


Monitoring program execution on ARM processors thanks to hardware components

Pascal Cotret, February 28 @ Evry






Hello!

I am Pascal Cotret

- Embedded software security engineer
- Research in my spare-time



HardBlare (2015/2018)

Funded by Labex CominLabs



Inria
INVENTEURS DU MONDE NUMÉRIQUE



- 3 labs (INRIA CIDRE, Lab-STICC, SCEE/IETR)
- 2 PhD students.
- 1 postdoc (from February 2018).

=> Heterogeneous information flow control

- DIFT: a short introduction
- Related works in hardware-assisted DIFT
- What can we do with ARM processors?
- Results
- Conclusion & perspectives

Dynamic Information Flow Tracking

Dynamic Information Flow Tracking (DIFT) is a promising technique for detecting software attacks.


Motivation: DIFT for security purposes => Integrity and Confidentiality

DIFT principle:

- We attach tags to containers and specify an information flow policy, i.e. relations between tags.
- At runtime, we propagate tags to reflect information flows that occurs.
- Allows to detect any security policy violation at run-time.

DIFT

Fine-grained DIFT



```
char buffer1[20], buffer2[20], buffer3[20];  
FILE *fpassword, findex, funauthorized;
```

```
fpassword = open("passwd.txt");  
findex = open("index.html");  
funauthorized = open("unauthorized.html");
```

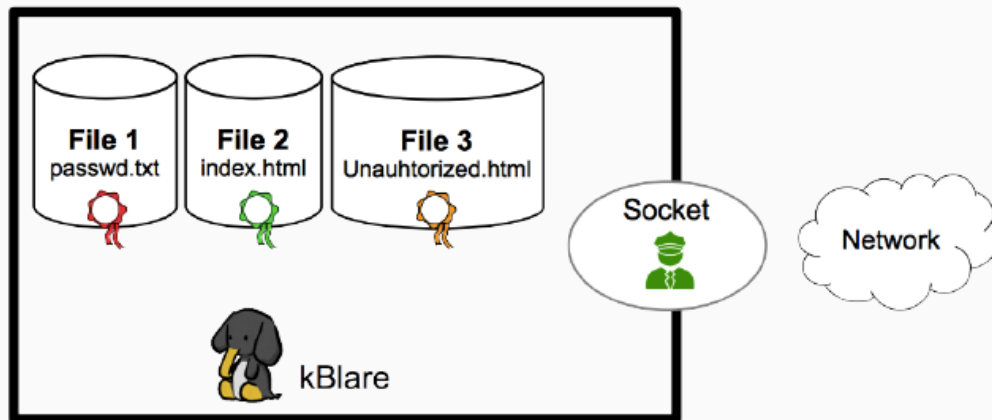
```
read(buffer1, fpassword)  
read(buffer2, findex)  
read(buffer3, funauthorized)
```

```
if(getuid()){  
    send_to_socket(buffer2);  
}  
else{  
    send_to_socket(buffer3);  
}
```

user land


kernel land

OS Support for DIFT



DIFT

Fine-grained DIFT



```
char buffer1[20], buffer2[20], buffer3[20];  
FILE *fpassword, findex, funauthorized;
```

```
fpassword = open("passwd.txt");  
findex = open("index.html");  
funauthorized = open("unauthorized.html");
```

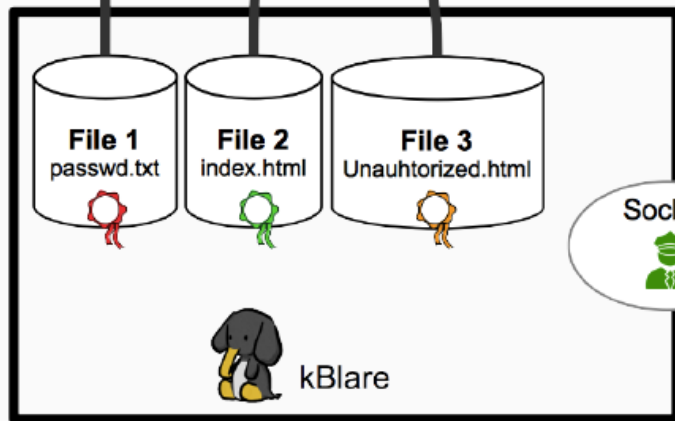
```
read(buffer1, fpassword)  
read(buffer2, findex)  
•▶ read(buffer3, funauthorized)
```

```
if(getuid()){  
    send_to_socket(buffer2);  
}  
else{  
    send_to_socket(buffer3);  
}
```

