The AMD EPYC Rome processor

Björn Dick (HLRS), Thomas Bönisch (HLRS)

Node - Overview

- 2 x AMD EPYC 7742 (Zen2 aka “Rome”)
  - 64 cores @ 2.25GHz, AVX2
  -

Cache	L1D	L1I	L2	L3
Size	32kB		512kB	16MB
Cache line size	64B			
Associativity	8-way			16-way
private/shared?	private			shared among 4 cores only!
Write policy	write-back			
Inclusion policy	inclusive			victim cache

- DRAM:
  - 256GB @ 380GB/s

Core

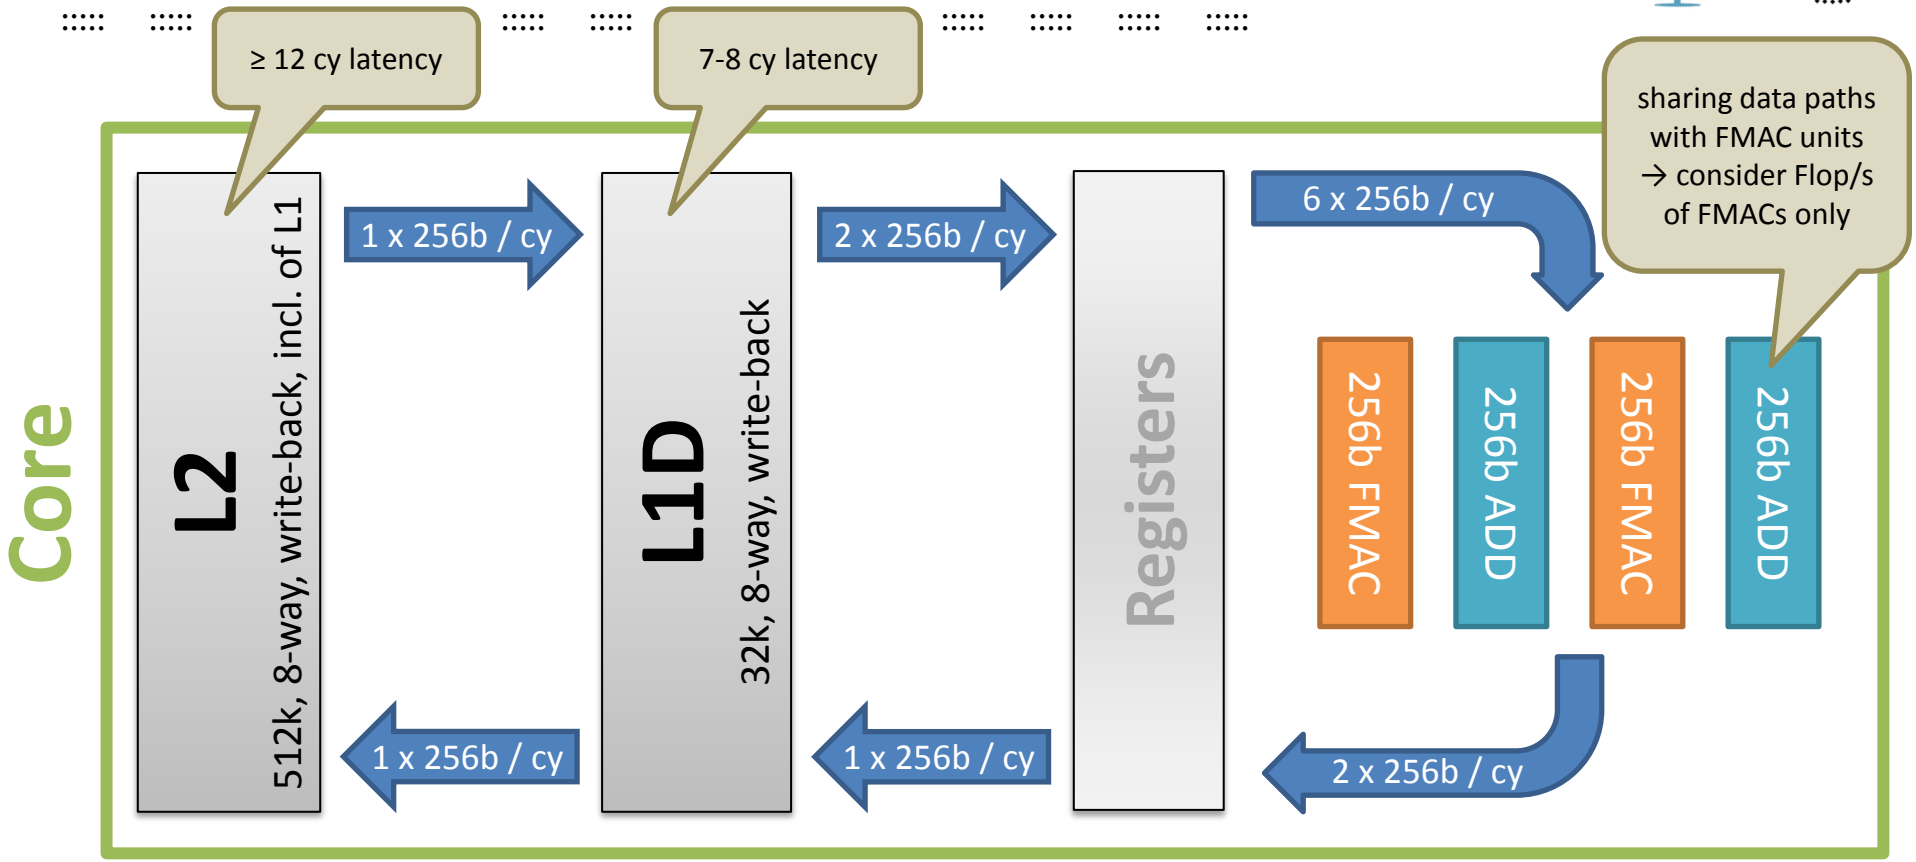
≥ 12 cy latency
7-8 cy latency
sharing data paths with FMAC units → consider Flop/s of FMACs only
Core
L2
512k, 8-way, write-back, incl. of L1
1 x 256b / cy
L1D
32k, 8-way, write-back
2 x 256b / cy
Registers
6 x 256b / cy
256b FMAC
256b ADD
256b FMAC
256b ADD
1 x 256b / cy
1 x 256b / cy
2 x 256b / cy

