BayLibre

Linux Embedded Technology Lab

OPTIMIZE DMA CONFIGURATION IN ENCRYPTION USE CASE

Guillène Ribière, CEO, System Architect

Problem Statement

- Low Performances on Hardware Accelerated Encryption: Max Measured 10MBps
- Expectations: 90 MBps
- Software Based Encryption Measured: 25 MBps

WHY IS HARDWARE ACCELERATED ENCRYPTION SO SLOW?

CONTEXT DESCRIPTION

Choice of Hardware or Software Encryption

BayLibre
Linux Embedded Technology Lab

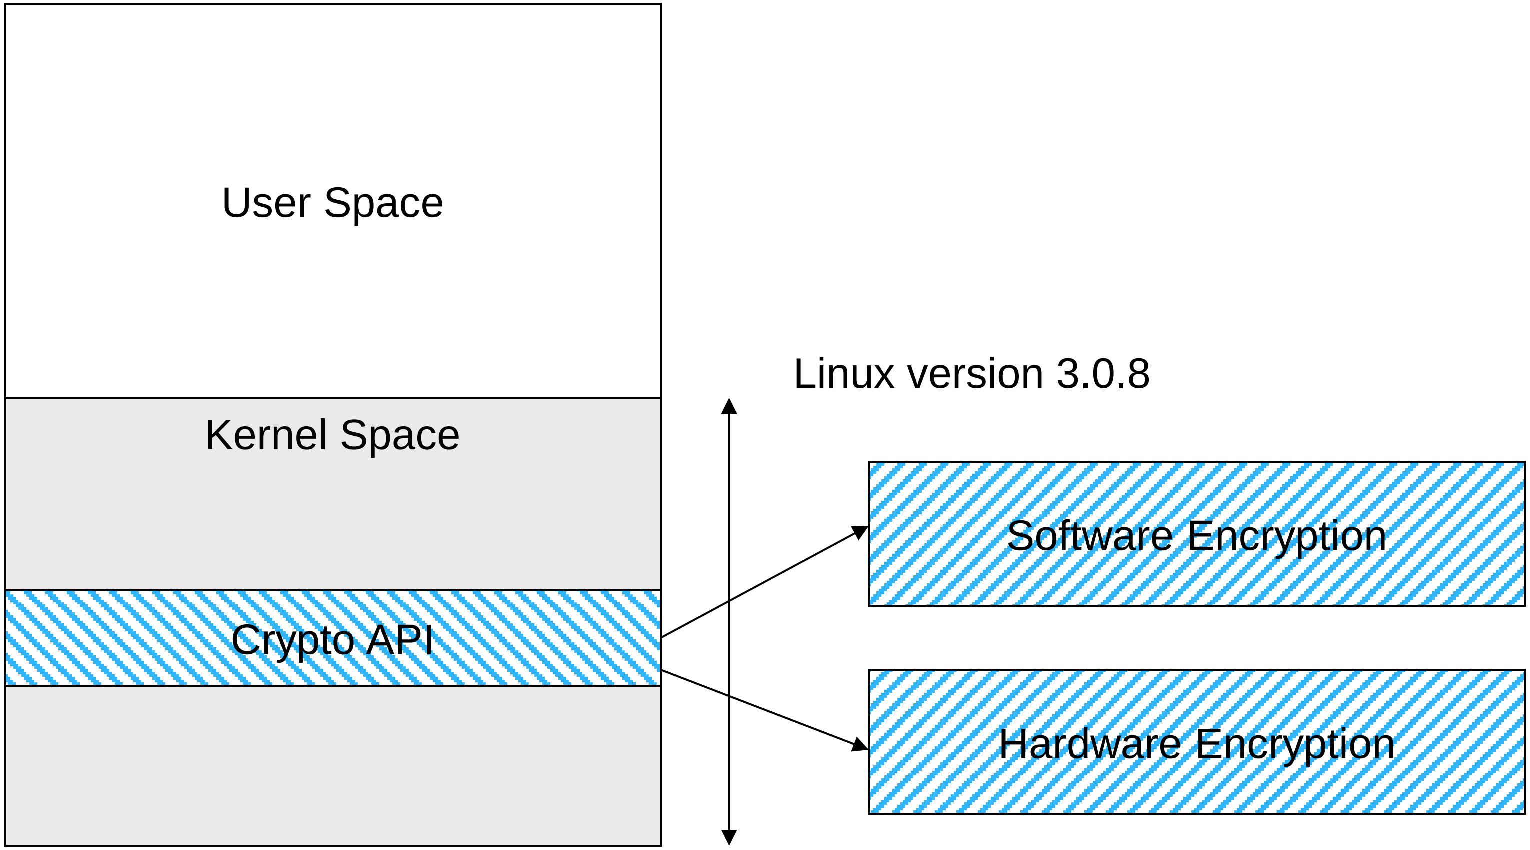

Kernel Knowledge of Encryption Algorithms

BayLibre
Linux Embedded Technology Lab

- Algorithm registration (AES, DES, CBC,…) in kernel,
- cat /proc/crypto shows registered drivers choice:

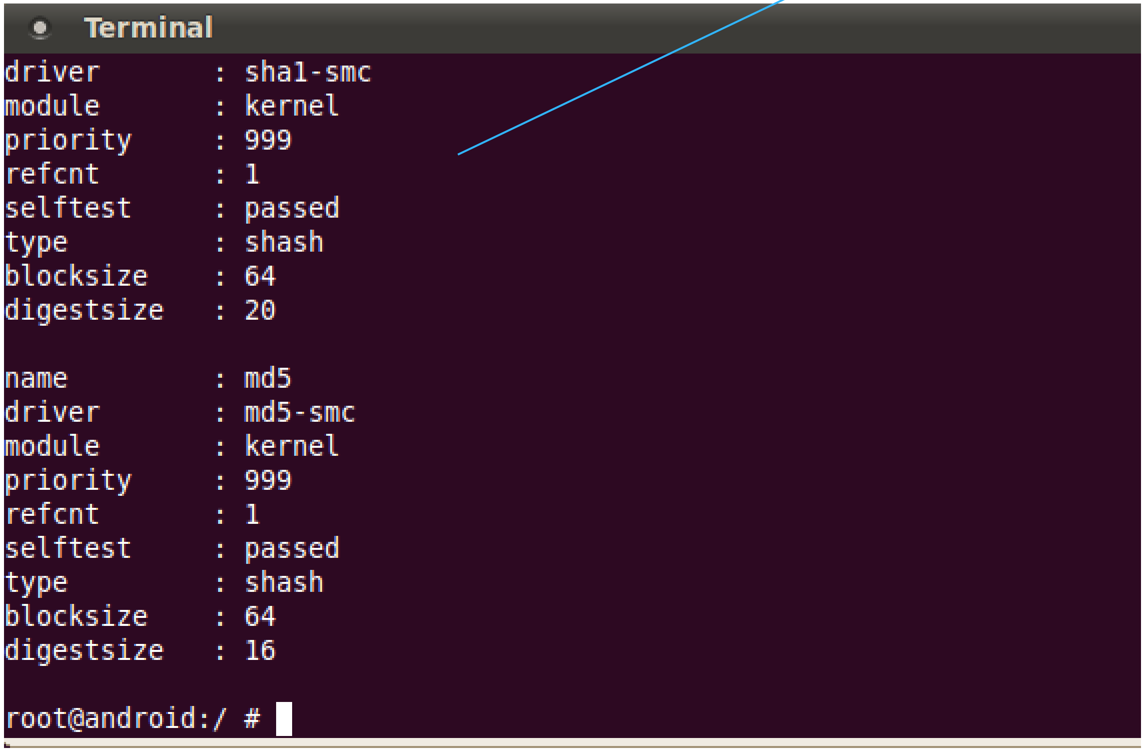

```
driver       : sha1-smc
module       : kernel
priority     : 999
refcnt       : 1
selftest     : passed
type         : shash
blocksize    : 64
digestsize   : 20

name         : md5
driver       : md5-smc
module       : kernel
priority     : 999
refcnt       : 1
selftest     : passed
type         : shash
blocksize    : 64
digestsize   : 16

root@android:/ #
```