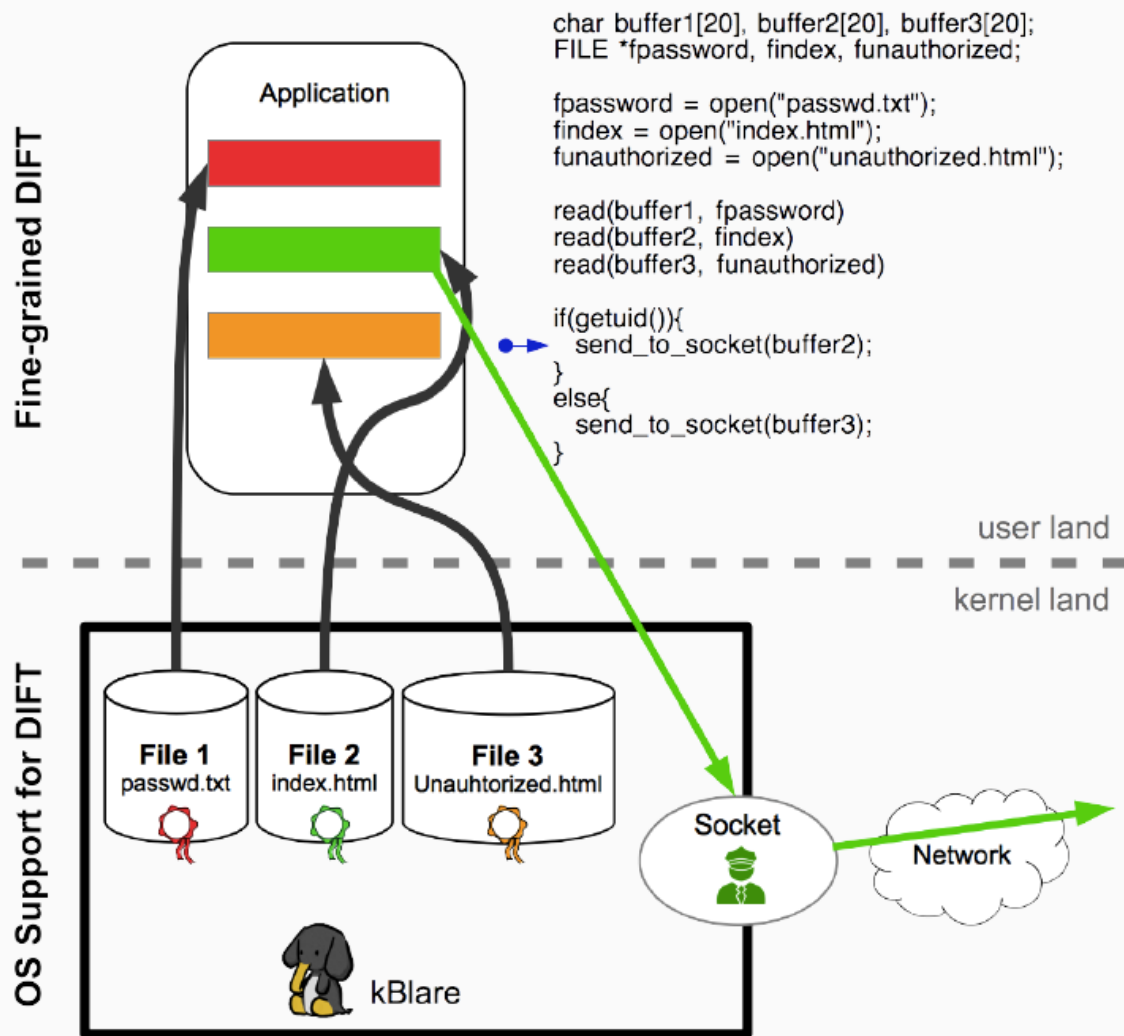
user land

kernel land

OS Support for DIFT



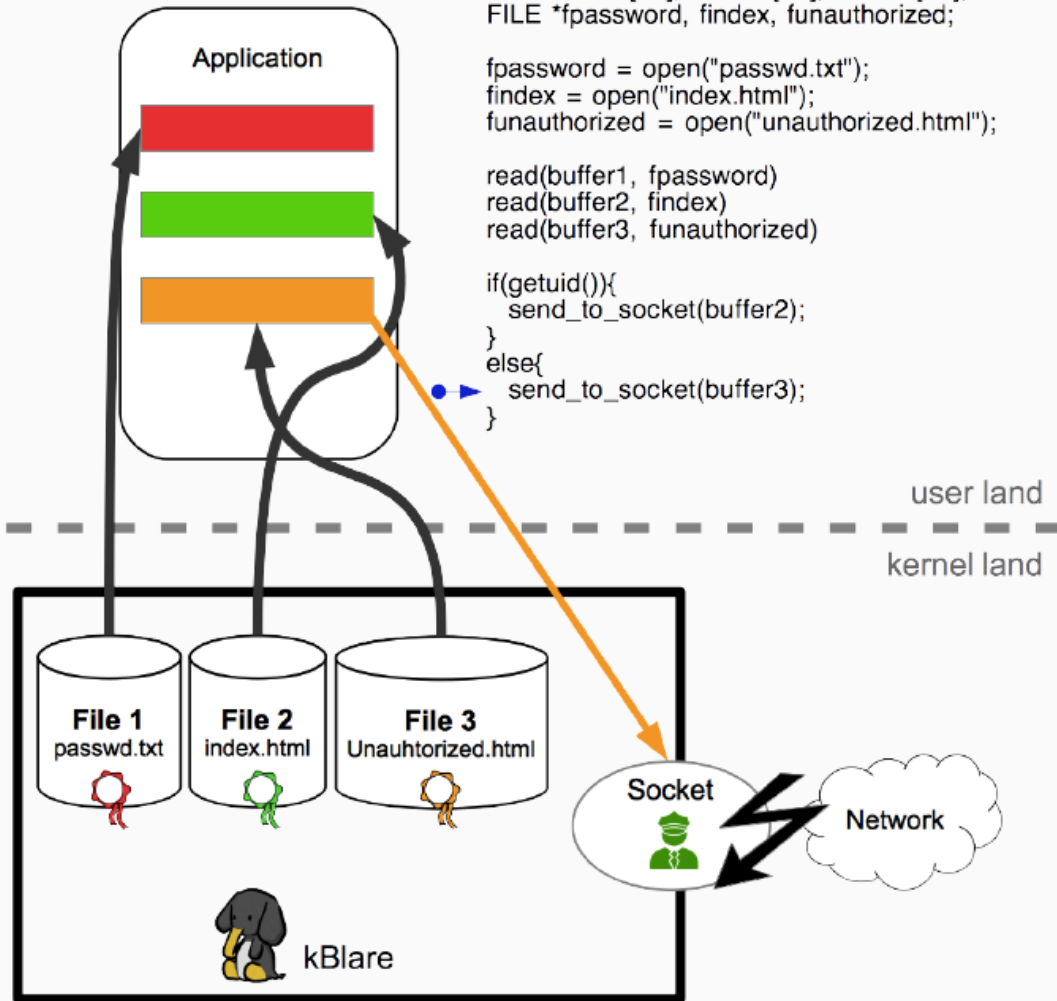
DIFT



DIFT

Fine-grained DIFT

OS Support for DIFT



DIFT – Three main steps

- **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

- Fine-grained (processor level)

