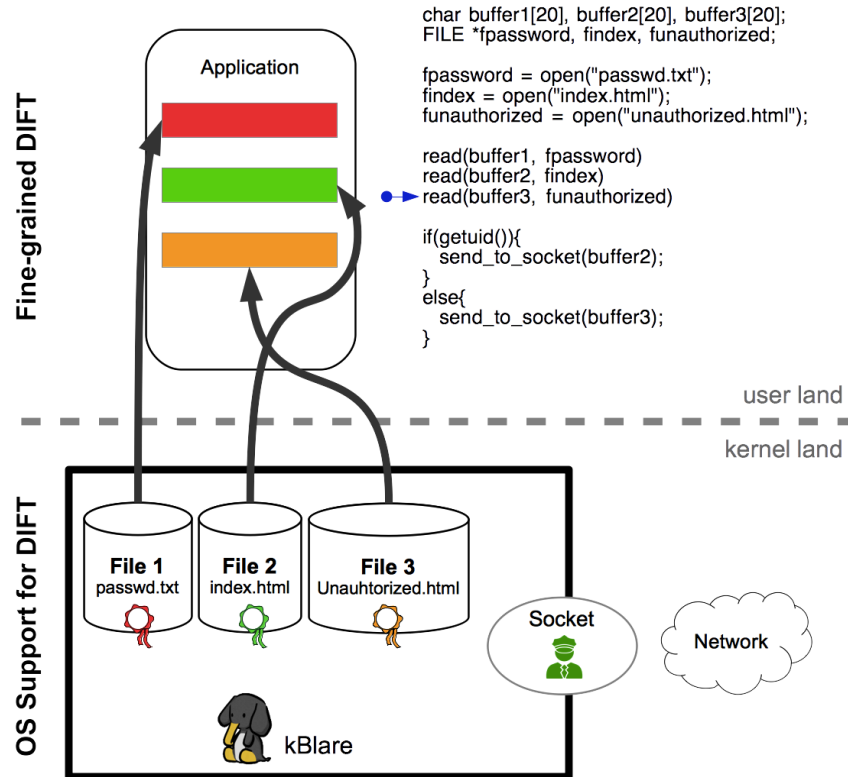
DIFT = Dynamic Information Flow Tracking

- DIFT => Detection of software attacks
 - Buffer overflow, Return Oriented Programming, etc.
- Security purposes => Integrity and Confidentiality
- Principle:
 - Tags attached to containers + relationship
 - At runtime, propagate tags
 - Detecting any violation at run-time asap

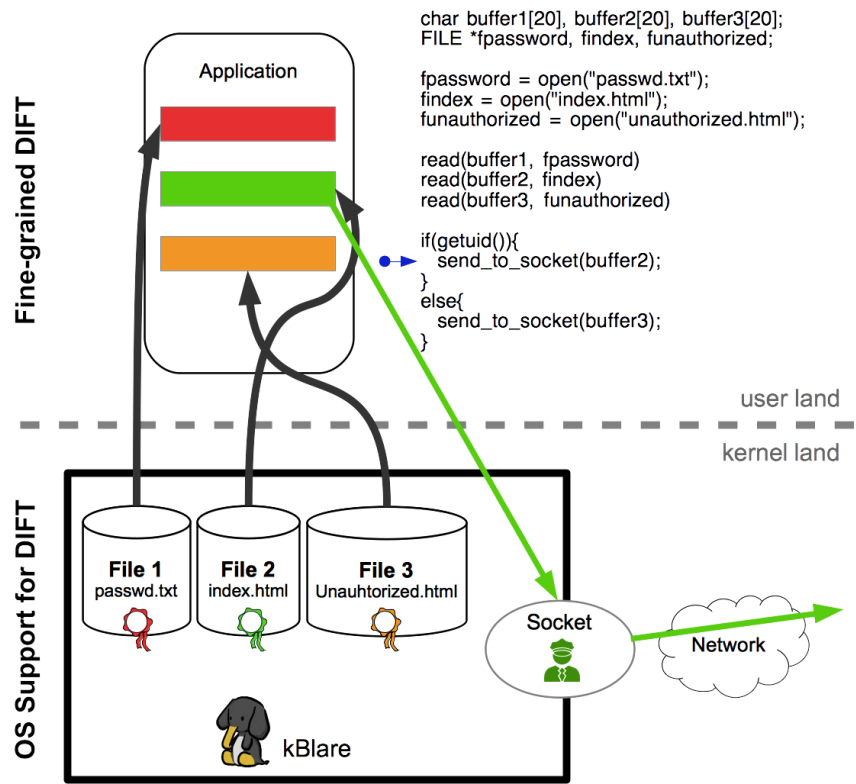
DIFT = Dynamic Information Flow Tracking



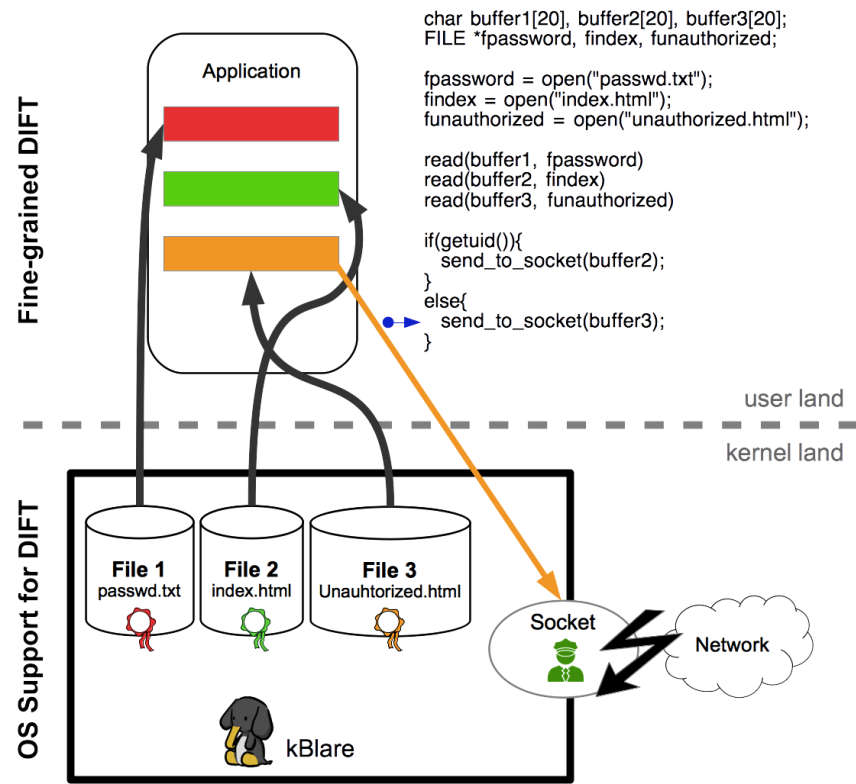
DIFT = Dynamic Information Flow Tracking