Driver registered

Use Case: Public AES Encryption

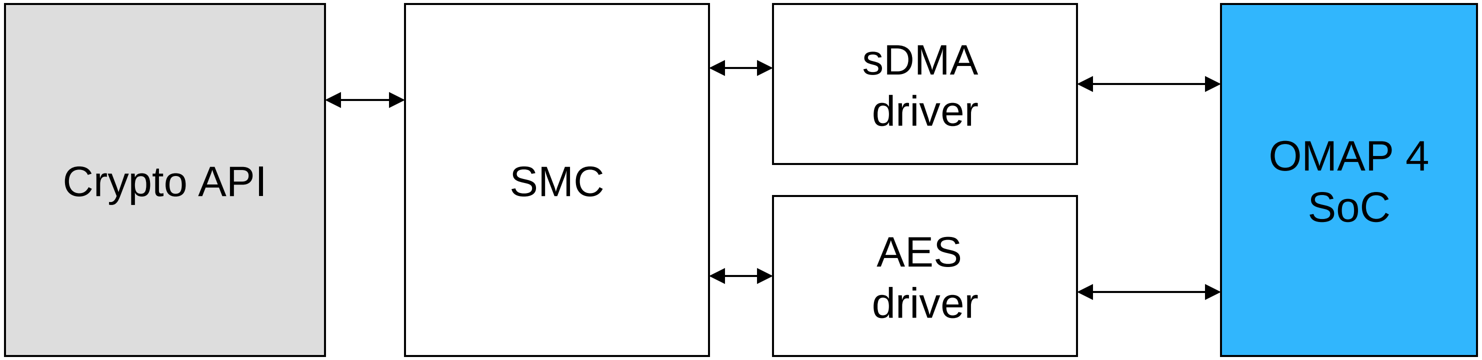

🟦 Hardware

□ Software Specific to OMAP Platform

▨ Software Generic to the Kernel

Use Case AES CBC Public Encryption Flow Single HiB 128-bit Key

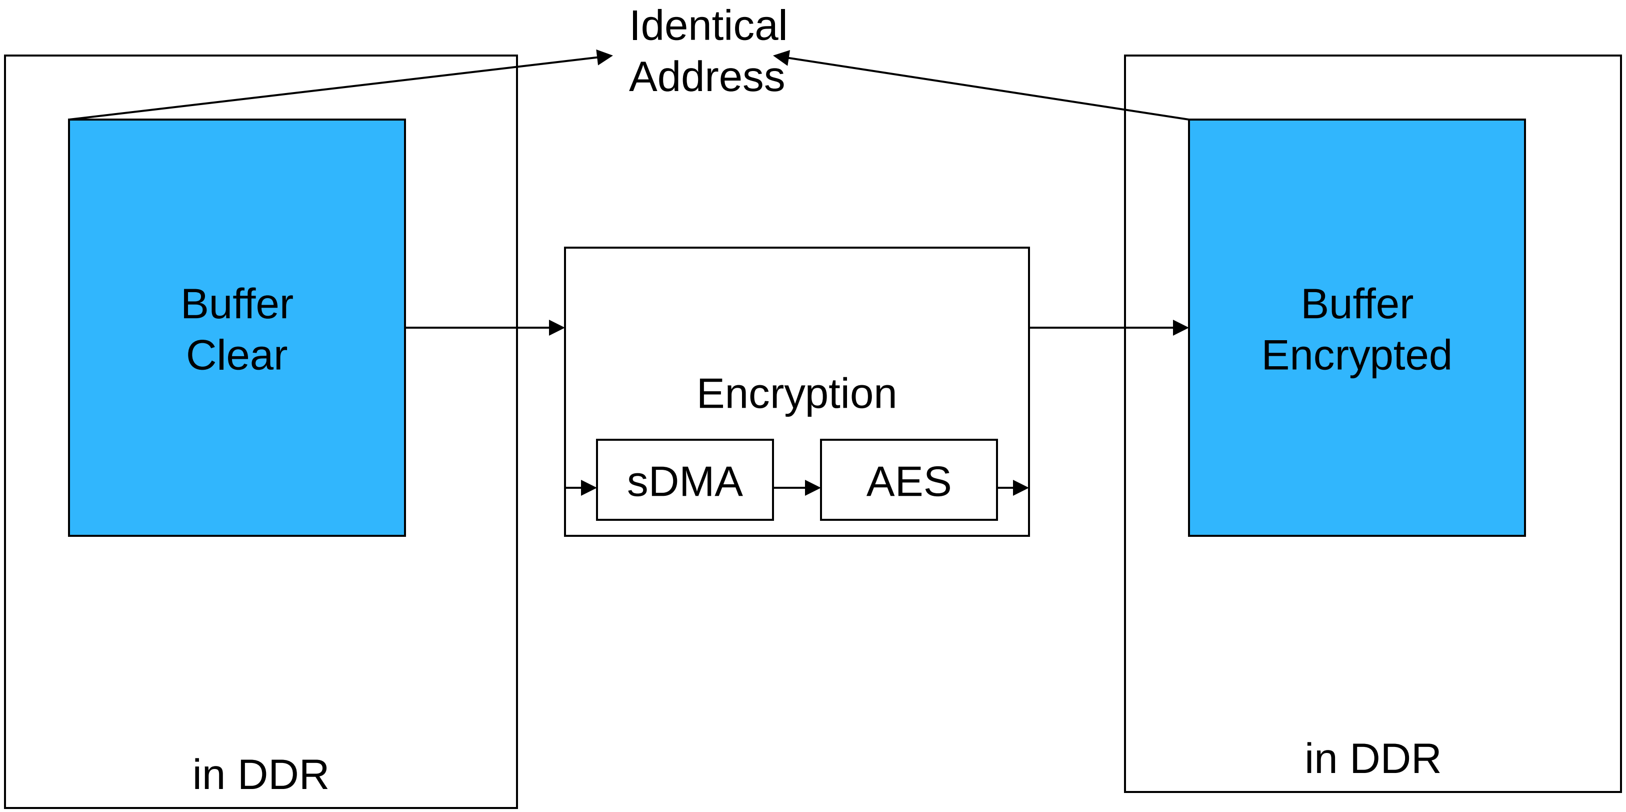

Metric: Number of Buffer Encryptions in 1 Second

Metric: Number of Encryptions Over 1 Second

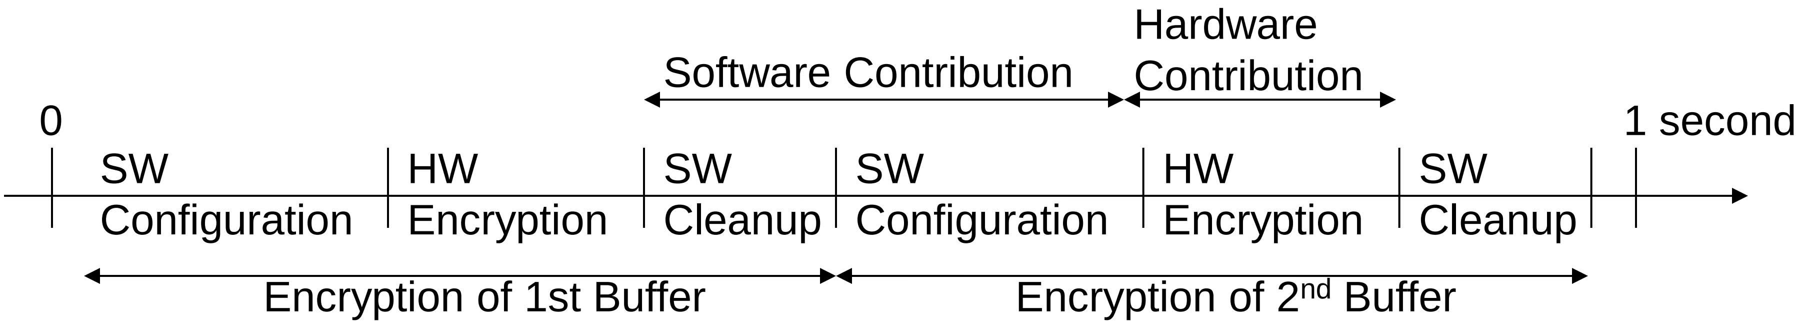

- Buffer Sizes: 64 Bytes / 256 Bytes / 512 Bytes / 1024 Bytes
- AES Block Size: 16 Bytes
- AES Input Buffer: 16 Bytes, Ping and Pong Buffer,
- AES Output Buffer: 16 Bytes, Ping and Pong Buffer.

AES Hardware Diagram

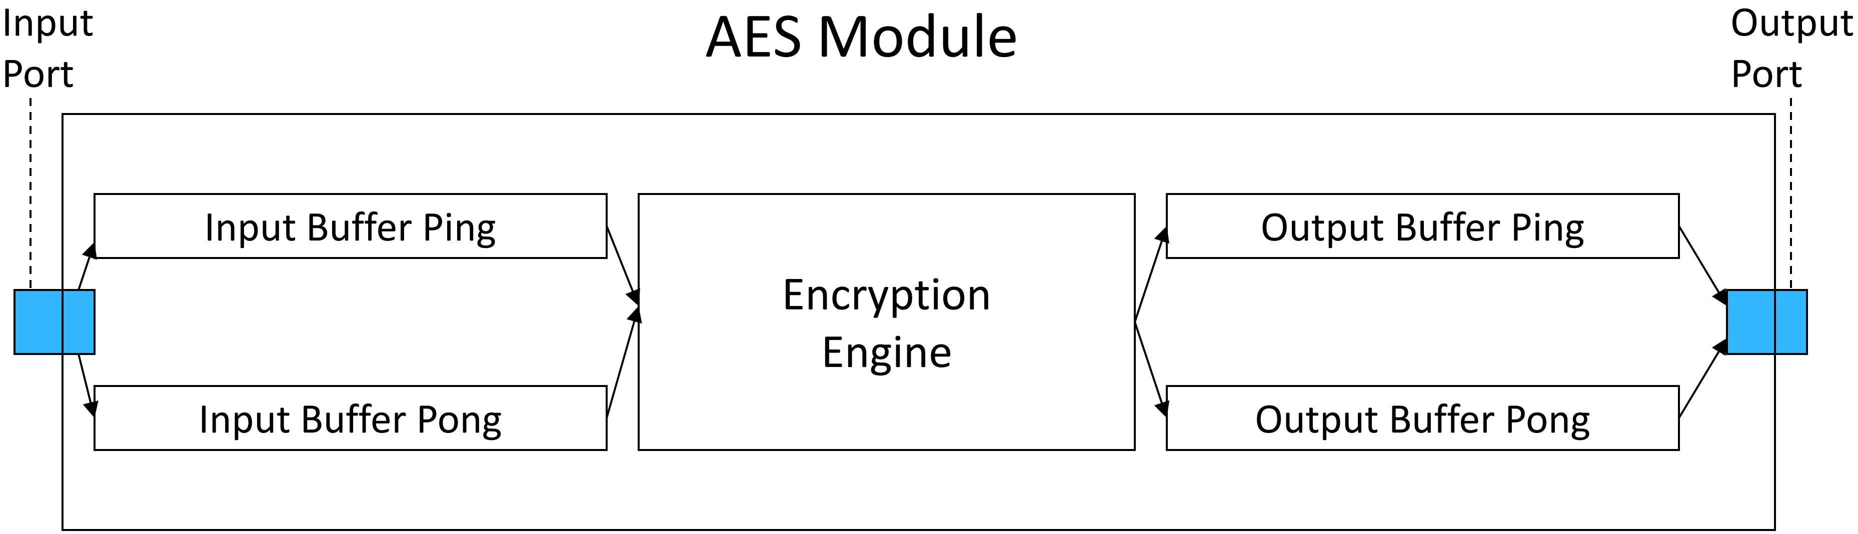

Software Contribution

- Buffer Allocation in cacheable bufferable memory area,
- sDMA configuration
- AES Configuration
- End of Encryption Interrupt Handling