`vi = addresses and registers; P = instructions`

- Medium-grained (language level)

`vi = variables; P = functions`

- Coarse-grained (operating system level)

`vi = files; P = executables`

DIFT

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

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

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

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
Red	r1:&input
Red	r2:idx=input
Green	r3:&buffer
Red	r4:&buffer+idx
Green	r5:x

T	Data
Green	Return Address
Green	int buffer[Size]

DIFT

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
Red	r1:&input
Red	r2:idx=input
Green	r3:&buffer
Red	r4:&buffer+idx
Red	r5:x

T	Data
Red	Return Address
Green	int buffer[Size]

Related work



Different levels:


- Application level
 - Java / Android, Javascript, C
- OS level
 - Laminar, HiStar
 - kBlare (1)
- Low level
 - Raksha (Kannan et al.)
 - Flexitaint (Ventakaramani et al.)
 - Flexcore (Deng et al.)
 - PAU (Heo et al.)

(1) Jacob Zimmermann, Ludovic Mé, and Christophe Bidan. Introducing Reference Flow Control for Detecting Intrusion Symptoms at the OS Level. In : RAID 2002.

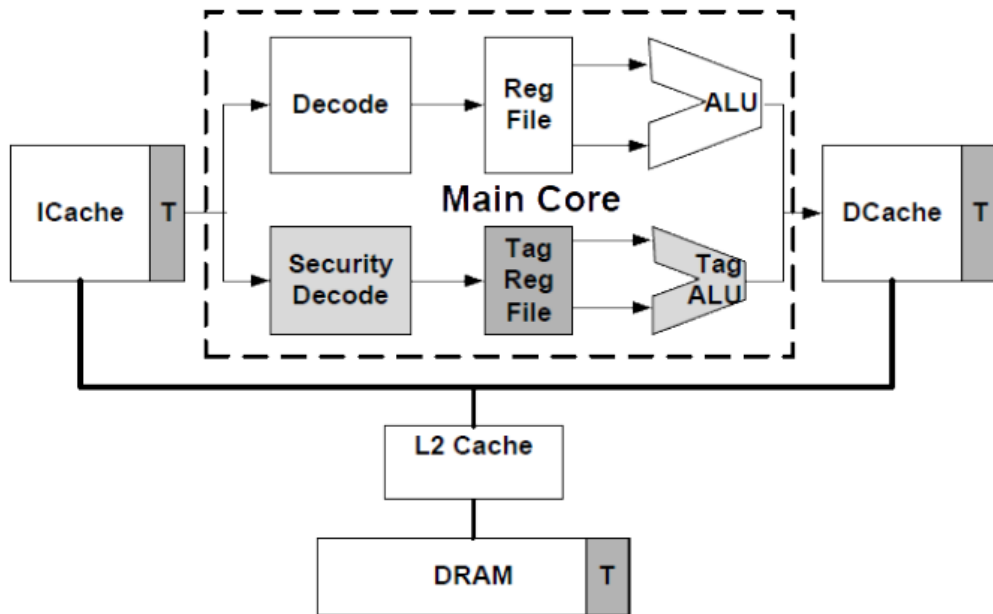
Target architecture

SoC = System-on-Chip

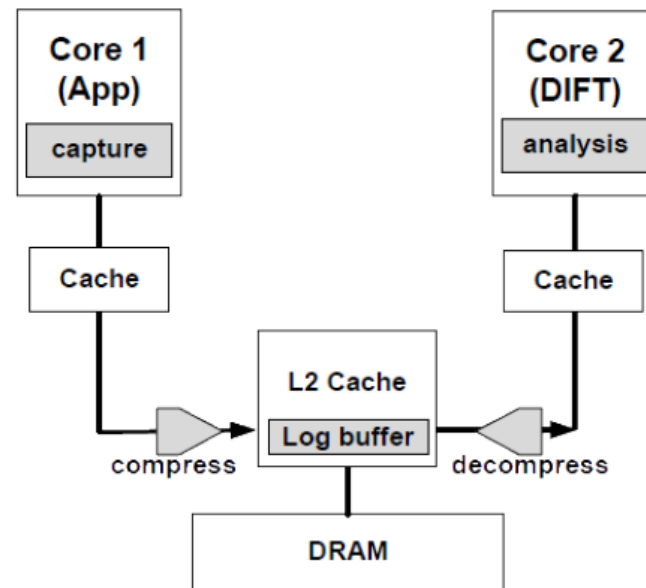
- 1+ processor(s)
- Configurable electronic chip (aka FPGA)



Related work

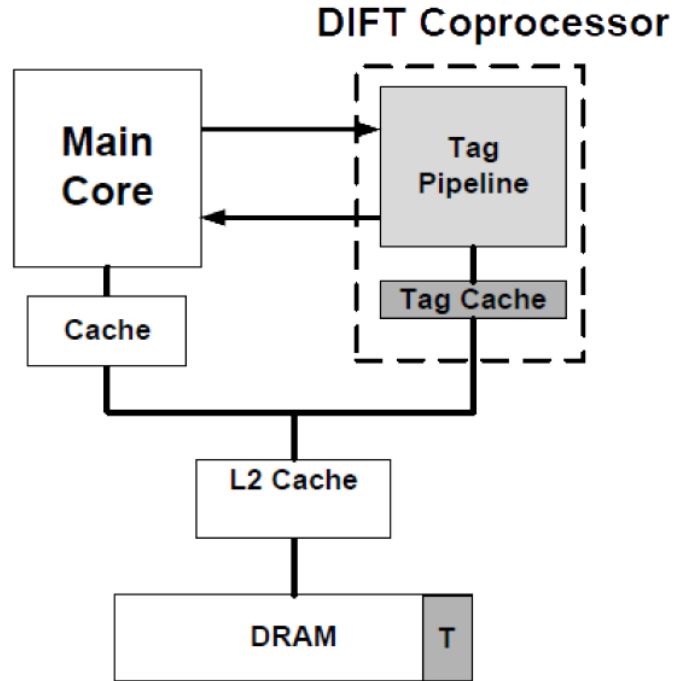


In-core DIFT



Offloading DIFT

Related work



Off-core DIFT

(2) Hari Kannan, Michael Dalton, and Christos ozyrakis. Decoupling dynamic information flow tracking with a dedicated coprocessor. In : Dependable Systems & Networks, 2009. IEEE. 2009, pp. 105-114.