DIFT = Dynamic Information Flow Tracking



DIFT = Dynamic Information Flow Tracking



Different levels for DIFT

- Operating system:

Files / Executables

- Language level:

Variables / Functions

- Processor level:

Address, registers / Instructions

DIFT – Memory corruption detection

Attacker overwrites return address and takes control

```
int idx = tainted_input; //stdin (> BUFFER SIZE)
buffer[idx] = x; // buffer overflow
```

set $r_1 \leftarrow \&\text{tainted_input}$
load $r_2 \leftarrow M[r_1]$
add $r_4 \leftarrow r_2 + r_3$
store $M[r_4] \leftarrow r_5$

T	Data
	$r_1:\&\text{input}$
	$r_2:\text{idx}=\text{input}$
	$r_3:\&\text{buffer}$
	$r_4:\&\text{buffer}+\text{idx}$
	$r_5:x$

T	Data
	Return Address
	$\text{int buffer}[\text{Size}]$

DIFT – Memory corruption detection

Attacker overwrites return address and takes control

```
int idx = tainted_input; //stdin (> BUFFER SIZE)
buffer[idx] = x; // buffer overflow
```

set $r1 \leftarrow \&tainted_input$
load $r2 \leftarrow M[r1]$
add $r4 \leftarrow r2 + r3$
store $M[r4] \leftarrow r5$

T	Data
	$r1:\&input$
	$r2:idx=input$
	$r3:\&buffer$
	$r4:\&buffer+idx$
	$r5:x$

T	Data
	Return Address
	$int\ buffer[Size]$

DIFT – Memory corruption detection

Attacker overwrites return address and takes control

```
int idx = tainted_input; //stdin (> BUFFER SIZE)
buffer[idx] = x; // buffer overflow
```

set $r_1 \leftarrow \&\text{tainted_input}$
load $r_2 \leftarrow M[r_1]$
add $r_4 \leftarrow r_2 + r_3$
store $M[r_4] \leftarrow r_5$

T	Data
Red	$r_1:\&\text{input}$
Red	$r_2:\text{idx}=\text{input}$
Green	$r_3:\&\text{buffer}$
Green	$r_4:\&\text{buffer}+\text{idx}$
Green	$r_5:x$

T	Data
Green	Return Address
Green	int buffer[Size]

DIFT – Memory corruption detection

Attacker overwrites return address and takes control

```
int idx = tainted_input; //stdin (> BUFFER SIZE)
buffer[idx] = x; // buffer overflow
```

set $r_1 \leftarrow \&\text{tainted_input}$
load $r_2 \leftarrow M[r_1]$
add $r_4 \leftarrow r_2 + r_3$
store $M[r_4] \leftarrow r_5$

T	Data
Red	$r_1:\&\text{input}$
Red	$r_2:\text{idx}=\text{input}$
Green	$r_3:\&\text{buffer}$
Red	$r_4:\&\text{buffer}+\text{idx}$
Green	$r_5:x$

T	Data
Green	Return Address
Green	int buffer[Size]

DIFT – Memory corruption detection

Attacker overwrites return address and takes control

```
int idx = tainted_input; //stdin (> BUFFER SIZE)
buffer[idx] = x; // buffer overflow
```

set $r1 \leftarrow \&tainted_input$
load $r2 \leftarrow M[r1]$
add $r4 \leftarrow r2 + r3$
store $M[r4] \leftarrow r5$

T	Data
■	$r1:\&input$
■	$r2:idx=input$
■	$r3:\&buffer$
■	$r4:\&buffer+idx$
■	$r5:x$

T	Data
■	Return Address
■	int buffer[Size]

Different levels for DIFT

- Tag initialization: data are tagged with their "security level"

`password="abcd" Tag(password)=secret`

- Tag propagation: any new data derived from the tagged data is also tagged

`log=err+password Tag(log)=Tag(err)+Tag(password)` 

- Tag check: raise an exception if an information flow doesn't respect a security policy

`write(log,network) Policy: (Tag(log)==public)`

Different levels for DIFT

- Application level
 - Java / Android, Javascript, C
- OS level
 - kBlare (Linux kernel w/ software IFT)
- Low level
 - Deeping into processor architecture maybe?

Different levels for DIFT

- Application level
 - Java / Android, Javascript, C
- OS level
 - kBlare (Linux kernel w/ software IFT)
- Low level
 - Deeping into processor architecture maybe?

Buying an ARM license => no way. Or...

FPGA => Programmable electronics

FIELD PROGRAMMABLE GATE ARRAYS (FPGA'S)

DISADVANTAGES

- * EXPENSIVE
- * HIGH POWER
- * VOLATILE / BOOT TIME
- * HIGH PIN COUNT / BGA
- * COMPLICATED
- * MANY TRAPS
- * COMPLEX TOOLS
- * HARD TO CHOOSE / COMPARE
- * HDL NOT EASY / INTUITIVE

IOB IOB IOB IOB IOB
IOB IOB IOB IOB IOB
IOB IOB IOB IOB IOB
IOB IOB IOB IOB IOB
IOB IOB IOB IOB IOB
IOB IOB IOB IOB IOB
IOB IOB IOB IOB IOB