sDMA to AES Data Path: Input Channel

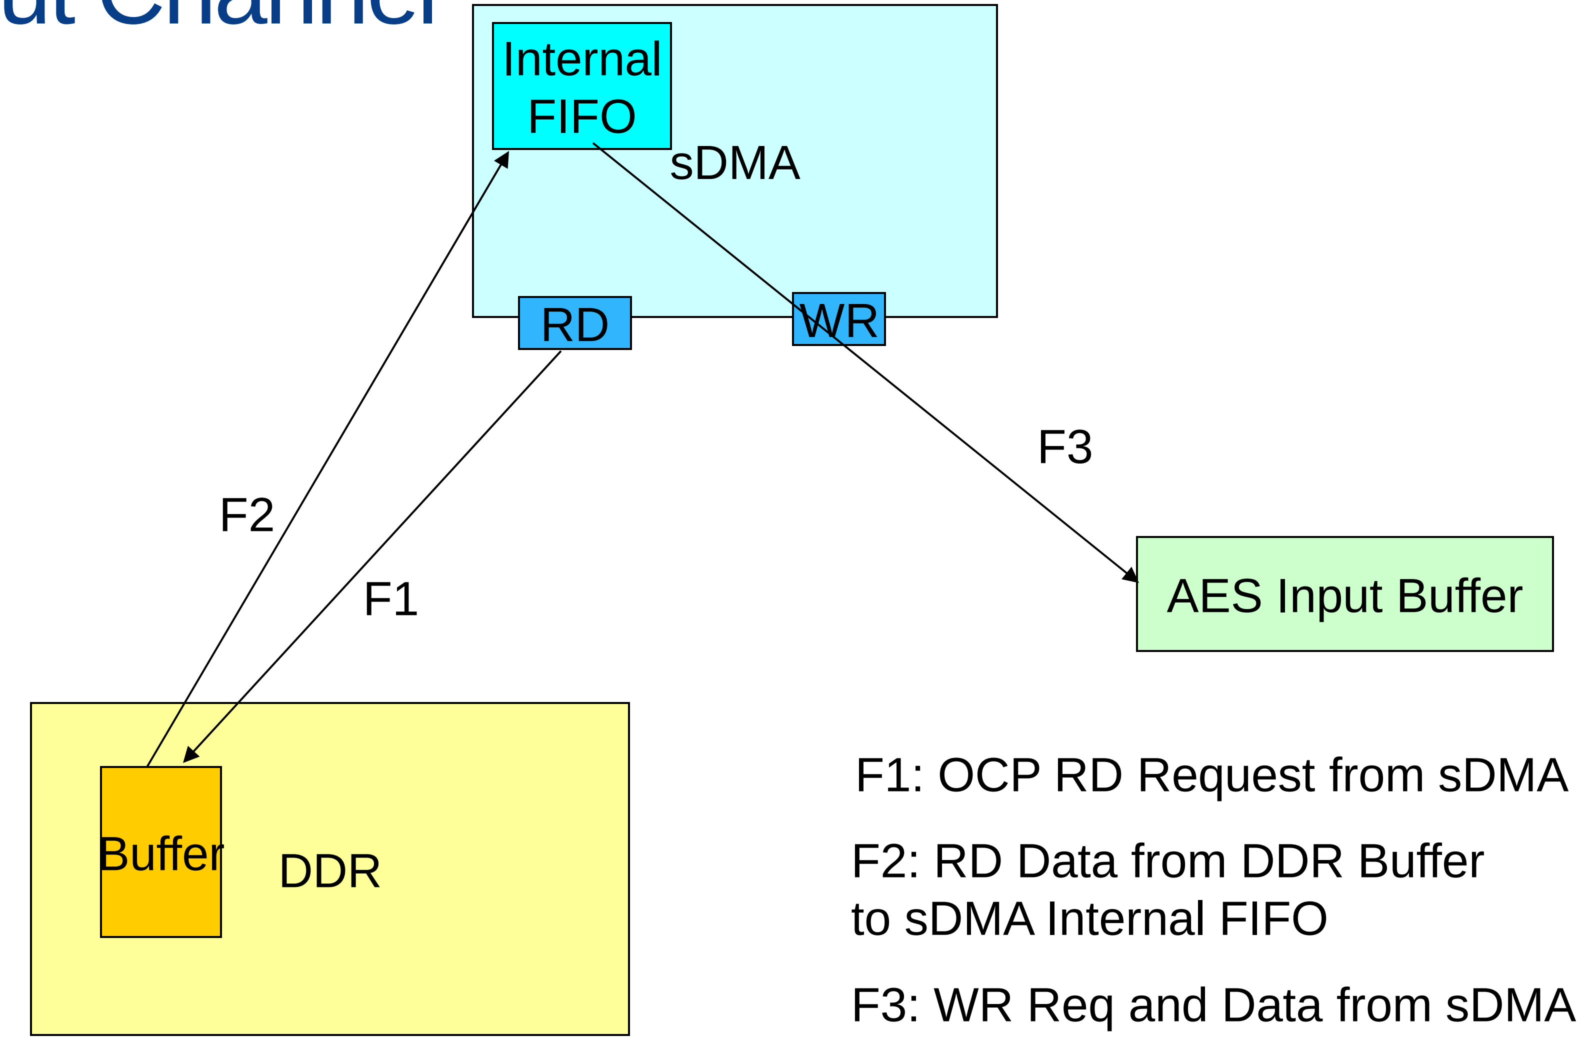

F1: OCP RD Request from sDMA

F2: RD Data from DDR Buffer to sDMA Internal FIFO

F3: WR Req and Data from sDMA Internal FIFO to AES Input Buffer

sDMA to AES Data Path: Output Channel

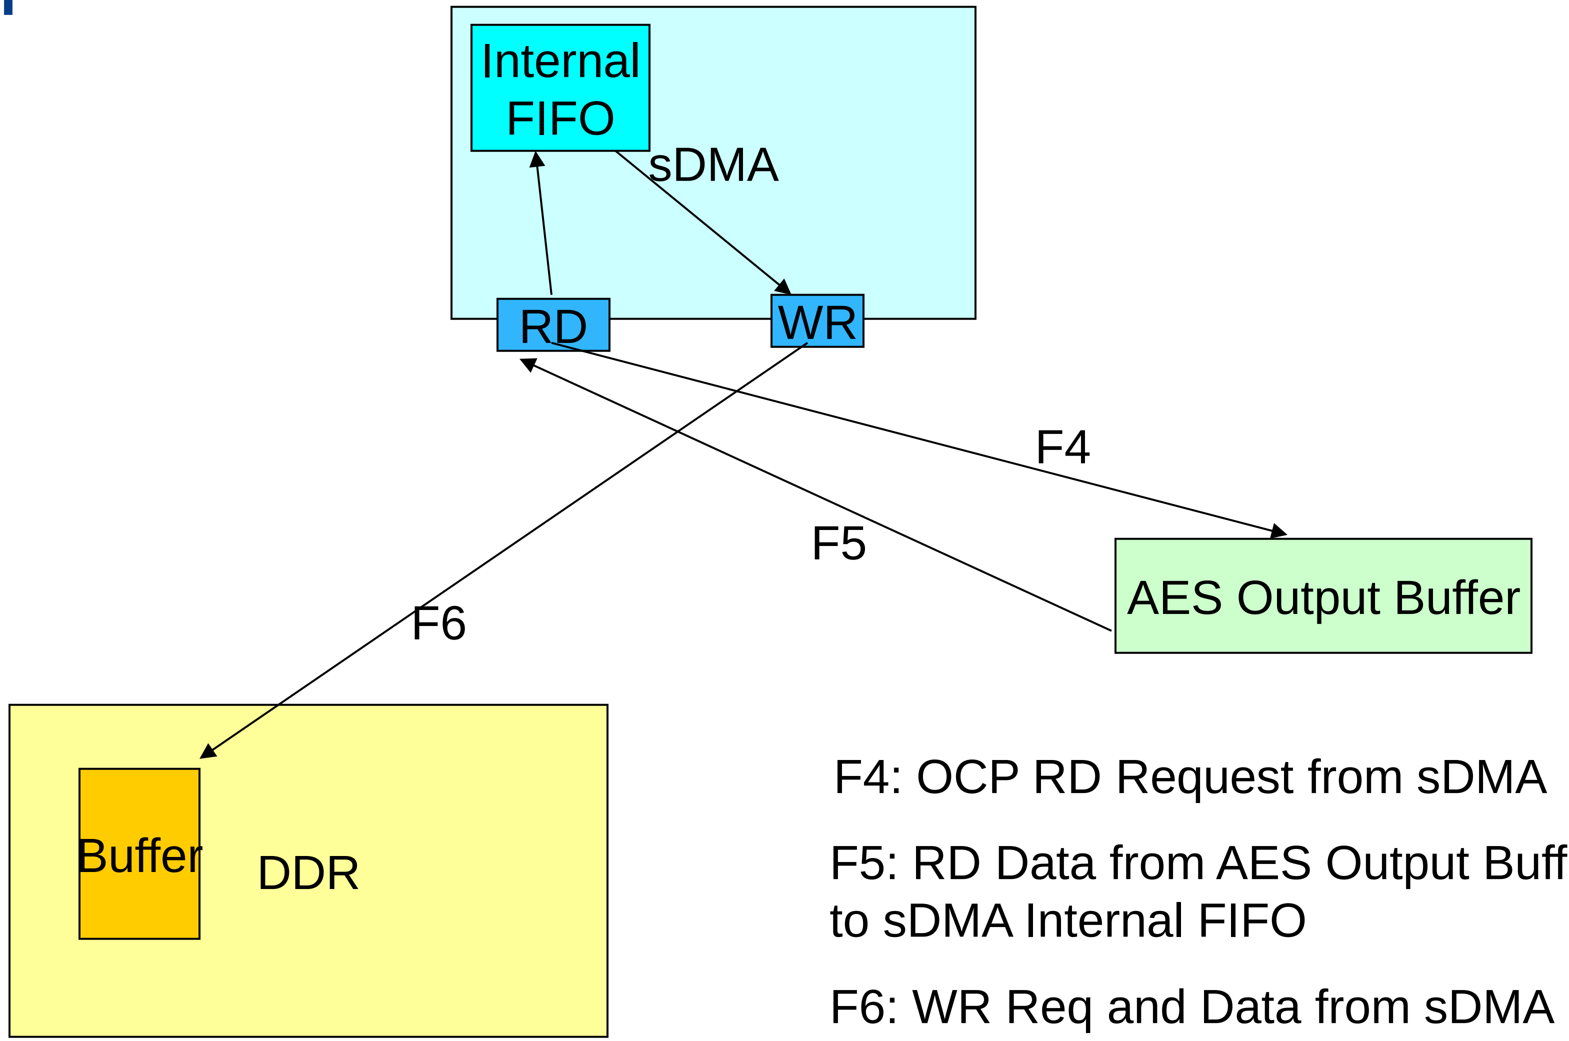

F4: OCP RD Request from sDMA

F5: RD Data from AES Output Buffer to sDMA Internal FIFO

F6: WR Req and Data from sDMA Internal FIFO to DDR Buffer

Single 16 Byte Buffer Encryption: Theory

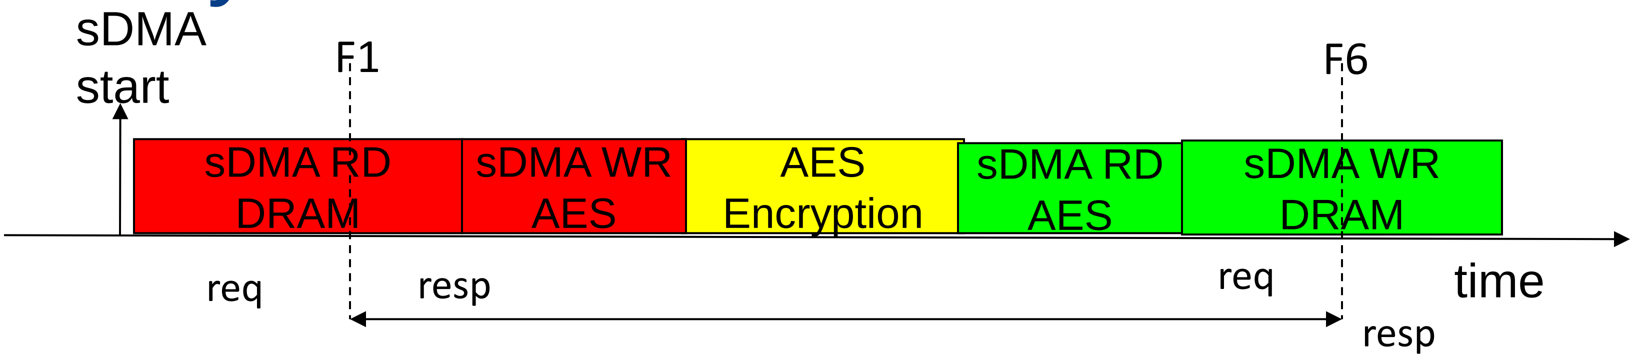

- Latency for sDMA RD expected to be around 50 L3 cycles round trip hence 25 cycles response only, input from simulation,
- sDMA WR to AES in: 20 cycles round trip,
- Latency for AES CBC 16-byte Encryption: 33 L3 cycles,
- sDMA RD to AES out: 20 cycles round trip,
- Latency for sDMA WR expected to be around 50 L3 cycles round trip hence 25 cycles response only
- **Total Latency Expected for Single 16 Byte Block Encryption 123 L3 cycles at L3 target agent to DMM Boundary, ballpark figure.**

Theory:
64 Byte Block Encryption = 4x16 Byte Bursts

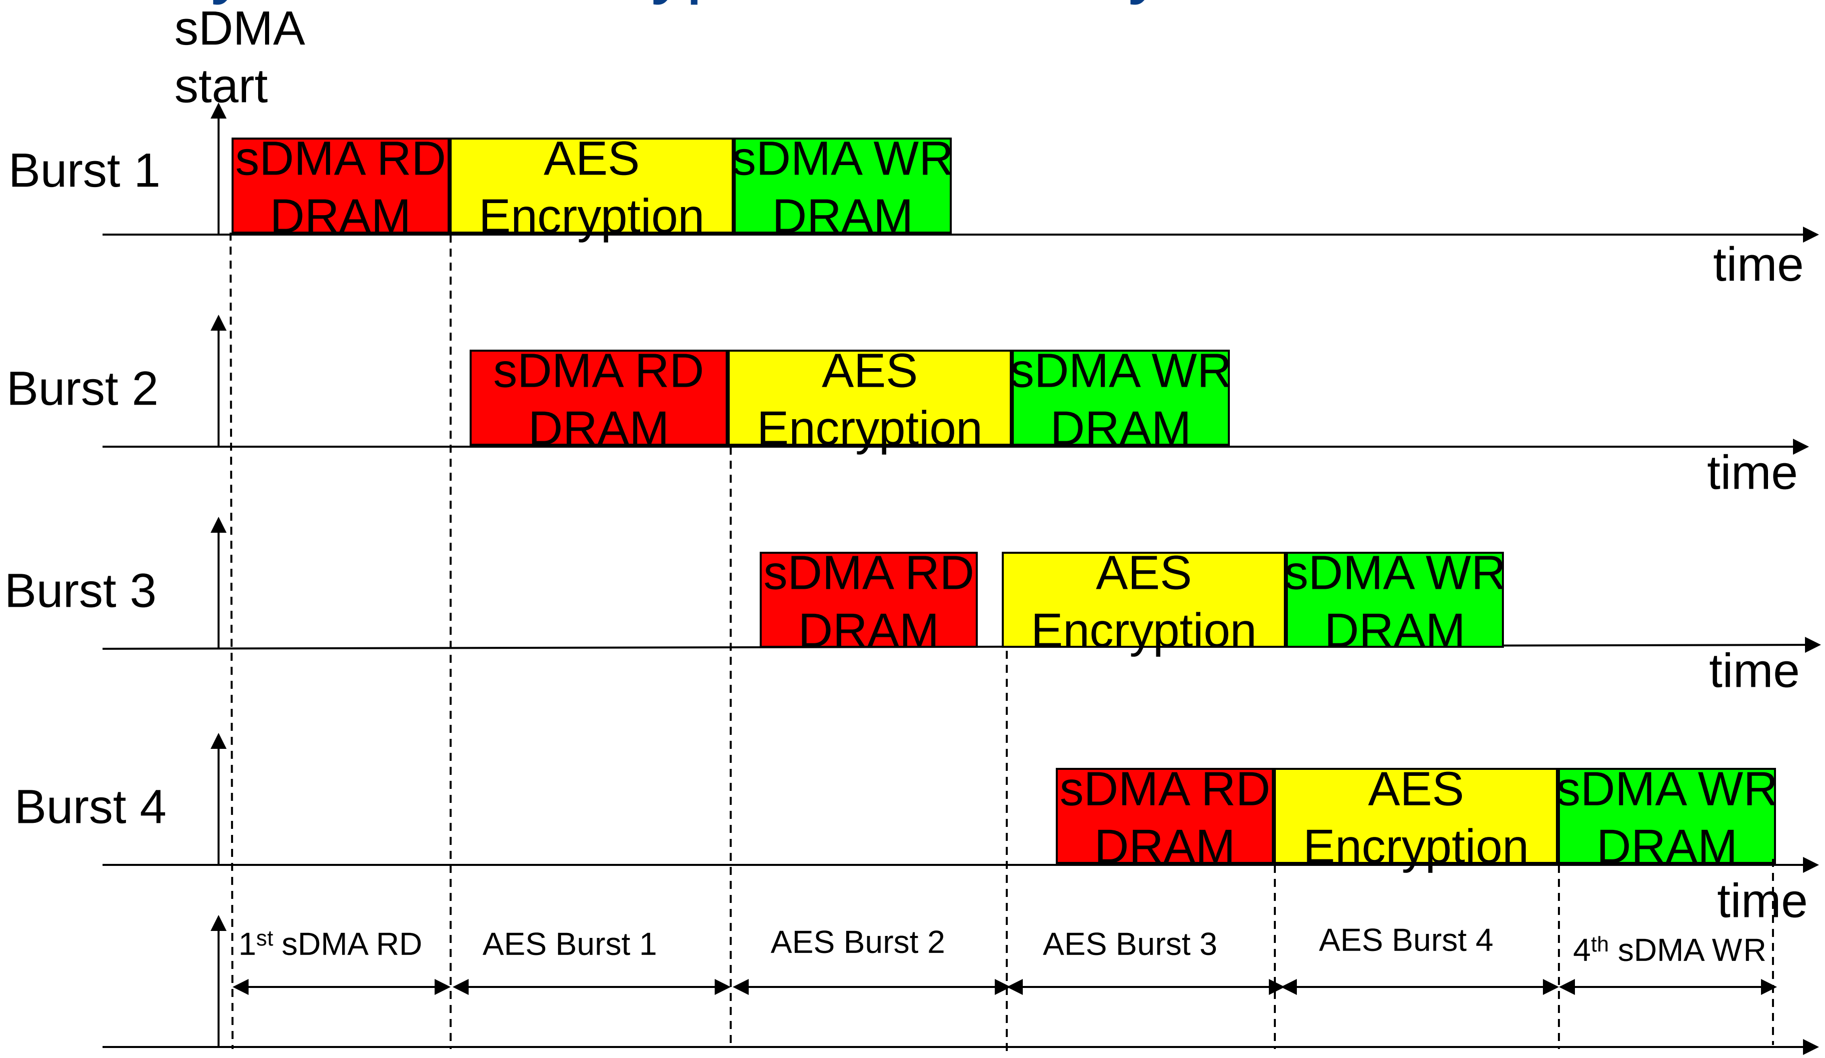

Theoretical Throughput: Expectations