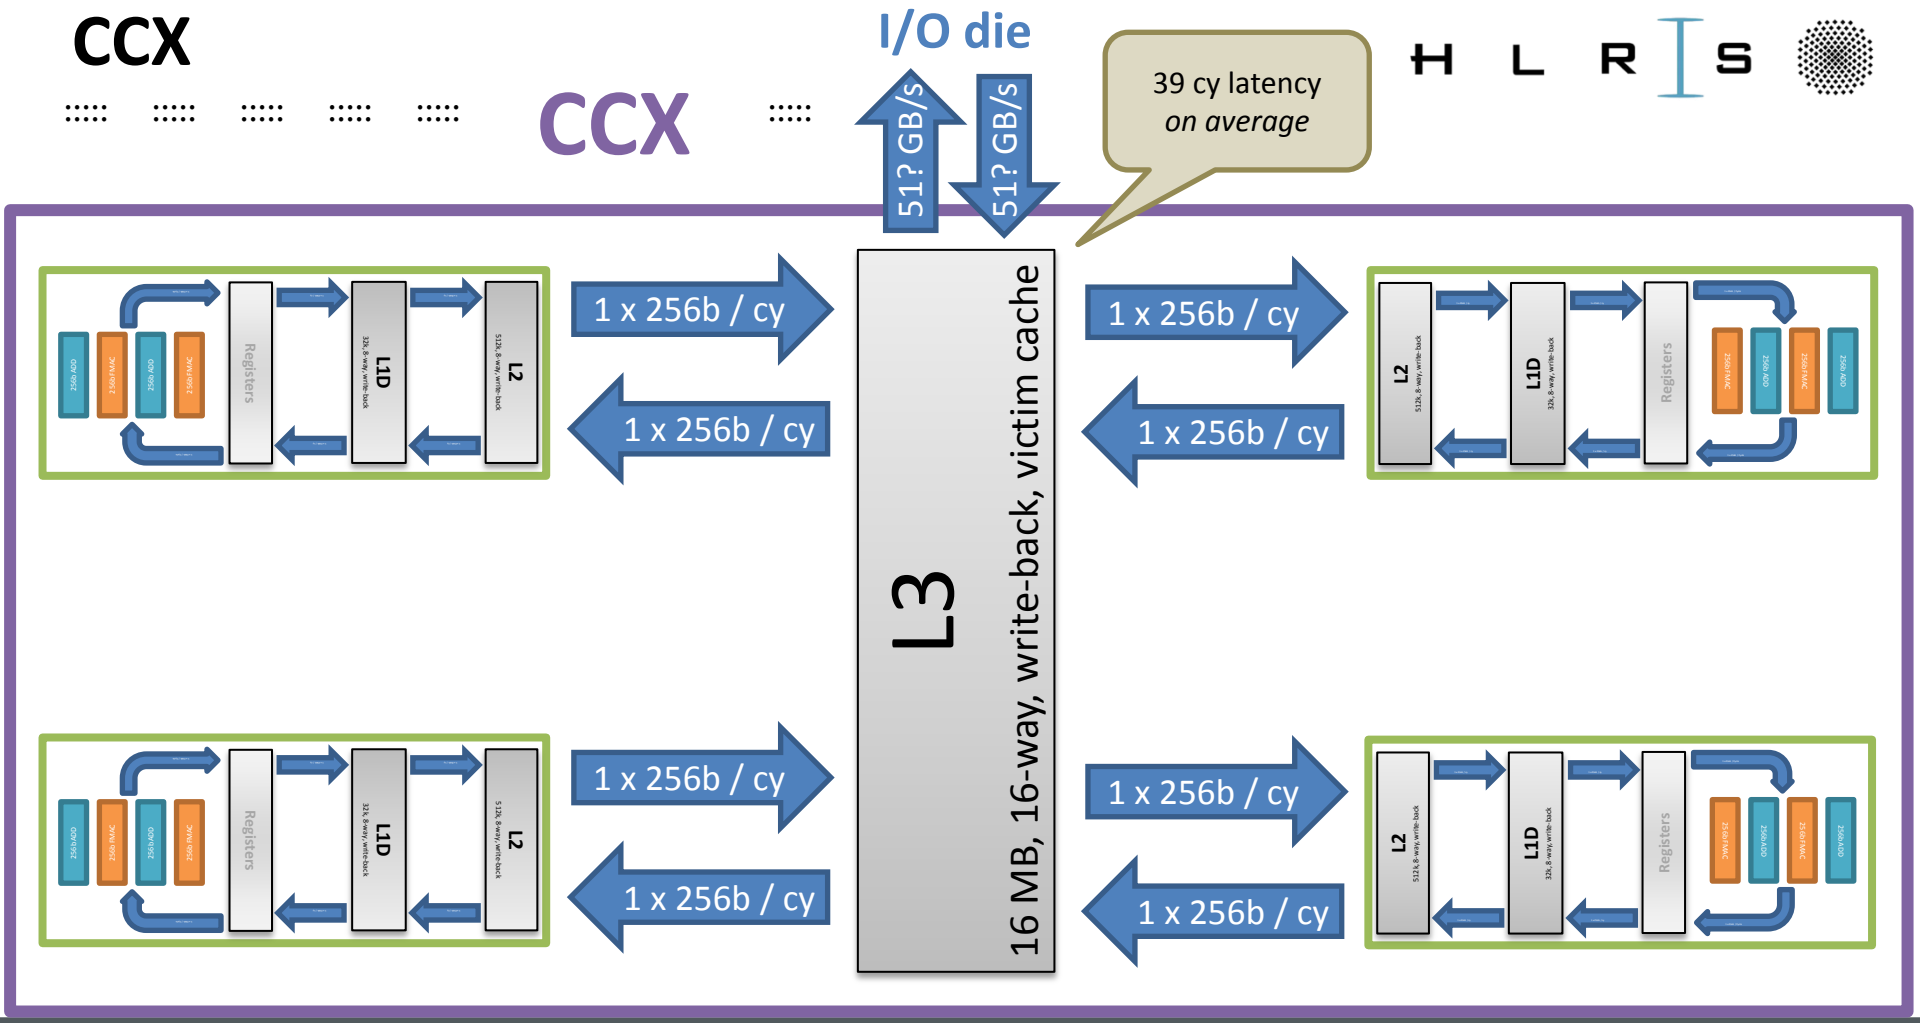
CCX
I/O die
51? GB/s
51? GB/s
39 cy latency
on average
1 x 256b / cy
1 x 256b / cy
1 x 256b / cy
1 x 256b / cy
L3
16 MB, 16-way, write-back, victim cache
1 x 256b / cy
1 x 256b / cy
1 x 256b / cy
1 x 256b / cy
Registers
L1D
32k, 8-way, write-back
L2
512k, 8-way, write-back