CONFIG FLASH MEMORY

CONFIG

CLB CLB CLB
CLB CLB CLB
CLB CLB

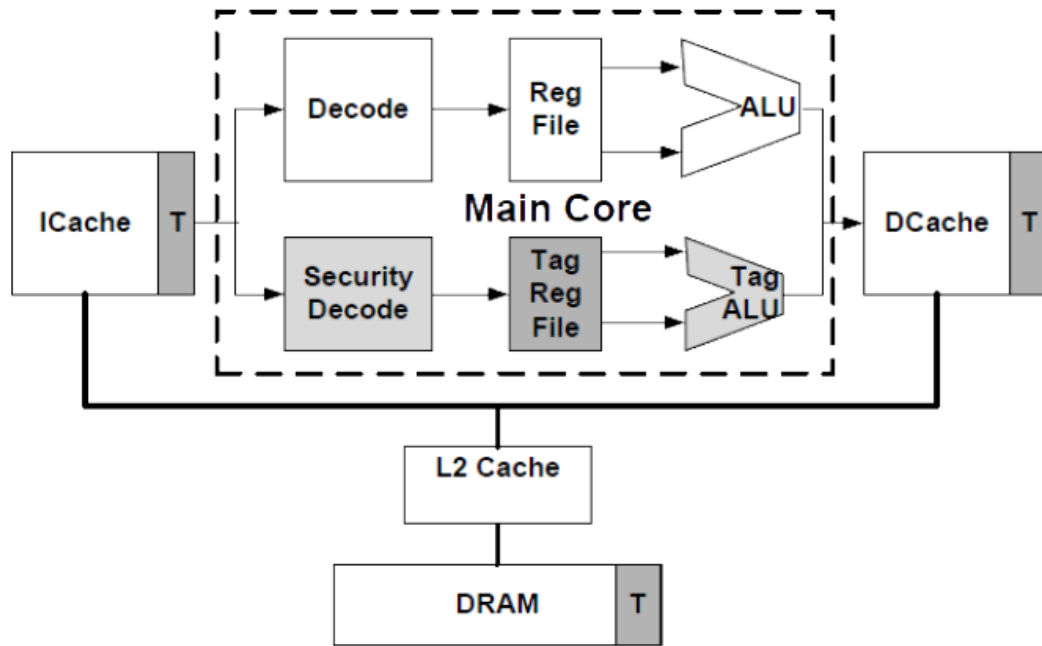
IOB IOB

IOB BLOCK

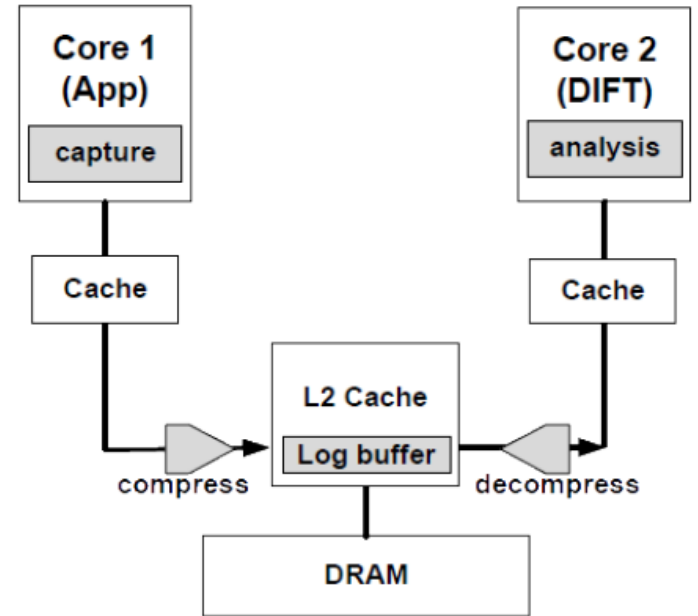
LOGIC BLOCK/ELEMENT

Source: EEVBlog #496 – What is an FPGA? (Youtube)