- SW overhead negligible
- Latencies to and from DDR hidden by pipelining
- Throughput should be close to 96MBps with L3@200MHz:
  - 33 L3 cycles for AES CBC encrypt
  - 16 Bytes per 165 ns (33 * 5 ns)
- For small buffer add cost of initial request and last request to DDR

Buffer size (Byte)	Theory (L3 cycles)	Theory Throughput (MBps)
16	123	26
64	222	57
256	618	82
512	1146	89
1024	2202	93

ON BOARD ANALYSIS

Default Configuration

Environment

- Blaze SEVM OMAP 4460 ES 1.1 HS
- Ice Cream Sandwich Daily Build 384
- MSHIELD-DK-LITE v1.7.5
- OPP 100
- MPU@700MHz, L3@200MHz
- Basic OS and Screen (On and OFF) Activity on Platform

Measurements Default Configuration

	64 Byte Buffer	256 Byte Buffer	512 Byte Buffer	1024 Byte Buffer
Number of Buffer Encryptions per Second	10278	10065	8377	7625
Time for a single Buffer Encryption (us)	97.29	99.35	119	131
Throughput (MBps)	0.65	2.57	4.28	7.8

OCP Watchpoint

- What is it?
  - Hardware Probes Logging OCP Transactions
- What Information can they Extract?
  - Transaction Type: RD/WR/WRNP
  - Address
  - Initiator
  - Time of Transaction Occurrence
- Where are they?
  - DDR Boundary, L4, GPMC

Actual Hardware Duration Measured !