Socket

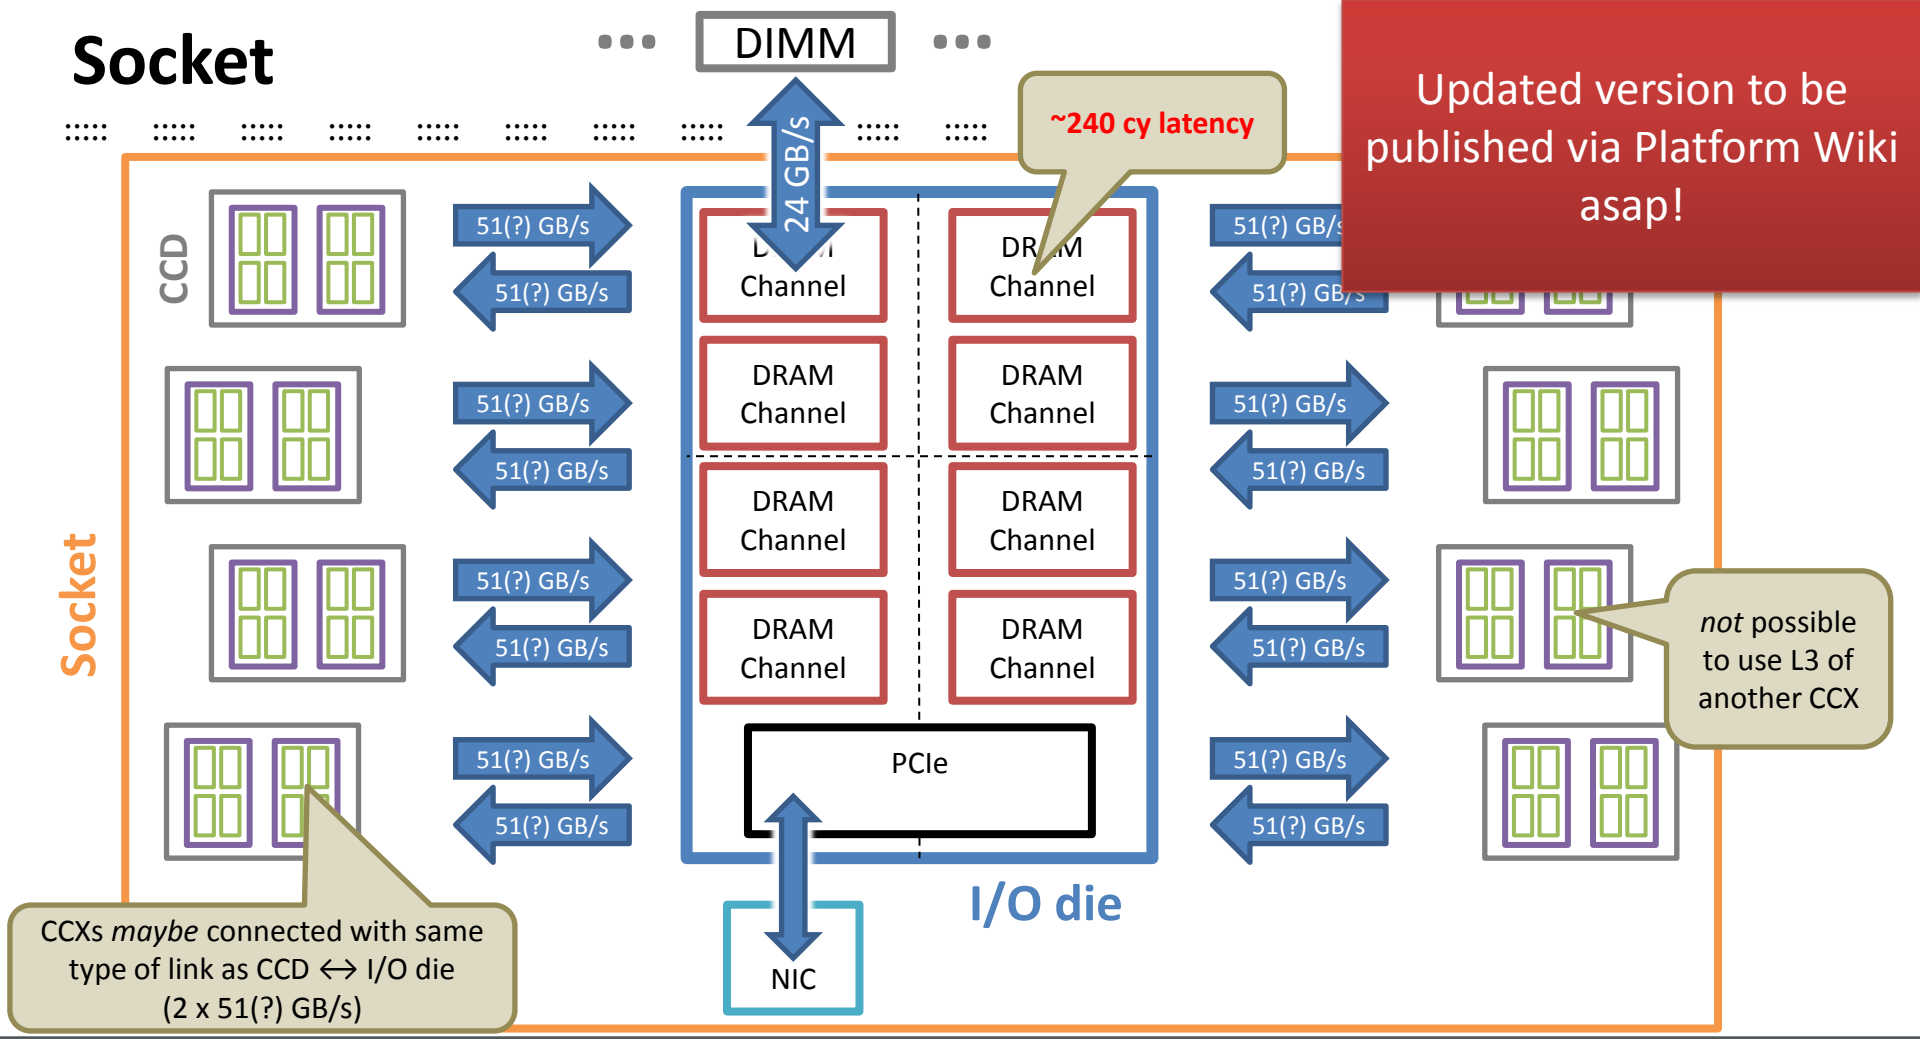
DIMM
24 GB/s
~240 cy latency
Updated version to be published via Platform Wiki asap!
CCD
51(?) GB/s
51(?) GB/s
DRAM Channel
DRAM Channel
DRAM Channel
DRAM Channel
DRAM Channel
DRAM Channel
DRAM Channel
DRAM Channel
PCIe
NIC
I/O die
Socket
not possible to use L3 of another CCX
CCXs *maybe* connected with same type of link as CCD ↔ I/O die (2 x 51(?) GB/s)