Different levels for DIFT

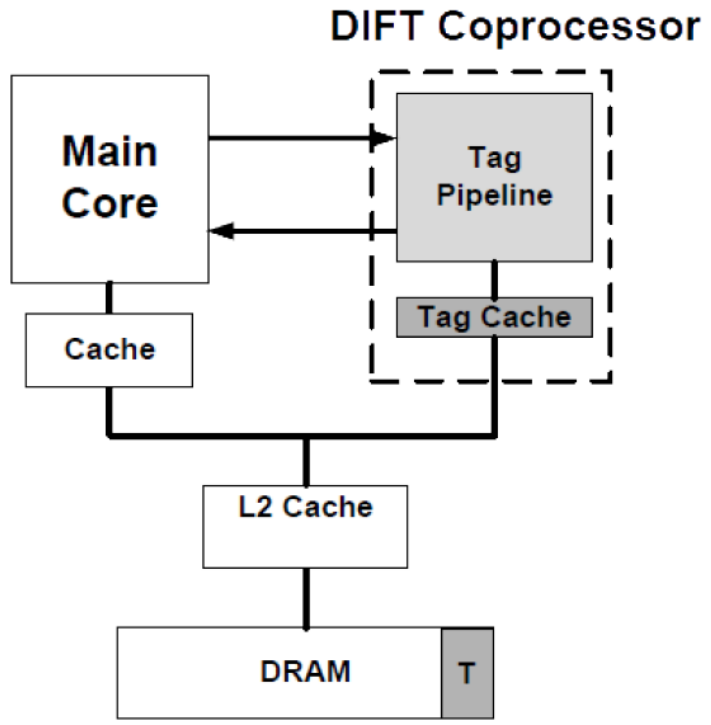


In-core DIFT



Offloading

Different levels for DIFT



Off-core DIFT

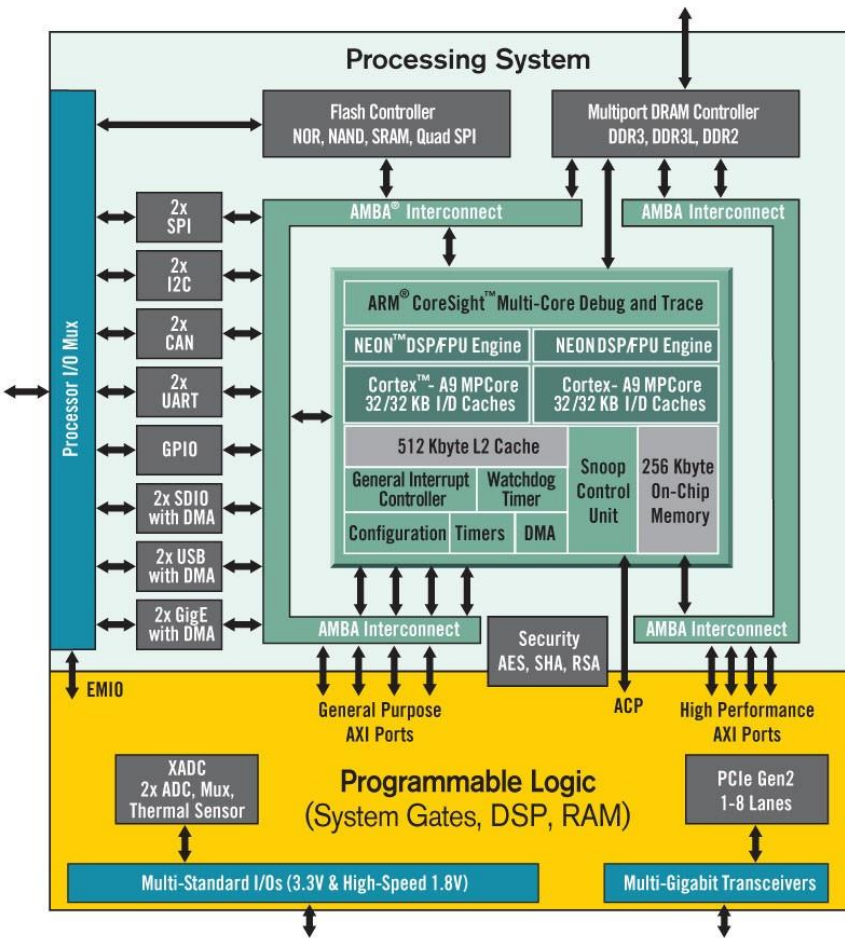
Related works

	Advantages	Disadvantages
Software	Flexible security policies	Overhead (300% at least...)
In-core DIFT	Low overhead (10%)	Invasive modifications
Dedicated CPU	Low overhead (10%)	Wasting resources
Dedicated coprocessor	Low overhead (10%) CPU not modified	CPU/coprocessor communication

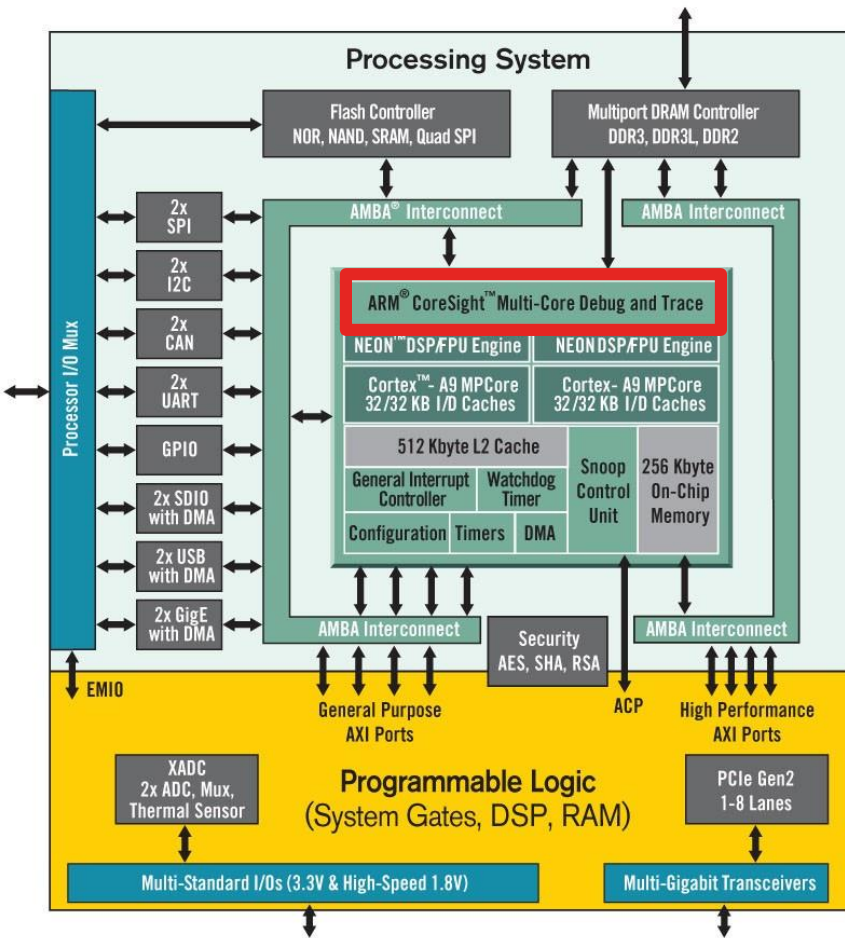
- Limiting the impact of software instrumentation
- Security of the coprocessor
- First work on ARM-based SoCs
- Additional challenges

- Limiting the impact of software instrumentation
- Security of the coprocessor
- **First work on ARM-based SoCs**
- Additional challenges

What can I do with my processor?