- Not Measured from MPU Perspective
- Measurement Made Using HW Instrumentation
- OCP Watchpoints embedded in L3 used
- OCP Watchpoint Probe to SDRAM used

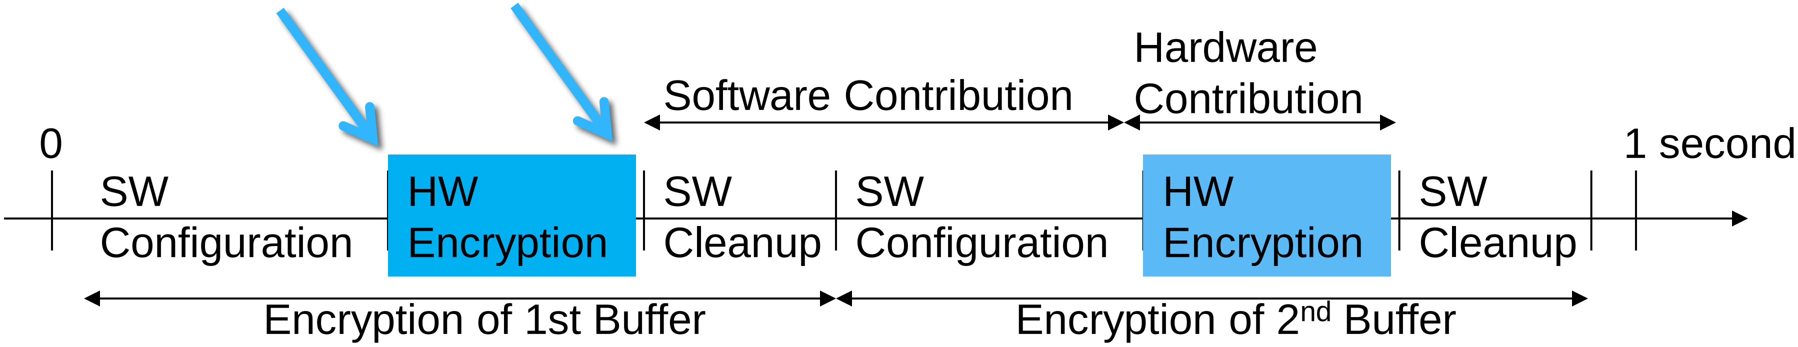

Single 16 Byte Buffer Encryption: Reality

BayLibre
Linux Embedded Technology Lab

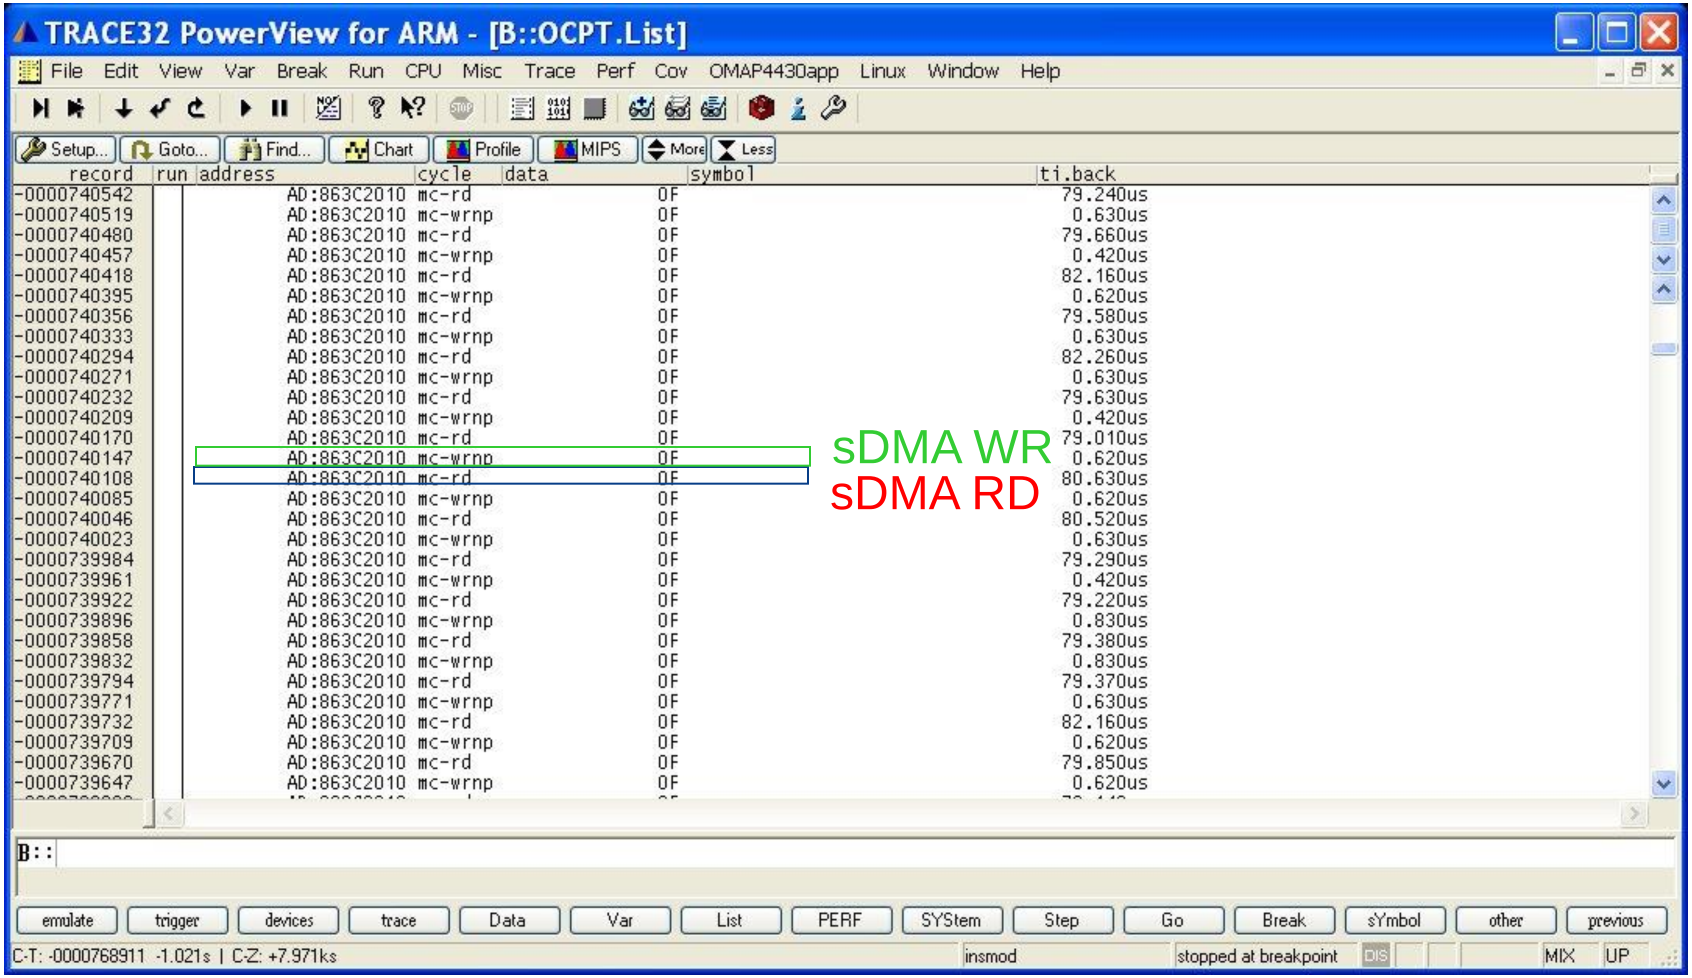

```
      record run address            cycle   data                 symbol                      ti.back
-0000740542            AD:863C2010 mc-rd                 0F                                  79.240us
-0000740519            AD:863C2010 mc-wrnp               0F                                   0.630us
-0000740480            AD:863C2010 mc-rd                 0F                                  79.660us
-0000740457            AD:863C2010 mc-wrnp               0F                                   0.420us
-0000740418            AD:863C2010 mc-rd                 0F                                  82.160us
-0000740395            AD:863C2010 mc-wrnp               0F                                   0.620us
-0000740356            AD:863C2010 mc-rd                 0F                                  79.580us
-0000740333            AD:863C2010 mc-wrnp               0F                                   0.630us
-0000740294            AD:863C2010 mc-rd                 0F                                  82.260us
-0000740271            AD:863C2010 mc-wrnp               0F                                   0.630us
-0000740232            AD:863C2010 mc-rd                 0F                                  79.630us
-0000740209            AD:863C2010 mc-wrnp               0F                                   0.420us
-0000740170            AD:863C2010 mc-rd                 0F                                  79.010us
-0000740147            AD:863C2010 mc-wrnp               0F                                   0.620us
-0000740108            AD:863C2010 mc-rd                 0F                                  80.630us
-0000740085            AD:863C2010 mc-wrnp               0F                                   0.620us
-0000740046            AD:863C2010 mc-rd                 0F                                  80.520us
-0000740023            AD:863C2010 mc-wrnp               0F                                   0.630us
-0000739984            AD:863C2010 mc-rd                 0F                                  79.290us
-0000739961            AD:863C2010 mc-wrnp               0F                                   0.420us
-0000739922            AD:863C2010 mc-rd                 0F                                  79.220us
-0000739896            AD:863C2010 mc-wrnp               0F                                   0.830us
-0000739858            AD:863C2010 mc-rd                 0F                                  79.380us
-0000739832            AD:863C2010 mc-wrnp               0F                                   0.830us
-0000739794            AD:863C2010 mc-rd                 0F                                  79.370us
-0000739771            AD:863C2010 mc-wrnp               0F                                   0.630us
-0000739732            AD:863C2010 mc-rd                 0F                                  82.160us
-0000739709            AD:863C2010 mc-wrnp               0F                                   0.620us
-0000739670            AD:863C2010 mc-rd                 0F                                  79.850us
-0000739647            AD:863C2010 mc-wrnp               0F                                   0.620us

B::

C-T: -0000768911  -1.021s | C-Z: +7.971ks        insmod        stopped at breakpoint   DIS   MIX   UP
```

sDMA WR

sDMA RD

Interpretation of OCP WP

- Path to End OCP Trace: L3->DebugSS->STP->PTI->Lauterbach
- Best case ~0.21 us between 2 traced ocp requests
- Big gaps are better represented than small ones
- When OCP Transactions Throughput > Throughput of OCP WP => overflow