Related work

	Advantages	Disadvantages
Software	Flexible security policies Multiple attacks detected	Overhead (from 300% to 3700%)
In-core DIFT	Low overhead (<10%)	Invasive modifications Few security policies
Dedicated CPU for DIFT	Low overhead (<10%) Few modifications to CPU	Wasting resources Energy consumption (x2)
Dedicated DIFT coprocessor	Flexible security policies Low overhead (<10%) CPU not modified	CPU/coprocessor communication

Related work

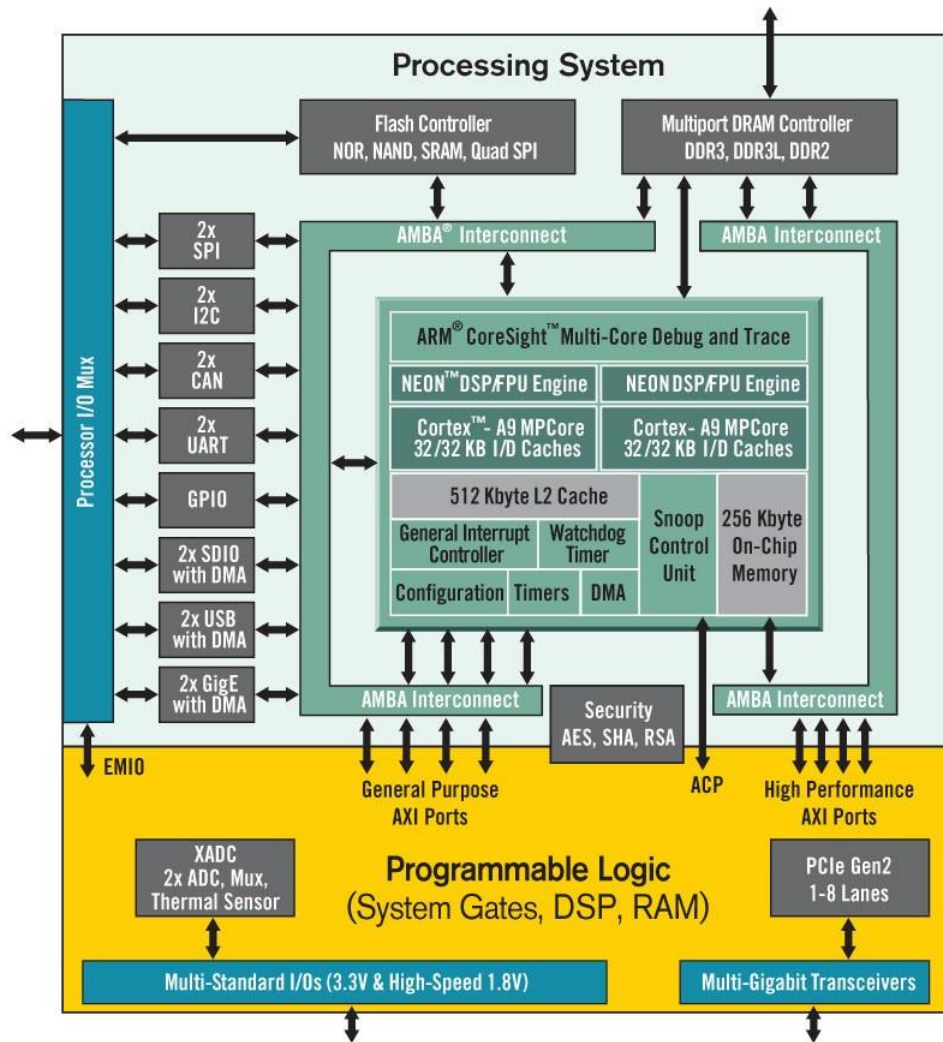
ARMHEx approach:

- Reduce overhead of software instrumentation as it represents the major portion of overall DIFT execution time overhead
- Lack of security of DIFT coprocessor
- No existing work targets ARM-based SoCs (related work implementations on softcores)
- Additional challenges
 - Limited visibility
 - Frequency gap between CPU and DIFT coprocessor
 - Communication interface...