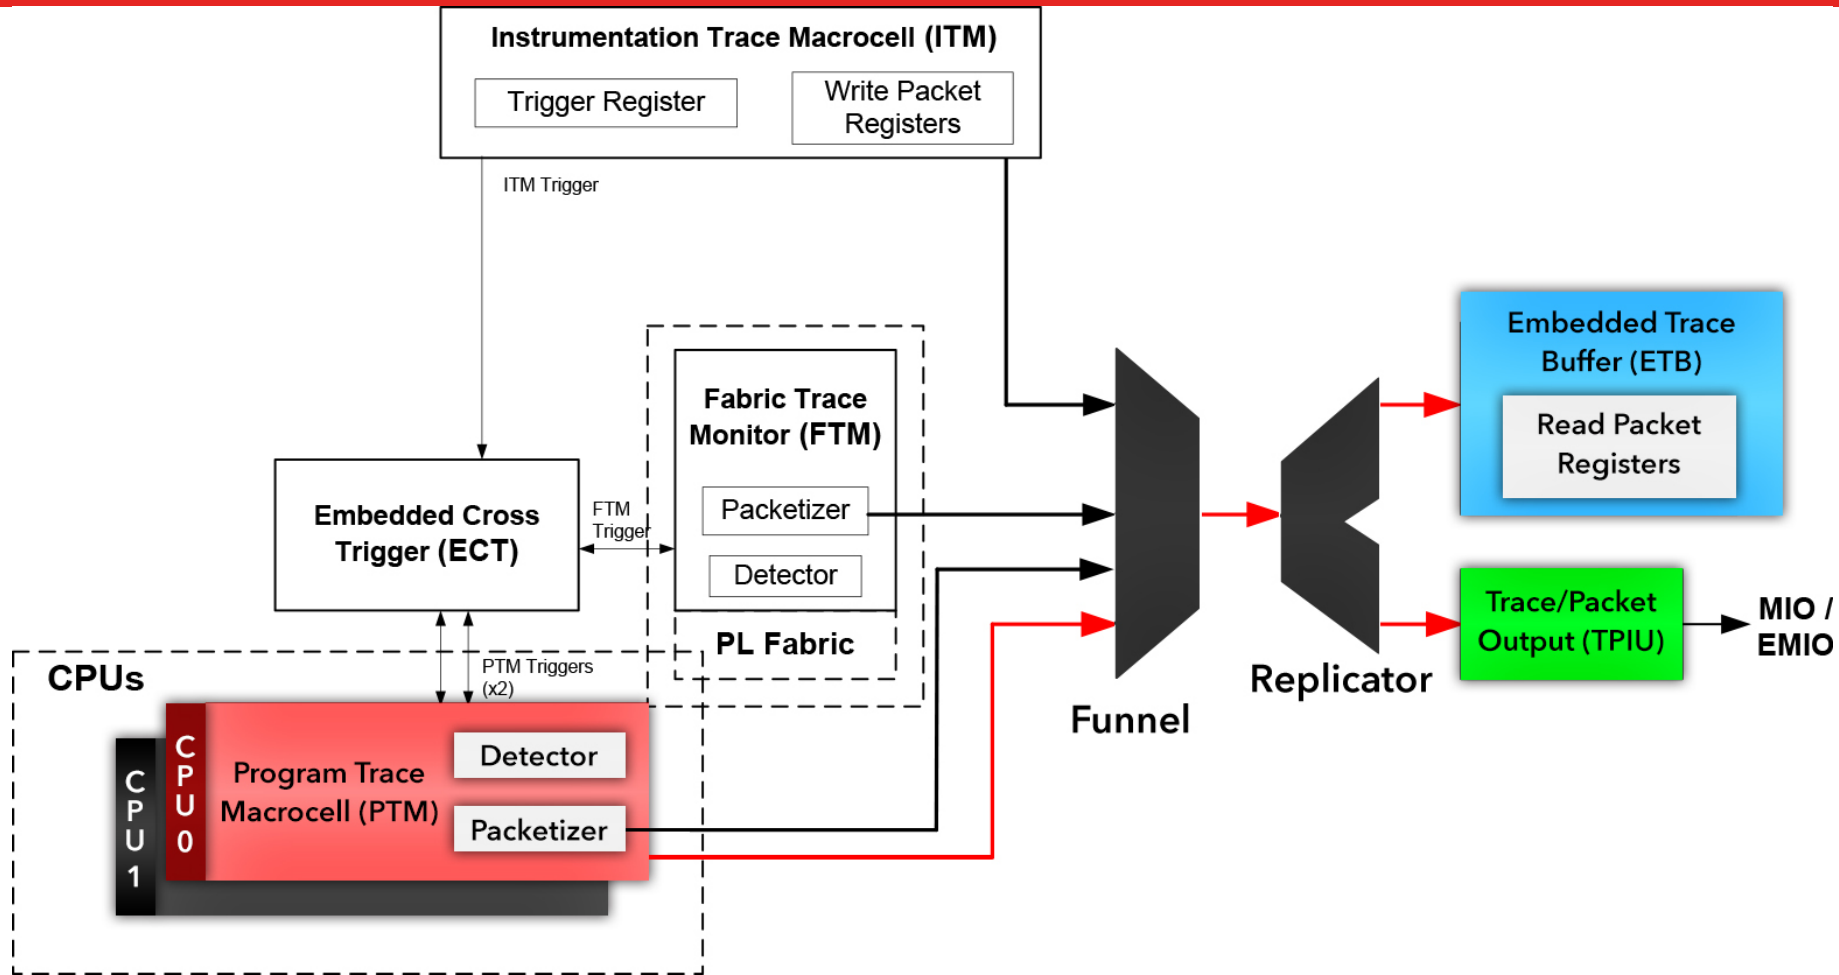


What can I do with my processor?