Node

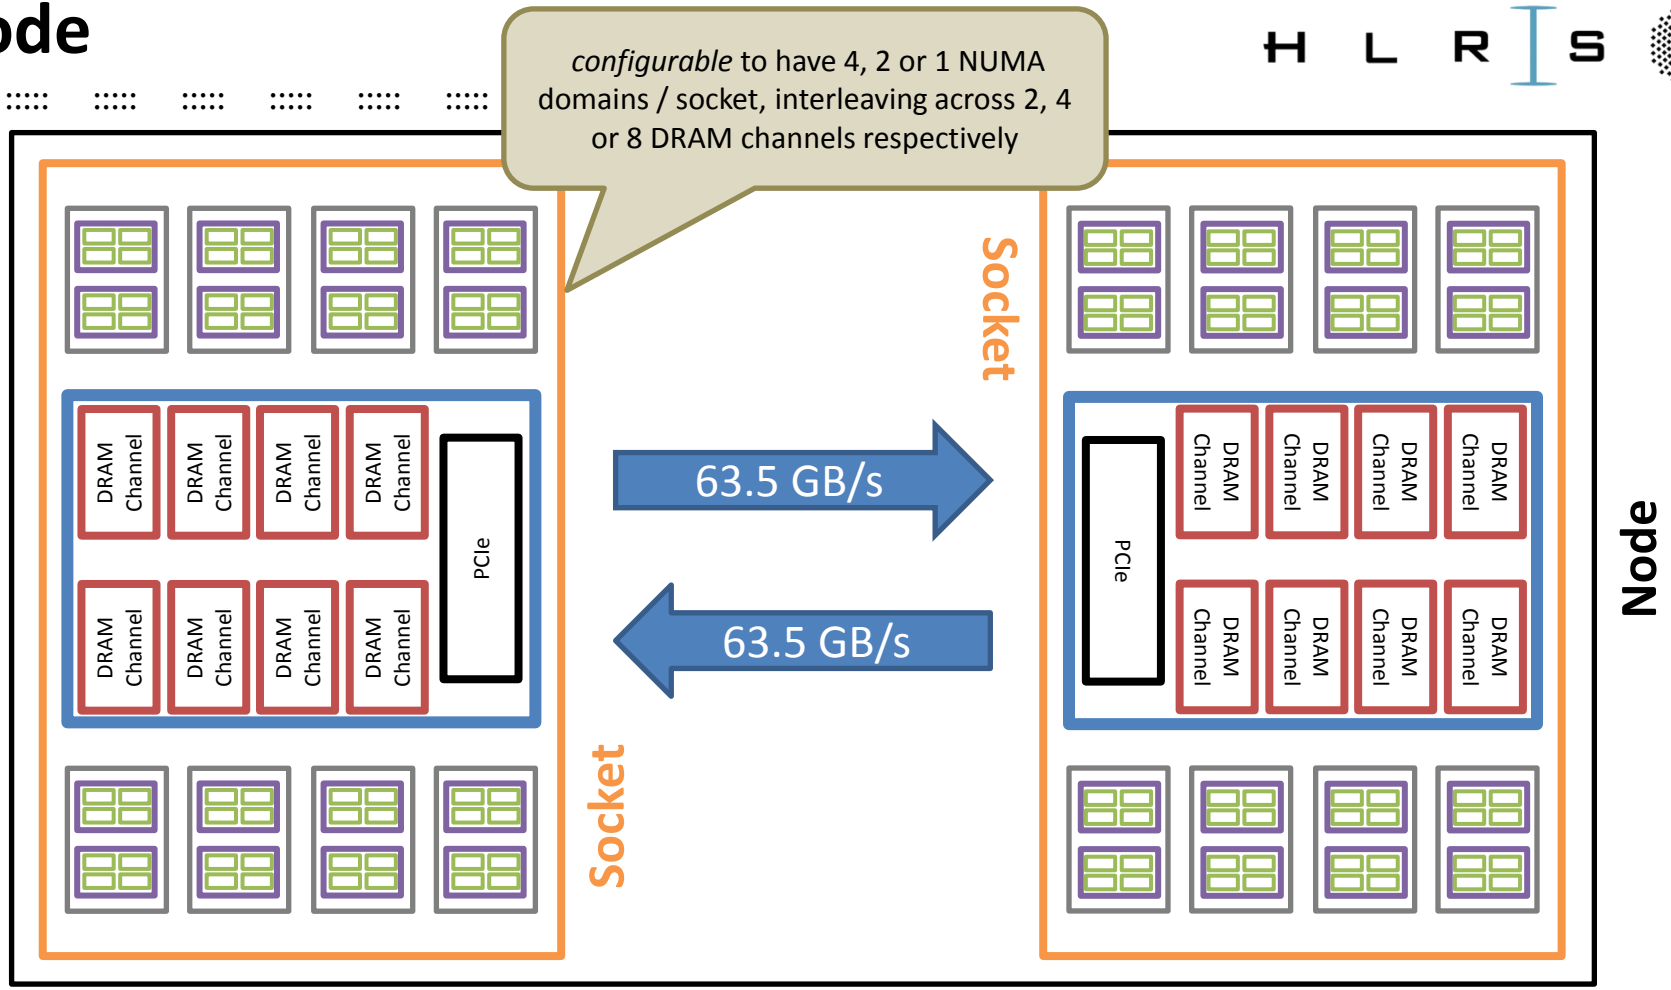
configurable to have 4, 2 or 1 NUMA domains / socket, interleaving across 2, 4 or 8 DRAM channels respectively
Socket
Socket
DRAM Channel
PCIe
63.5 GB/s
63.5 GB/s
Node