Target architecture

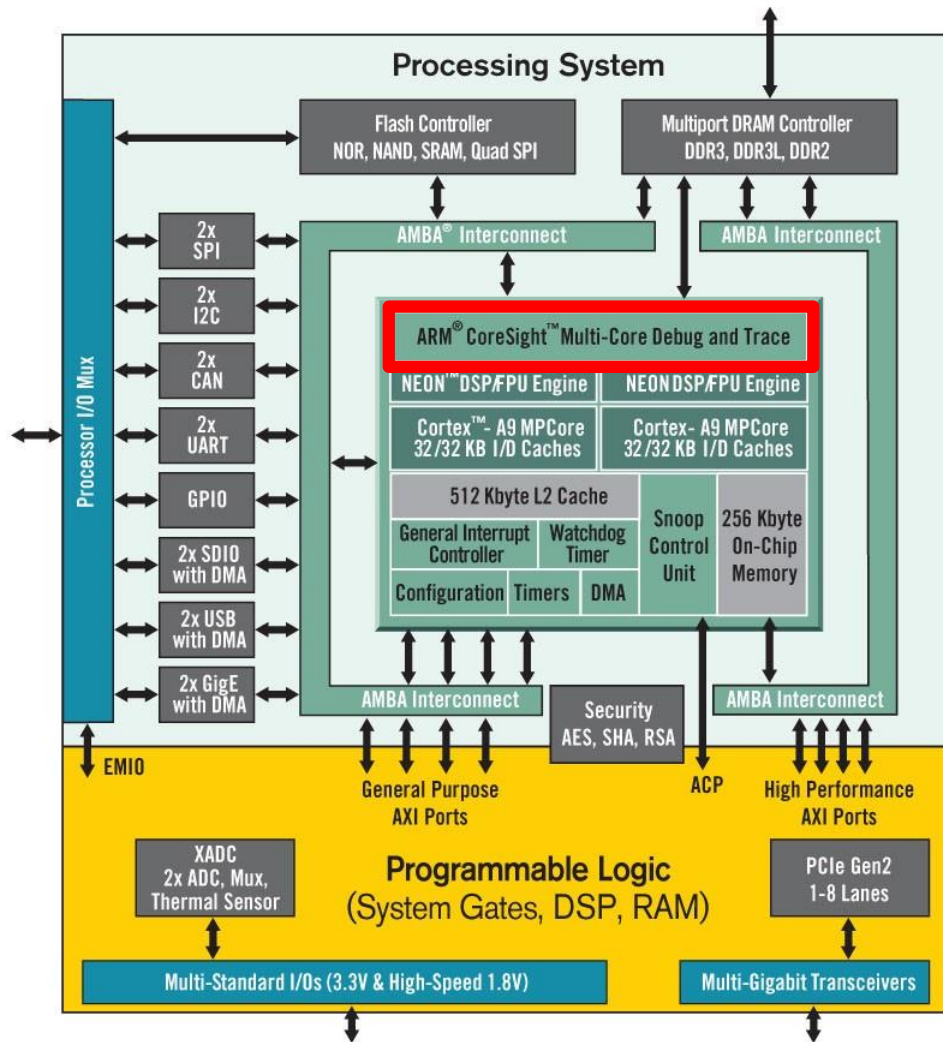
SoC = System-on-Chip

- 1+ processor(s)
- Configurable electronic chip (aka FPGA)