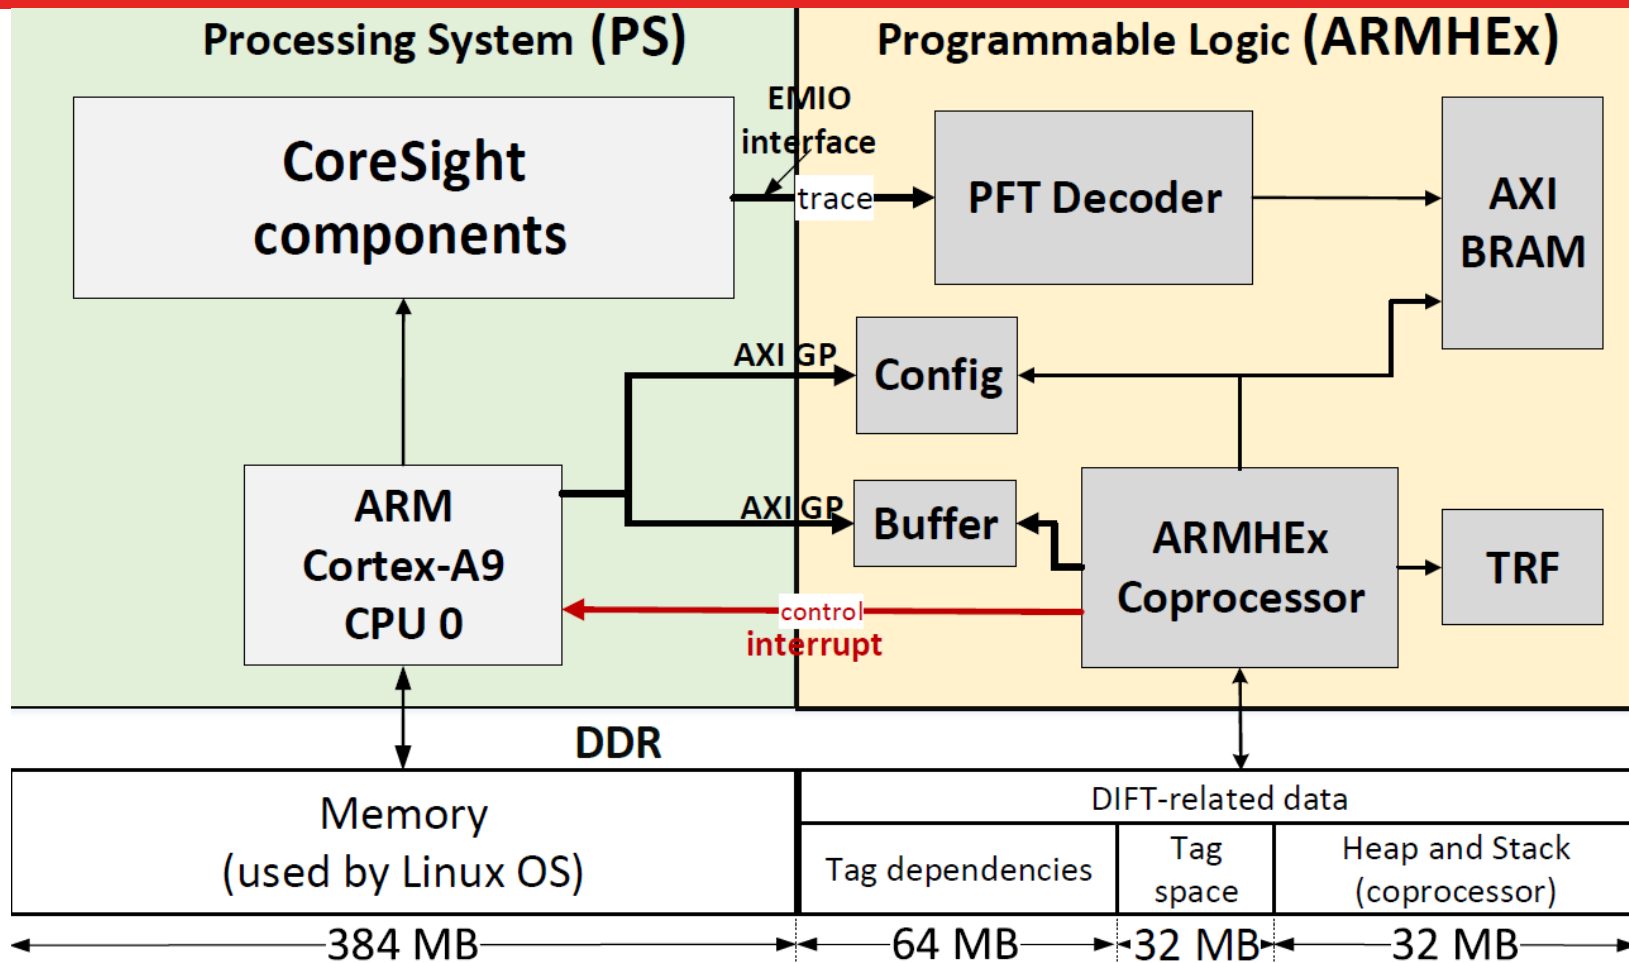


- CoreSight: debug components
- Available in most of Cortex-A + Cortex-M3 (for ARM)
- Can export stuff

CoreSight components

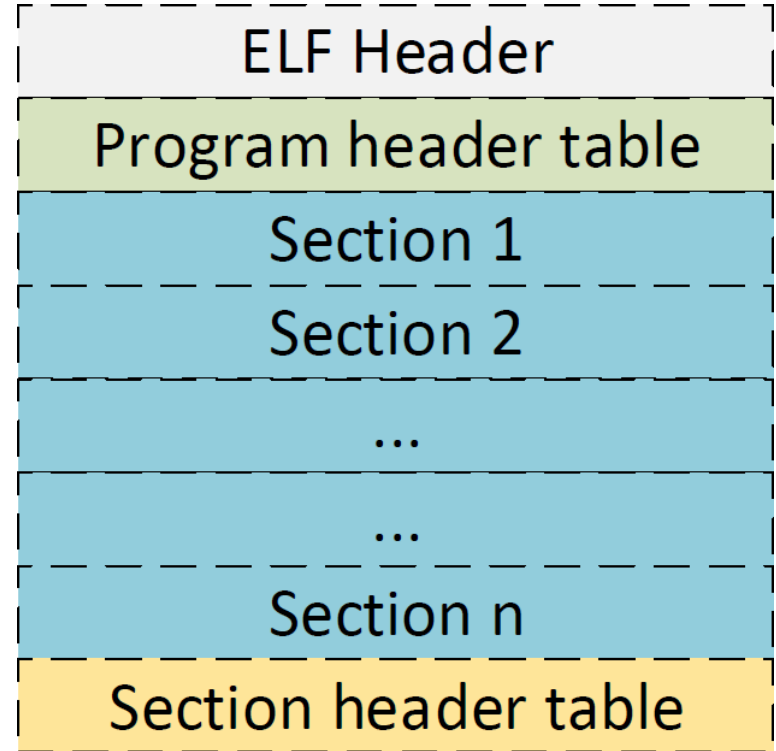


CoreSight components – Where should I export my metadata?



Features:

- Trace filter
- Branch Broadcast
- Timestamping
- Etc, etc.



What does a trace look like?