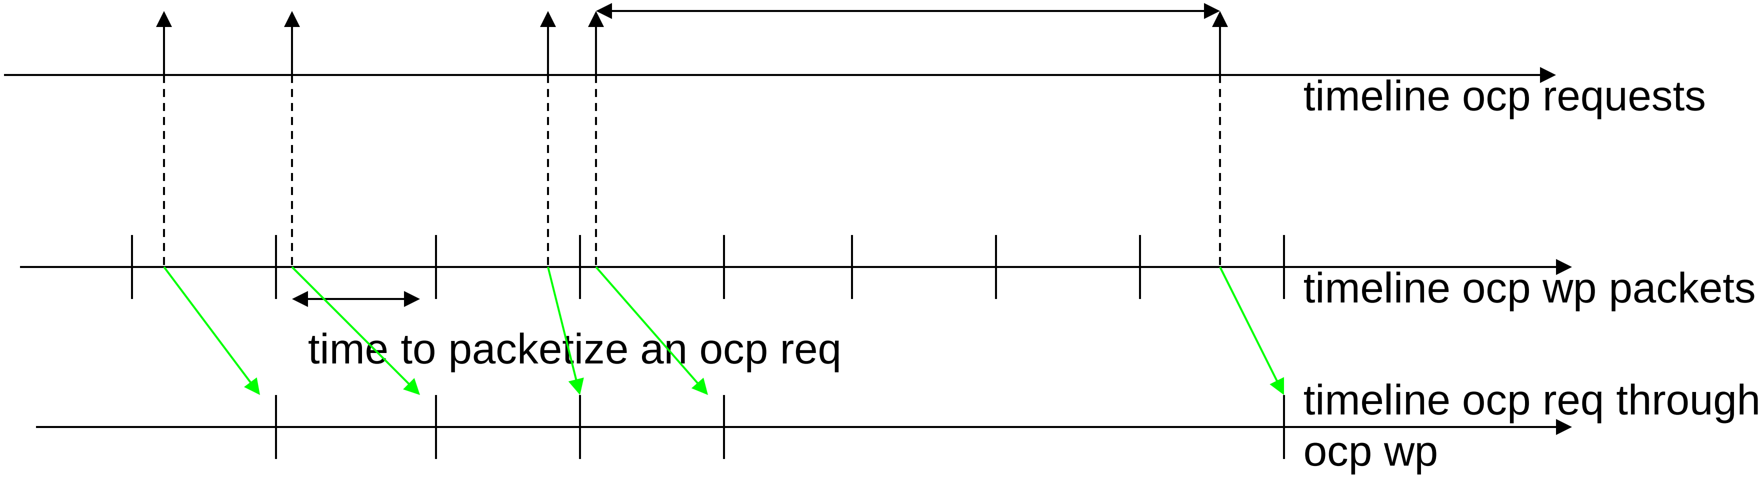

256 Byte Buffer Encryption: Reality

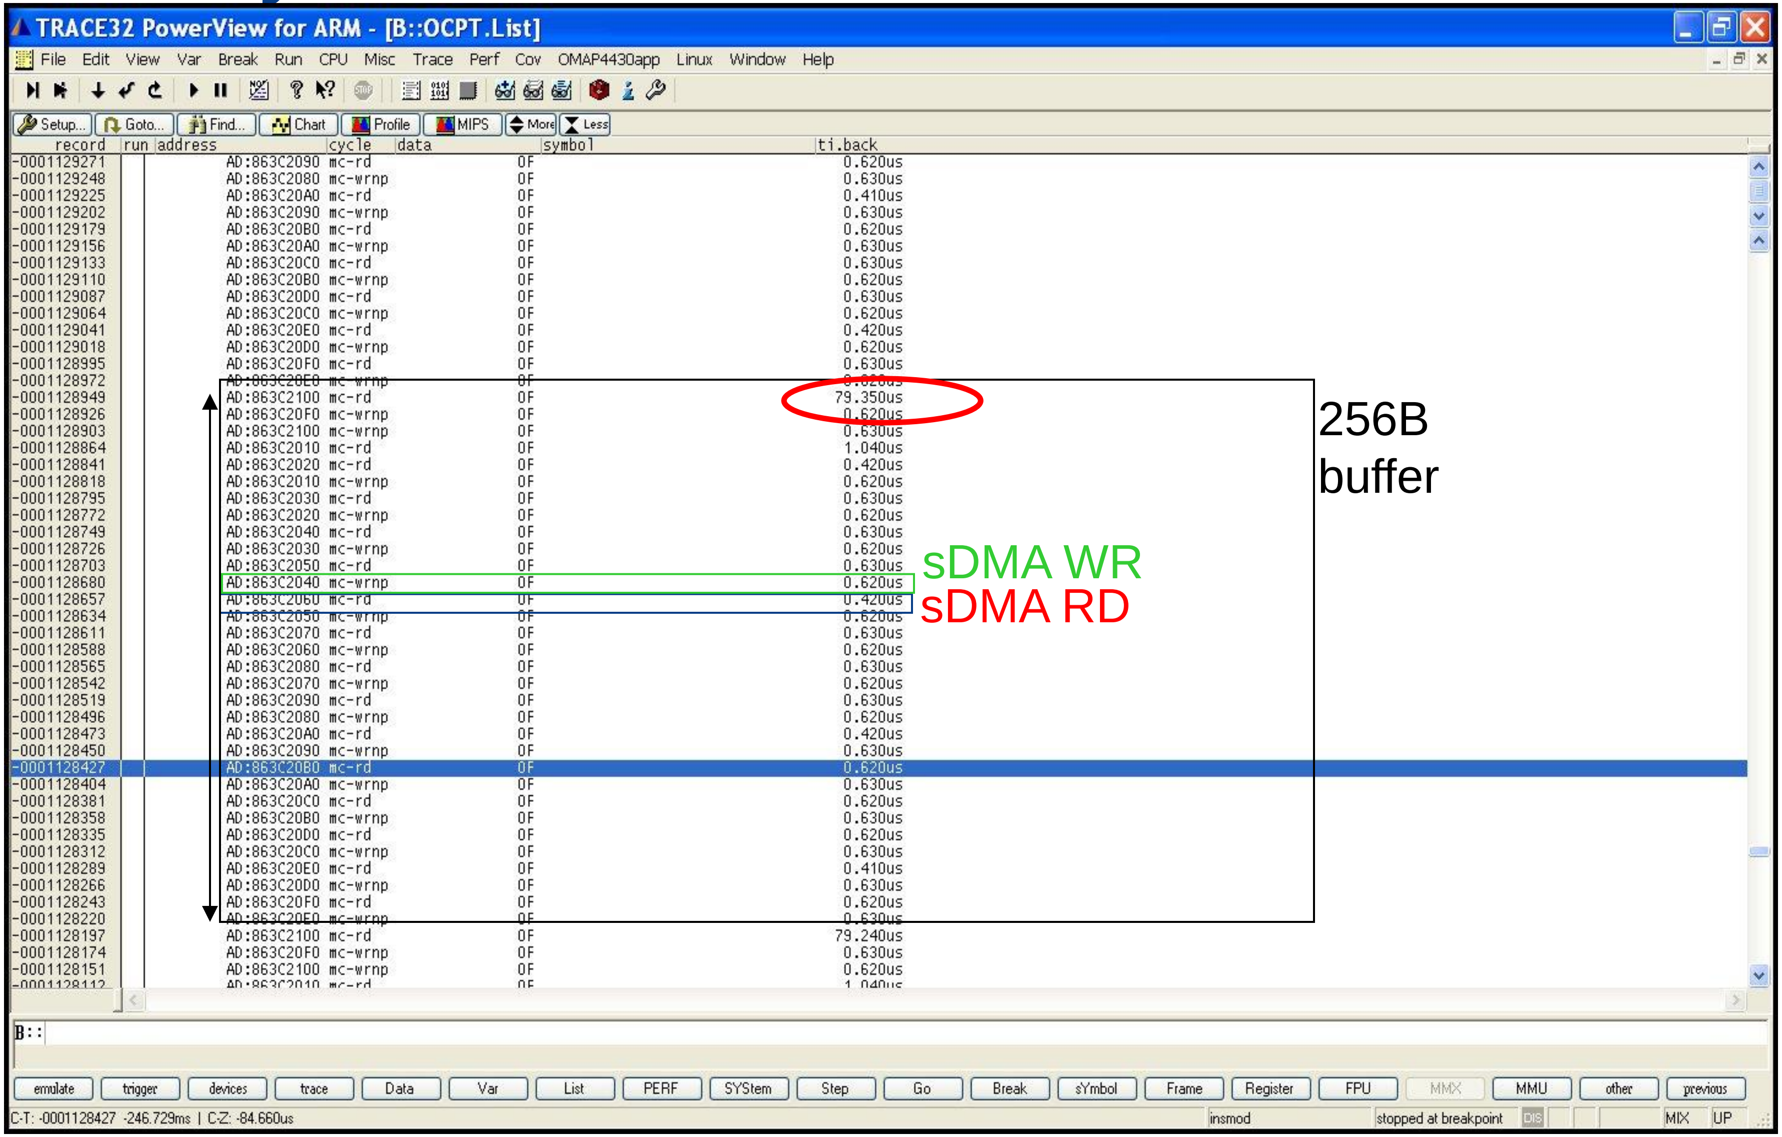
TRACE32 PowerView for ARM - [B::OCPT.List]
File Edit View Var Break Run CPU Misc Trace Perf Cov OMAP4430app Linux Window Help
Setup... Goto... Find... Chart Profile MIPS More Less
B::
emulate trigger devices trace Data Var List PERF SYStem Step Go Break sYmbol Frame Register FPU MMX MMU other previous

```
   record  run address           cycle   data         symbol                     ti.back
-0001129271        AD:863C2090 mc-rd        0F                                   0.620us
-0001129248        AD:863C2080 mc-wrnp      0F                                   0.630us
-0001129225        AD:863C20A0 mc-rd        0F                                   0.410us
-0001129202        AD:863C2090 mc-wrnp      0F                                   0.630us
-0001129179        AD:863C20B0 mc-rd        0F                                   0.620us
-0001129156        AD:863C20A0 mc-wrnp      0F                                   0.630us
-0001129133        AD:863C20C0 mc-rd        0F                                   0.630us
-0001129110        AD:863C20B0 mc-wrnp      0F                                   0.620us
-0001129087        AD:863C20D0 mc-rd        0F                                   0.630us
-0001129064        AD:863C20C0 mc-wrnp      0F                                   0.620us
-0001129041        AD:863C20E0 mc-rd        0F                                   0.420us
-0001129018        AD:863C20D0 mc-wrnp      0F                                   0.620us
-0001128995        AD:863C20F0 mc-rd        0F                                   0.630us
-0001128972        AD:863C20E0 mc-wrnp      0F                                   0.620us
-0001128949        AD:863C2100 mc-rd        0F                                  79.350us
-0001128926        AD:863C20F0 mc-wrnp      0F                                   0.620us
-0001128903        AD:863C2100 mc-wrnp      0F                                   0.630us
-0001128864        AD:863C2010 mc-rd        0F                                   1.040us
-0001128841        AD:863C2020 mc-rd        0F                                   0.420us
-0001128818        AD:863C2010 mc-wrnp      0F                                   0.620us
-0001128795        AD:863C2030 mc-rd        0F                                   0.630us
-0001128772        AD:863C2020 mc-wrnp      0F                                   0.620us
-0001128749        AD:863C2040 mc-rd        0F                                   0.630us
-0001128726        AD:863C2030 mc-wrnp      0F                                   0.620us
-0001128703        AD:863C2050 mc-rd        0F                                   0.630us
-0001128680        AD:863C2040 mc-wrnp      0F                                   0.620us
-0001128657        AD:863C2060 mc-rd        0F                                   0.420us
-0001128634        AD:863C2050 mc-wrnp      0F                                   0.620us
-0001128611        AD:863C2070 mc-rd        0F                                   0.630us
-0001128588        AD:863C2060 mc-wrnp      0F                                   0.620us
-0001128565        AD:863C2080 mc-rd        0F                                   0.630us
-0001128542        AD:863C2070 mc-wrnp      0F                                   0.620us
-0001128519        AD:863C2090 mc-rd        0F                                   0.630us
-0001128496        AD:863C2080 mc-wrnp      0F                                   0.620us
-0001128473        AD:863C20A0 mc-rd        0F                                   0.420us
-0001128450        AD:863C2090 mc-wrnp      0F                                   0.630us
-0001128427        AD:863C20B0 mc-rd        0F                                   0.620us
-0001128404        AD:863C20A0 mc-wrnp      0F                                   0.630us
-0001128381        AD:863C20C0 mc-rd        0F                                   0.620us
-0001128358        AD:863C20B0 mc-wrnp      0F                                   0.630us
-0001128335        AD:863C20D0 mc-rd        0F                                   0.620us
-0001128312        AD:863C20C0 mc-wrnp      0F                                   0.630us
-0001128289        AD:863C20E0 mc-rd        0F                                   0.410us
-0001128266        AD:863C20D0 mc-wrnp      0F                                   0.630us
-0001128243        AD:863C20F0 mc-rd        0F                                   0.620us
-0001128220        AD:863C20E0 mc-wrnp      0F                                   0.630us
-0001128197        AD:863C2100 mc-rd        0F                                  79.240us
-0001128174        AD:863C20F0 mc-wrnp      0F                                   0.630us
-0001128151        AD:863C2100 mc-wrnp      0F                                   0.620us
-0001128112        AD:863C2010 mc-rd        0F                                   1.040us
```

```
C-T: -0001128427 -246.729ms | C-Z: -84.660us        insmod        stopped at breakpoint        MIX  UP
```