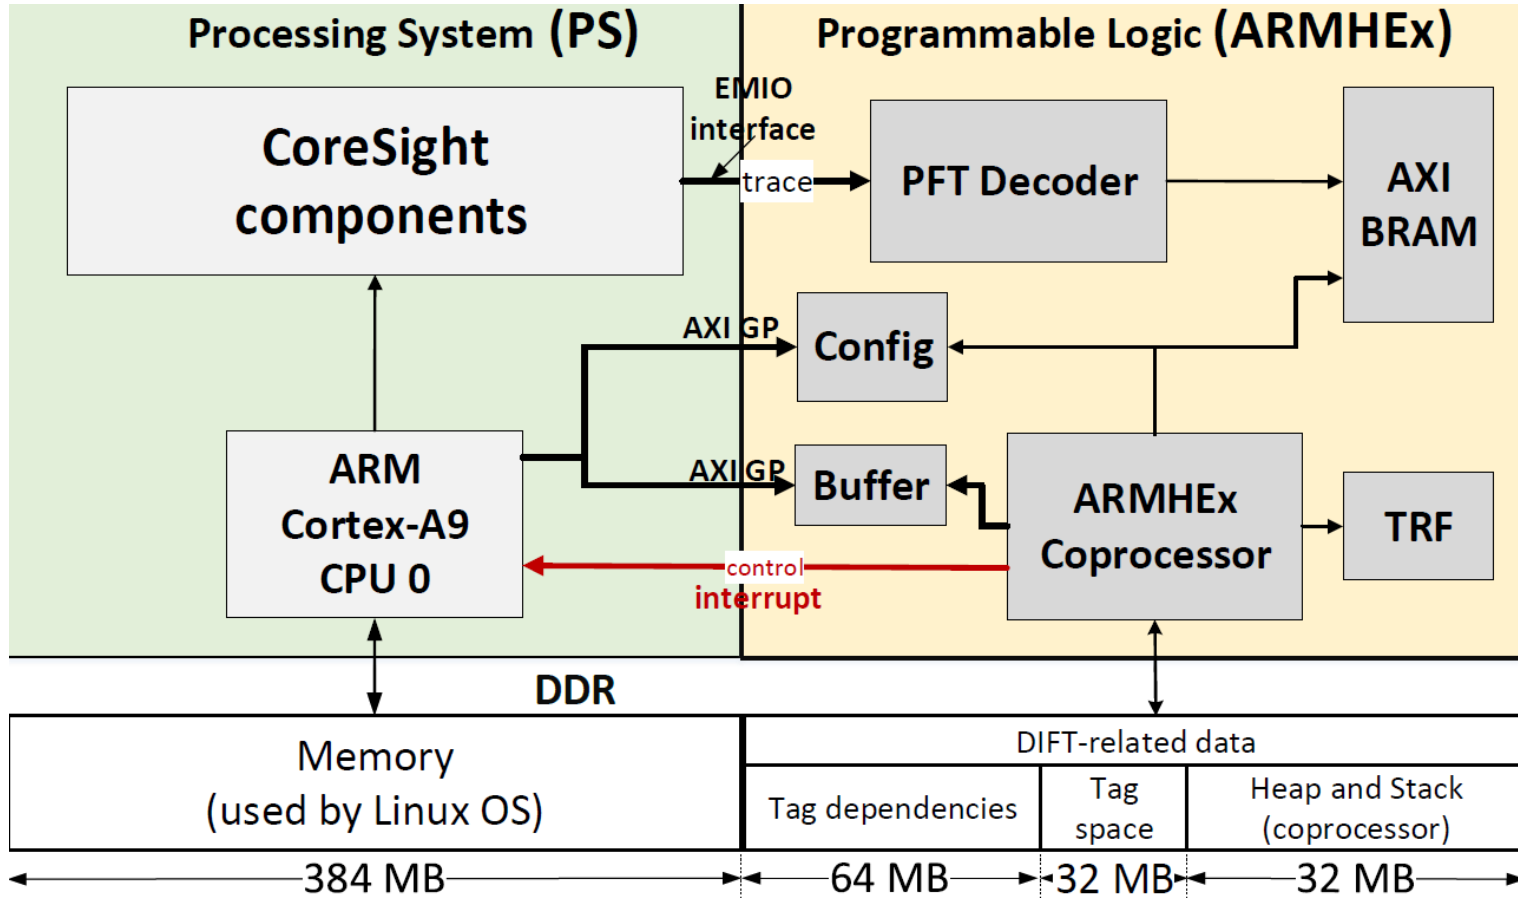


CoreSight components

A set of IP blocks providing HW-assisted system tracing

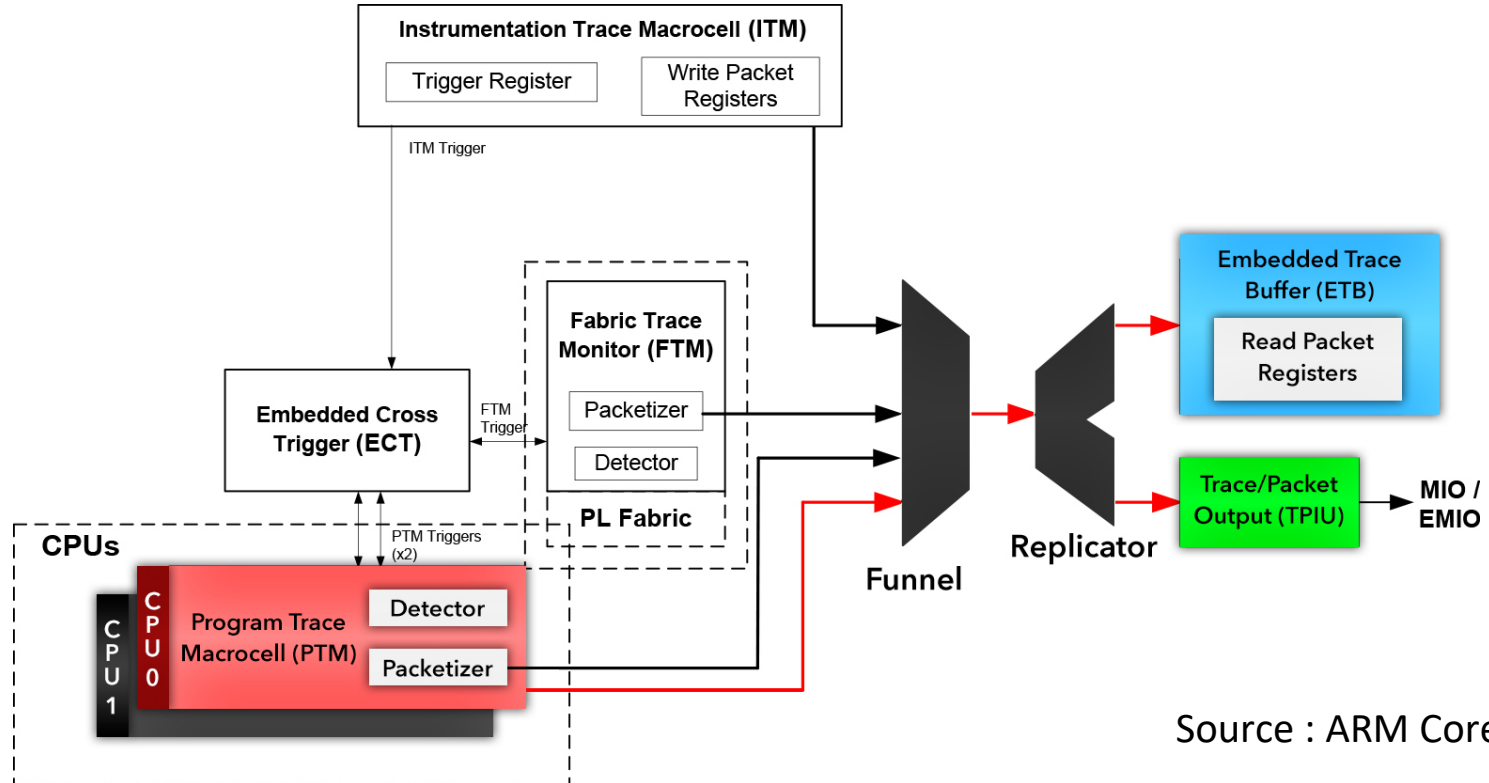


Overall architecture



CoreSight components

A set of IP blocks providing HW-assisted system tracing




Source : ARM CoreSight TRM

CoreSight components

A set of IP blocks providing HW-assisted system tracing

C11.11.31 DBGOSLAR, OS Lock Access Register

The DBGOSLAR bit assignments are:



OS Lock Access, bits[31:0]

Writing the key value `0xC5ACCE55` to this field locks the debug registers. In v7 Debug, the write also resets the internal counter for the OS Save or OS Restore operation.

Writing any other value to this register unlocks the debug registers if they are locked.

See [The OS Save and Restore mechanism](#) on page C7-2154 for a description of using the OS Save and Restore mechanism registers, including the behavior when the OS Lock is set.

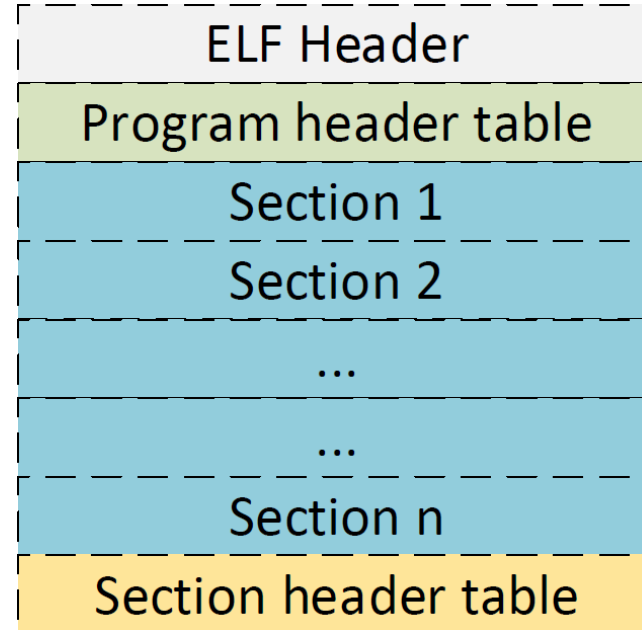
In v7 Debug, it is IMPLEMENTATION DEFINED whether Software debug events are not permitted when the OS Lock is set. See [About invasive debug authentication](#) on page C2-2030.