Source code

```
int i;  
for(i=0;i<10;i++)
```

Assembly

```
8638 for_loop:  
...  
b 8654 :  
...  
866c : bcc 8654
```

Trace

```
00 00 00 00 00 80 08 38 86 00 00 21  
2a 2a 2a 2a 2a 2a 2a 2a 2a 2a 86 01  
00 00 00 00 00 00 00 00
```

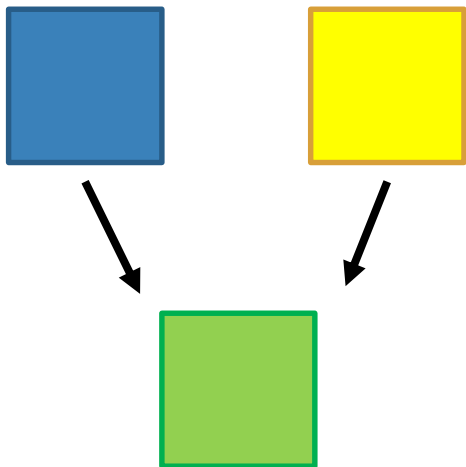
Decoded trace

```
A-sync  
Address 00008638, (I-sync Context  
00000000, IB 21)  
Address 00008654, Branch Address  
packet (x 10)
```

Our case:

- We want to store tags and initialize tags from the operating system:
 - Modified kBlare (based on a **Linux Kernel 4.9**)
- We don't want to lose information (no over-approximation):
 - **Dynamic approach:** Instrumentation + PTM traces
- Extract some information about the data flow (for tag propagation):
 - **Static Analysis:** Generating annotations.

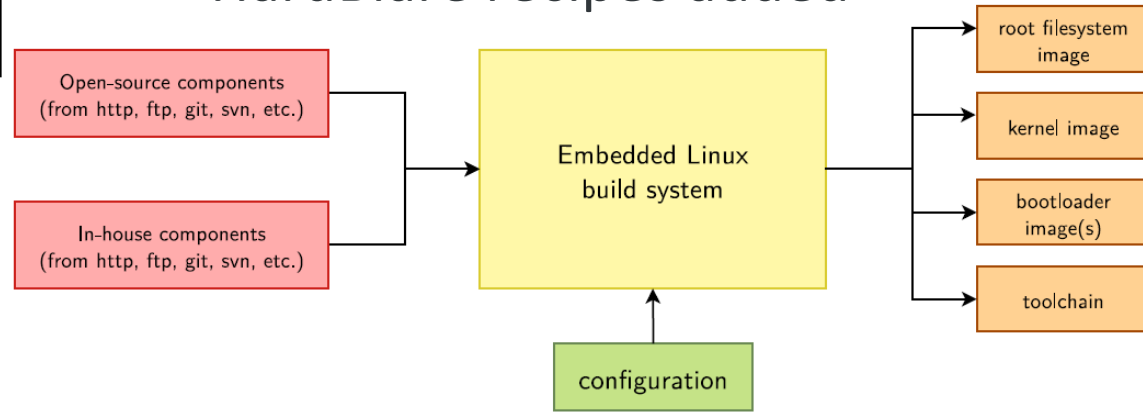
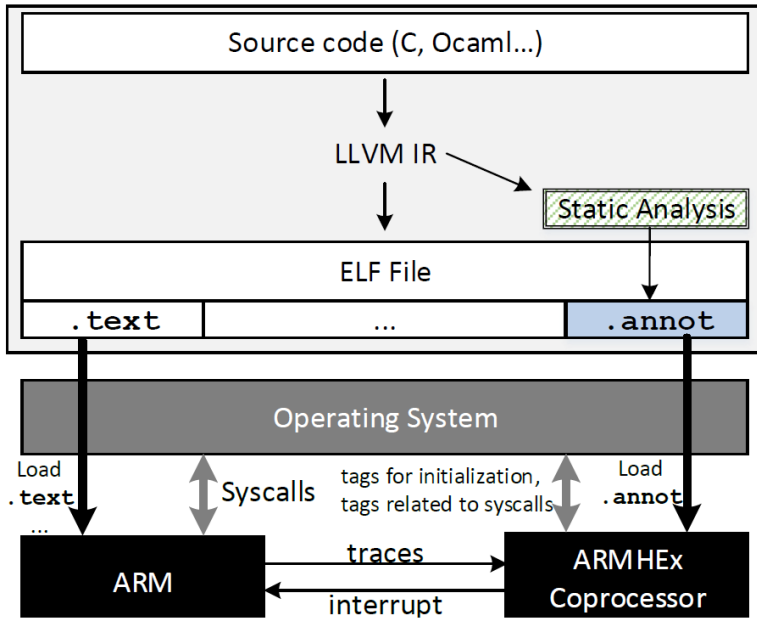
Generating annotations



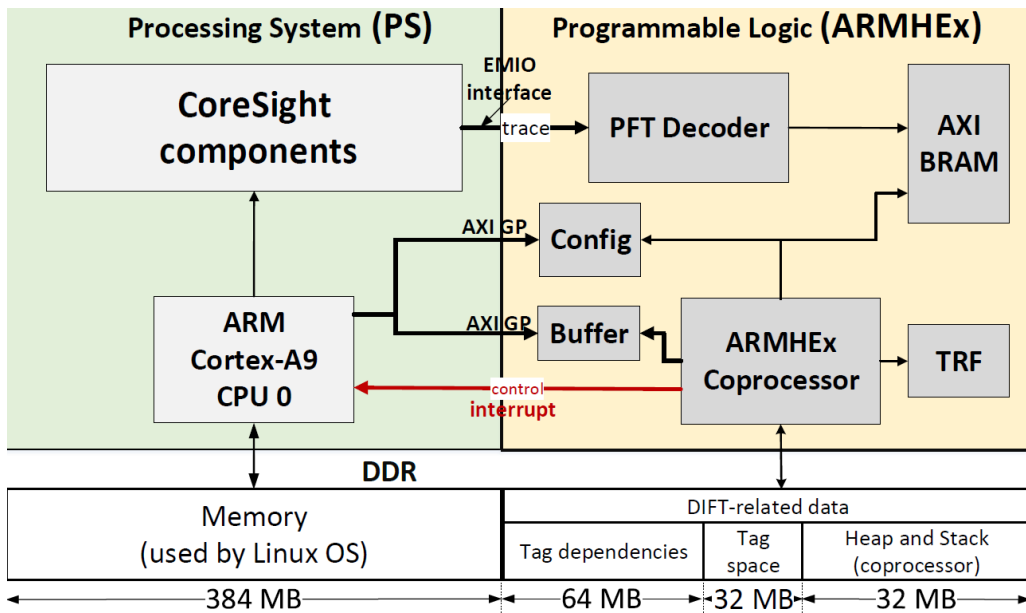
(status on late February)

- 200 instructions done:
 - LLVM meta-instructions
 - « Basic » stuff: add, compare, load/store, etc.
- TODO: 200 instructions left (at least...)
 - Parallel additions/subtractions features
 - Advanced SIMD instructions