256B buffer

sDMA WR

sDMA RD

Interpretation of OCP WP Trace: Differentiation Between SW Contribution and HW Contribution

- **SW Contribution ~ 80 us**
- SW Contribution Big, Measurement Through OCP WP Relevant
- HW Contribution: the more transaction, the more the average is relevant
- 1024 Byte Buffer provokes OCP WP Overflow
- Trace shows that RD and WR requests alternate one to one
- sDMA prefetch not enabled

sDMA Input Channel: Reality with 256 Bytes Buffer

```
      sample #          ocp transaction    ocp req
                        hex address        type
-0012437933 | |   AD:85C84100 mc-wrnp     Last Transaction Previous Block
-0012437894 | |   AD:85C84010 mc-rd       First Transaction Current Block
                                          RD Burst 1 to ping input buffer
-0012437871 | |   AD:85C84020 mc-rd       RD Burst 2 to pong input buffer
-0012437848 | |   AD:85C84010 mc-wrnp     WR Burst 1
-0012437825 | |   AD:85C84030 mc-rd       RD Burst 3
-0012437802 | |   AD:85C84020 mc-wrnp     WR Burst 2
-0012437779 | |   AD:85C84040 mc-rd
-0012437756 | |   AD:85C84030 mc-wrnp
```

Trace Extracted Through OCP WP Activated on sDMA RD and sDMA WR to DDR

Measurements Default Configuration (2)

	64 Byte Buffer	256 Byte Buffer	512 Byte Buffer	1024 Byte Buffer
Number of Buffer Encryptions per Second	10278	10065	8377	7625
Time for a single Buffer Encryption (us)	97.29	99.35	119	131
Throughput (MBps)	0.65	2.57	4.28	7.8
Hardware Throughput (MBps)*	3.7	13.23	13	20

*Buffer size / (time per Buffer – 80us)

*16 byte buffer jittery measurement

BayLibre
Linux Embedded Technology Lab

SDMA CONFIGURATION MODIFICATION

Goal: Improving Hardware Contribution

sDMA Configuration Modification

- Prefetch enabled
- Logical Channel Fifo Size Increase
- Move from Write posted to Write posted with last non posted
- Setup stays Identical

sDMA Input Channel Config: Prefetch ON and FIFO size Increased with 256 Bytes Buffer

```
                         ocp transaction       ocp req
  sample #               hex address           type

-0009281077 | |       AD:863C8100 mc-wrnp   ←  Last Transaction Previous Block
-0009281038 | |       AD:863C8010 mc-rd     ←  RD Burst 1
-0009281015 | |       AD:863C8020 mc-rd     ←  RD Burst 2
-0009280992 | |       AD:863C8030 mc-rd     ←  RD Burst 3
-0009280969 | |       AD:863C8040 mc-rd     ←  RD Burst 4
-0009280946 | |       AD:863C8050 mc-rd     ←  RD Burst 5
-0009280923 | |       AD:863C8060 mc-rd     ←  RD Burst 6
-0009280900 | |       AD:863C8010 mc-wrnp   ←  WR Burst 1
-0009280877 | |       AD:863C8070 mc-rd     ←  RD Burst 7
-0009280854 | |       AD:863C8020 mc-wrnp   ←  WR Burst 2
-0009280831 | |       AD:863C8080 mc-rd
-0009280808 | |       AD:863C8030 mc-wrnp
```

Trace Extracted Through OCP WP Activated on sDMA RD and sDMA WR to DDR

Interpretation of OCP WP Trace Prefetch ON

- 6 RD Transactions at start of Buffer Encryption
- 2 RD Transactions go into AES Input Buffer: Ping and Pong
- 4 are stored in sDMA FIFO
- Address Difference between RD and WR shows that sDMA contains Data to write to AES in advance

Raw Results with **Prefetch On**

sDMA FIFO in 64-bit words	Prefetch	Tcrypt: number of buffers per second (always same conditions)			
		64 B Buffer	256 B Buffer	512 B Buffer	1024 B Buffer
16	OFF	10278	10065	8377	7625
16	ON	11049	10074	8364	8312
64	ON	11076	10144	8411	8330

+10% overall for 1024 Bytes Blocks
other Block Sizes unchanged

Interpreted Result with Prefetch ON

Metric	Prefetch	sDMA FIFO Size (64 bit words)	64 Byte Buffer	256 Byte Buffer	512 Byte Buffer	1024 Byte Buffer
Number of Buffers Encrypted in 1 second	ON	64	11076	10144	8411	8330
Time per Buffer HW Encryption (us)	ON	64	10.29	18.58	38.89	40.05
Hardware Throughput	ON	64	6.22	13.78	13.16	25.57