In v7.1 Debug, Software debug events are not permitted when the OS Lock is set.

CoreSight PTM

Features:

- Trace filter (all code or regions of code)
- Branch Broadcast
- Context ID comparator
- Cycle accurate tracing
- Timestamping



Example trace

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)
```

Static analysis – Tag dependencies

Static Analysis



ARM Assembly Code

Tags Dependencies

@BasicBlock 1

```
00010508 <main>:  
10508: e92d4800  push{fp, lr}  
1050c: e1a0b00d  mov fp, sp  
10510: e24dd030  sub sp, sp, #48 ; 0x30  
10514: e3a00000  mov r0, #0  
10518: e50b0004  str r0, [fp, #-4]  
1051c: ebfff9a  bl 1038c <geteuid@plt>  
10520: e3500000  cmp r0, #0  
10524: 0a00000a  beq 10554 <main+0x4c>
```

@BasicBlock 1

```
tag_of(SP, #4) = tag_of(FP);  
tag_of(SP, #8) = tag_of(LR)  
tag_of(SP) = tag_of(SP, #48)  
tag_of(FP) = tag_of(SP);  
tag_of(r0) = 0  
tag_of(SP, #44) = tag_of(r0)  
KBLARE ... geteuid call
```

DIFT stack

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 loose information (no over-approximation) :
Dynamic approach: Instrumentation + PTM traces
- Extract some informations about the data flow (for tag propagation) :
Static Analysis: Generating annotations.

Overall architecture

```
char buffer1[20], buffer2[20], buffer3[20];  
FILE *fpassword, findex, funauthorized;
```

```
fpassword = open("passwd.txt");  
findex = open("index.html");  
funauthorized = open("unauthorized.html");
```

```
read(buffer1, fpassword)  
read(buffer2, findex)  
read(buffer3, funauthorized)
```