Remarks

- Don’t use it as a black-box!
- Core is quite similar to Haswell
- SoC however is quite different
- Rome core @ 2.25 GHz vs.
  Haswell core @ 2.5 GHz
- Haswell node: 2 x 70 GB/s
  Rome node: 2 x 190 GB/s
- → it’s ***not*** fair to compare
  128 Haswell cores
  (= 6 x 2 x 70 GB/s = 840 GB/s)
  to 128 Rome cores
  (= 2 x 190 GB/s = 380GB/s)
  if your code is bound by DRAM B/W
- Compare node vs. node instead

- It might be beneficial to use less than 128 cores per node!

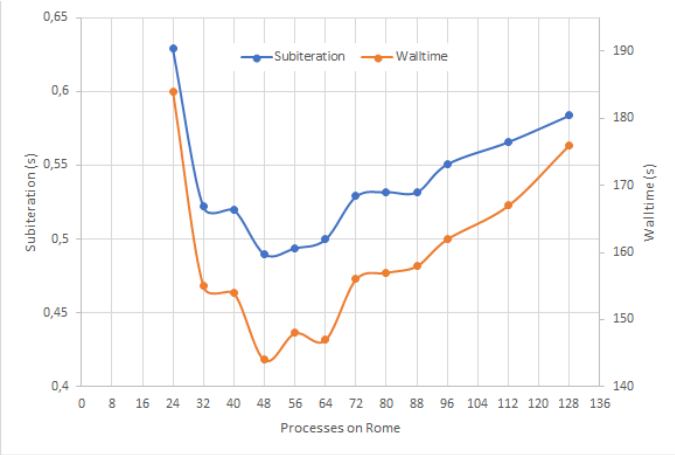

- Hardware counters available for Flop/s, cache and DRAM B/W

Software :: Compiler

GNU

- currently 9.1.0
- will be used as default for a start

PGI

- currently 19.5 / 19.7

AMD (called AOCC)

- currently 2.1.0
- LLVM clone, including flang (now) / F18 (future)
- Give it a try, HLRS will be able to influence further development of this compiler!

Intel

- currently 18.0.2 / 19.0.4 / 19.1.0
- Give it a try, icc generated highly competitive code for benchmarks on EPYC Naples!