+25% Hardware Throughput for 1024 Bytes Blocks
other Block Sizes unchanged

Trial: sDMA started before AES

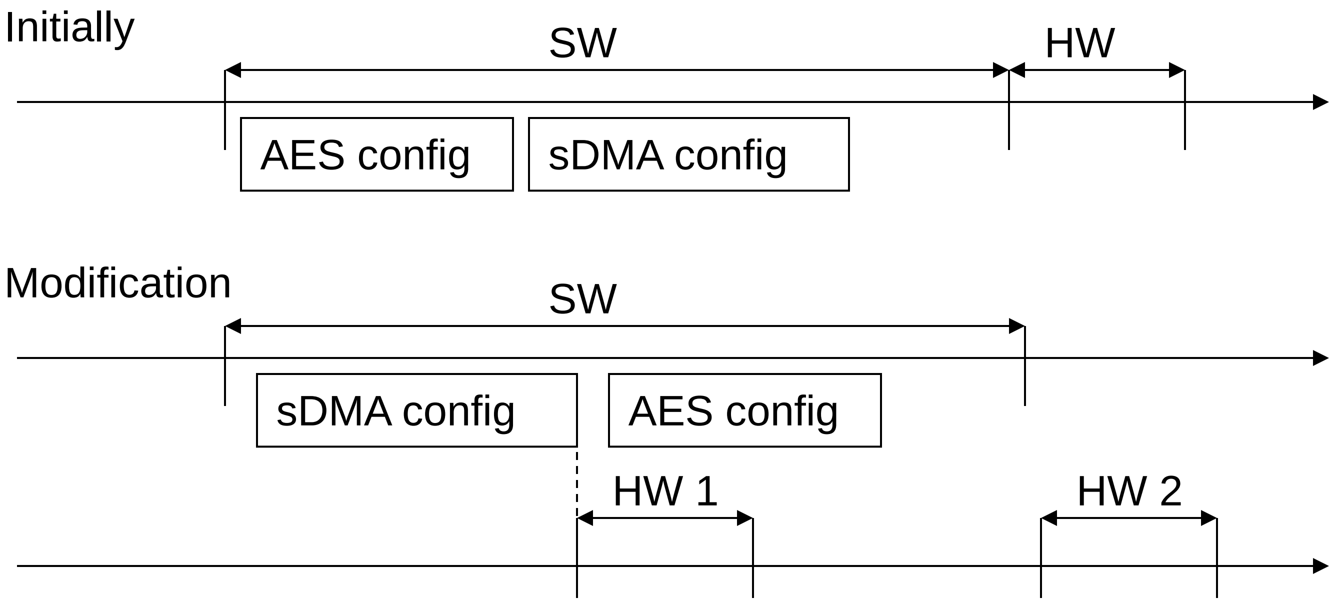

- sDMA early start allows more time for prefetch

OCP WP Trace with sDMA Early Start 256 Bytes

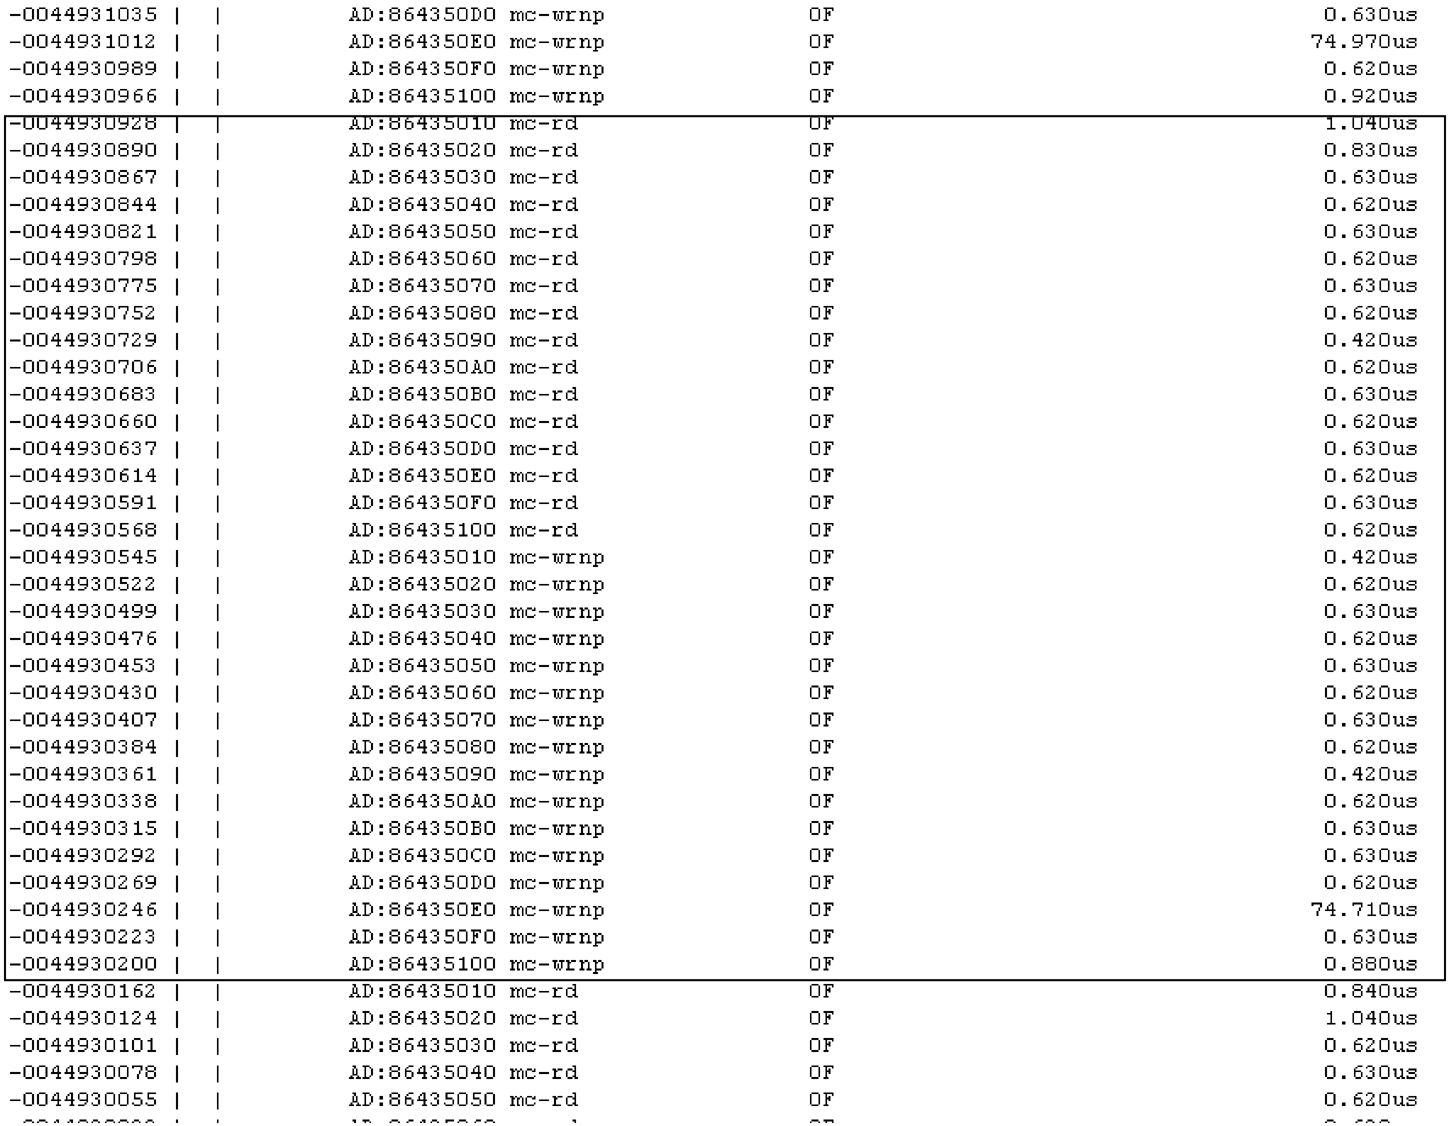

```
-0044931035 |  |         AD:864350D0 mc-wrnp          0F                             0.630us
-0044931012 |  |         AD:864350E0 mc-wrnp          0F                            74.970us
-0044930989 |  |         AD:864350F0 mc-wrnp          0F                             0.620us
-0044930966 |  |         AD:86435100 mc-wrnp          0F                             0.920us
-0044930928 |  |         AD:86435010 mc-rd            0F                             1.040us
-0044930890 |  |         AD:86435020 mc-rd            0F                             0.830us
-0044930867 |  |         AD:86435030 mc-rd            0F                             0.630us
-0044930844 |  |         AD:86435040 mc-rd            0F                             0.620us
-0044930821 |  |         AD:86435050 mc-rd            0F                             0.630us
-0044930798 |  |         AD:86435060 mc-rd            0F                             0.620us
-0044930775 |  |         AD:86435070 mc-rd            0F                             0.630us
-0044930752 |  |         AD:86435080 mc-rd            0F                             0.620us
-0044930729 |  |         AD:86435090 mc-rd            0F                             0.420us
-0044930706 |  |         AD:864350A0 mc-rd            0F                             0.620us
-0044930683 |  |         AD:864350B0 mc-rd            0F                             0.630us
-0044930660 |  |         AD:864350C0 mc-rd            0F                             0.620us
-0044930637 |  |         AD:864350D0 mc-rd            0F                             0.630us
-0044930614 |  |         AD:864350E0 mc-rd            0F                             0.620us
-0044930591 |  |         AD:864350F0 mc-rd            0F                             0.630us
-0044930568 |  |         AD:86435100 mc-rd            0F                             0.620us
-0044930545 |  |         AD:86435010 mc-wrnp          0F                             0.420us
-0044930522 |  |         AD:86435020 mc-wrnp          0F                             0.620us
-0044930499 |  |         AD:86435030 mc-wrnp          0F                             0.630us
-0044930476 |  |         AD:86435040 mc-wrnp          0F                             0.620us
-0044930453 |  |         AD:86435050 mc-wrnp          0F                             0.630us
-0044930430 |  |         AD:86435060 mc-wrnp          0F                             0.620us
-0044930407 |  |         AD:86435070 mc-wrnp          0F                             0.630us
-0044930384 |  |         AD:86435080 mc-wrnp          0F                             0.620us
-0044930361 |  |         AD:86435090 mc-wrnp          0F                             0.420us
-0044930338 |  |         AD:864350A0 mc-wrnp          0F                             0.620us
-0044930315 |  |         AD:864350B0 mc-wrnp          0F                             0.630us
-0044930292 |  |         AD:864350C0 mc-wrnp          0F                             0.630us
-0044930269 |  |         AD:864350D0 mc-wrnp          0F                             0.620us
-0044930246 |  |         AD:864350E0 mc-wrnp          0F                            74.710us
-0044930223 |  |         AD:864350F0 mc-wrnp          0F                             0.630us
-0044930200 |  |         AD:86435100 mc-wrnp          0F                             0.880us
-0044930162 |  |         AD:86435010 mc-rd            0F                             0.840us
-0044930124 |  |         AD:86435020 mc-rd            0F                             1.040us
-0044930101 |  |         AD:86435030 mc-rd            0F                             0.620us
-0044930078 |  |         AD:86435040 mc-rd            0F                             0.630us
-0044930055 |  |         AD:86435050 mc-rd            0F                             0.620us
```

256 byte
Buffer RD and WR

Complete 256
Byte Buffer
Prefetched

Results various sDMA Configurations

- sDMA early start: No performance improvement
- Set channel in and channel out to high priority: gain for 512 bytes buffer and 1024 bytes buffer
- Thread reservation:
  - channels high priority
  - one thread reserved read and one thread reserved write
  - arbitration rate of 1
  - No Benefit
- Write posted (all except last of transfer) instead of write non posted for ALL logical channels: no benefit

End Result sDMA Configurations

	64 Byte Buffer	256 Byte Buffer	512 Byte Buffer	1024 Byte Buffer
Number of Buffer Encryption per Second	11426	10709	10696	8813
Time for a single Buffer Encryption (us)	87.52	93.38	93.49	113.47
Throughput (MBps)	0.73	2.74	5.47	9.02
Gain from Default Config	12%	6%	28%	15%

Note Hardware and Software Contributions cannot be differentiated because sDMA is started before AES is enabled.

Conclusion sDMA Configuration

Configuration	Used in Optimal Configuration on Board with no Concurrent Traffic	Recommended to use in Production Software	Positive Impact Anticipated in Loaded Platform
sDMA early start	Yes	Yes	Yes
Input and output channel high priority	Yes	Yes	Yes
Thread Reservation	No	Yes	Yes
Write Posted except Last	No	Yes	Yes
Prefetch ON	Yes	Yes	Yes
FIFO Size @ 32	No	Yes	Yes
Packet Synchronization	No	Yes	Yes

🟩 Strongly recommended modifications

BACKUP

References

- OMAP4460 ES1 Public TRM v0
- OCP Watchpoint Chapter 28.8.3 of TRM

Acronyms

- AES: Advanced Encryption Standard
- CBC: Cipher Block Chaining
- DDR: Double Data Rate
- DMA: Direct Memory Access
- DMM: Dynamic Memory Management
- L3: Interconnect Level 3 (Level 1 and 2 being caches)
- OCP: Open Core Protocol