```
if(getuid()){  
    send_to_socket(buffer2);  
}  
else{  
    send_to_socket(buffer3);  
}
```



Modified Clang/LLVM

- **LLVM modification:**
 - Static Analysis
 - Generating annotations
 - Code instrumentation
 - Modifying the ARM Backend (excl. conditional exec, ...)
- **kBlare:** loading annotations into the co-processor dedicated memory.
- **Dynamic linker:** resolving addresses of instrumentation IP address, merging annotations.
- **Yocto:** Generating the complete distribution and the SDK suitable for HardBlare.

Instrumentation

Recover memory addresses

Two possible strategies:

- Recover all memory addresses through instrumentation
- Recover only register-relative memory address through instrumentation

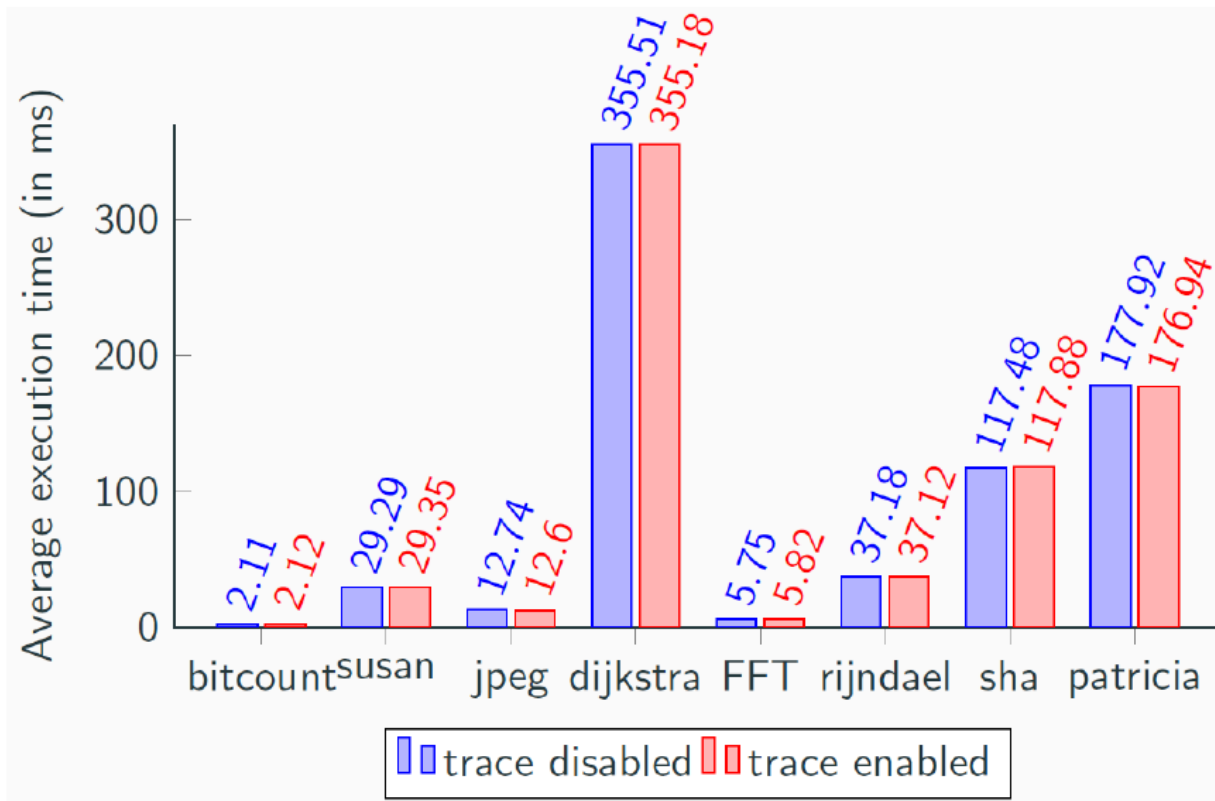
Instrumentation strategy 1

Example Instructions	Tag dependencies	Memory address recovery
sub r0, r1, r2	$\underline{r0} = \underline{r1} + \underline{r2}$	
mov r3, r0	$\underline{r3} = \underline{r0}$	
str r1, [PC, #4]	$\underline{@Mem(PC+4)} = \underline{r1}$	instrumented
ldr r3, [SP, #-8]	$\underline{r3} = \underline{@Mem(SP-8)}$	instrumented
str r1, [r3, r2]	$\underline{@Mem(r3+r2)} = \underline{r1}$	instrumented

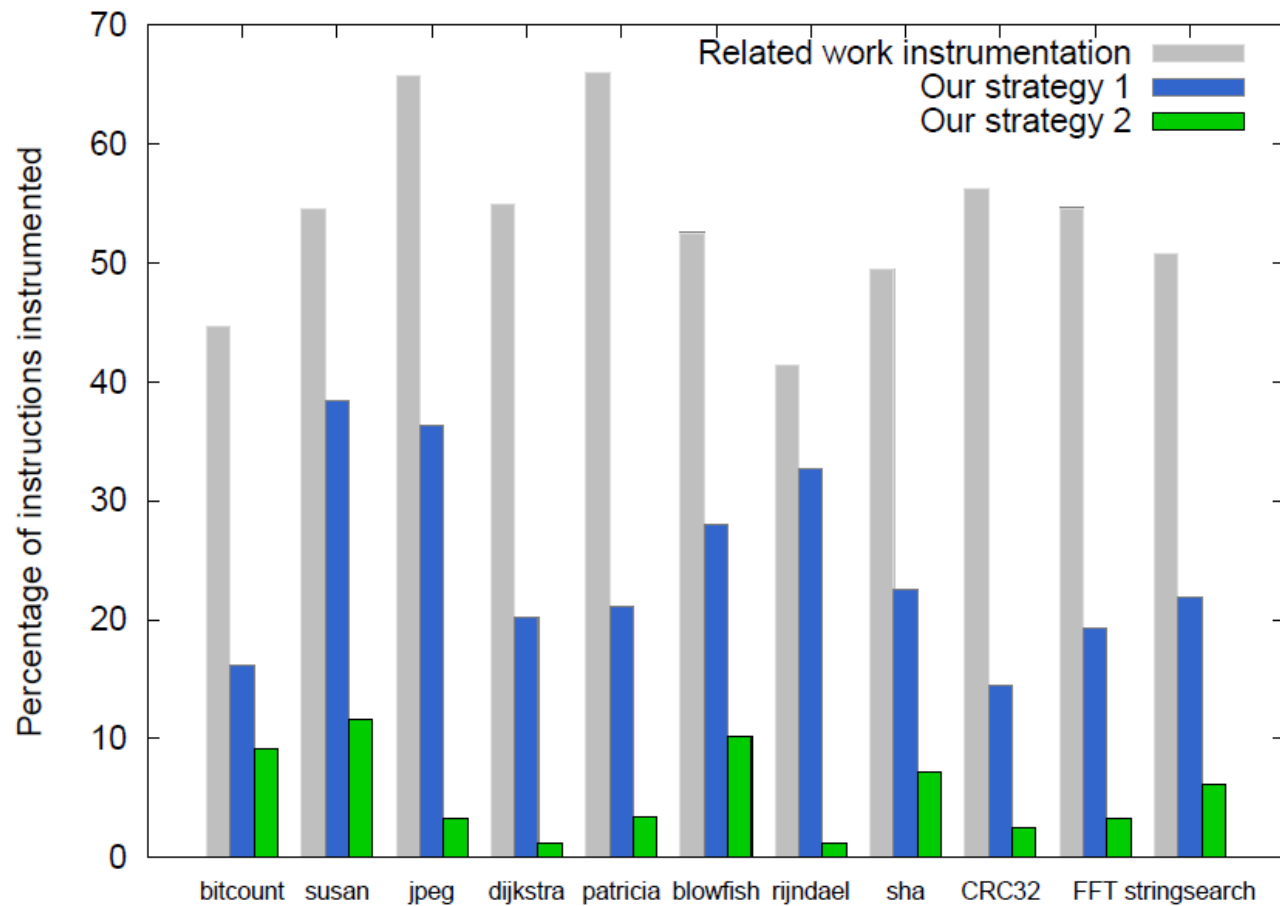
Instrumentation strategy 2

Example Instructions	Tag dependencies	Memory address recovery
sub r0, r1, r2 mov r3, r0 str r1, [PC, #4] ldr r3, [SP, #-8] str r1, [r3, r2]	$\underline{r0} = \underline{r1} + \underline{r2}$ $\underline{r3} = \underline{r0}$ $\underline{@Mem(PC+4)} = \underline{r1}$ $\underline{r3} = \underline{@Mem(SP-8)}$ $\underline{@Mem(r3+r2)} = \underline{r1}$	CoreSight PTM Static analysis instrumented

CoreSight components – Performance overhead



Instrumentation time overhead



Comparison with existing works – Hardware view

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

Conclusion

As well as some perspectives 😊

Conclusion - Perspectives

Take away:

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

Perspectives:

- Mid-term: releasing a PoC of the whole system
- Combine Low-level and OS-level DIFT
- Extend DIFT on multicore CPU
- Take use of other debug components for security
- What about Intel, ST ?



Thanks!

Any questions?

You can find me at: pascal.cotret@gmail.com