- Templates/tools/methods
- Custom embedded Linux
- HardBlare recipes added

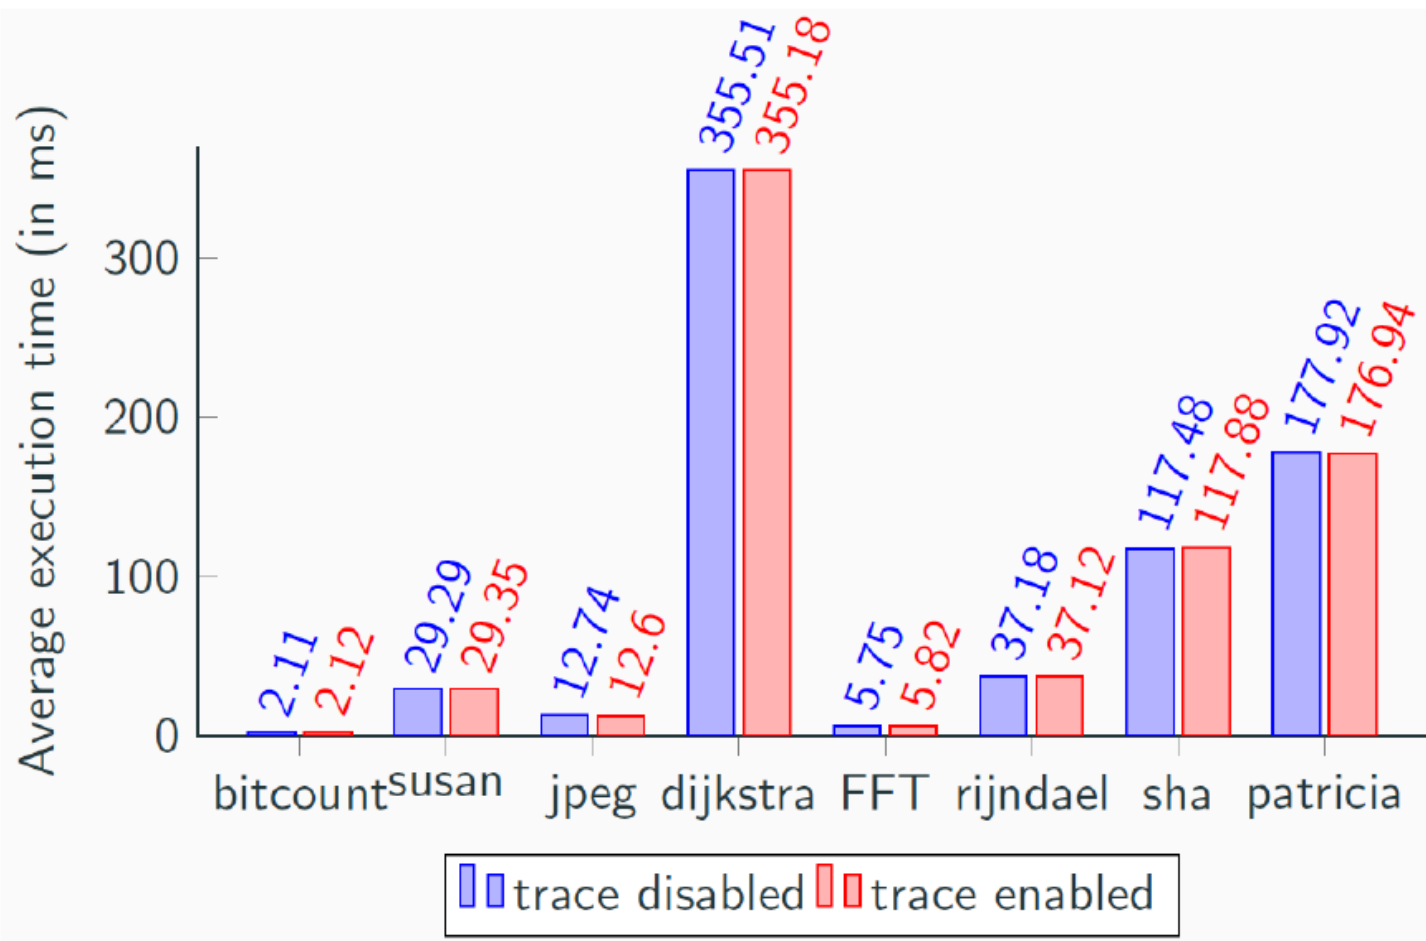


Coprocessor – Quick hints



- DIFT metadata protection
 - TrustZone + secure world
- Main challenge: speed!

Some latency results




Comparison w/ existing works

Approaches	Kannan	Deng	Heo	ARMHEX
Hardcore portability	No	No	Yes	Yes
Main CPU	Softcore	Softcore	Softcore	Hardcore
Communication overhead	N/A	N/A	60%	5.4%
Area overhead	6.4%	14.8%	14.47%	0.47%
Area (Gate Counts)	N/A	N/A	256177	128496
Power overhead	N/A	6.3%	24%	16%
Max frequency	N/A	256 MHz	N/A	250 MHz
Isolation	No	No	No	Yes

Take away:

- CoreSight PTM allows to obtain runtime information (Program Flow)
- Non-intrusive tracing => Negligible performance overhead



RaspberryPi PoC (hopefully March)

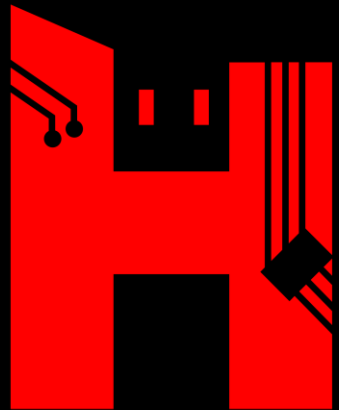
Full PoC later this year (SoC files + Yocto)

Intel / ST? (study)

Multicore multi-thread IFT

Full-speed IFT

Monitoring program execution (and more!) on ARM processors



Toulouse
Hacking
Convention_

Pascal Cotret
pascal.cotret@gmail.com / @Pascal_r2

Many thanks to Muhammad, Mounir, Guy,
Guillaume, Vianney and Arnab