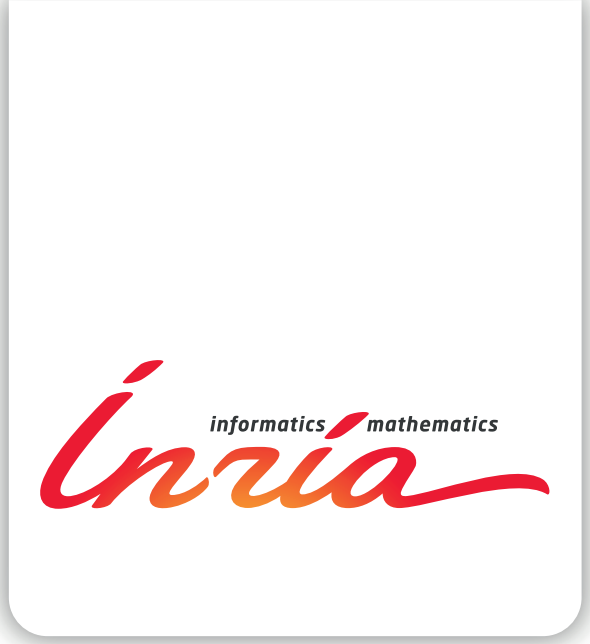

Performance models of parallel applications and machines

INTRODUCTION

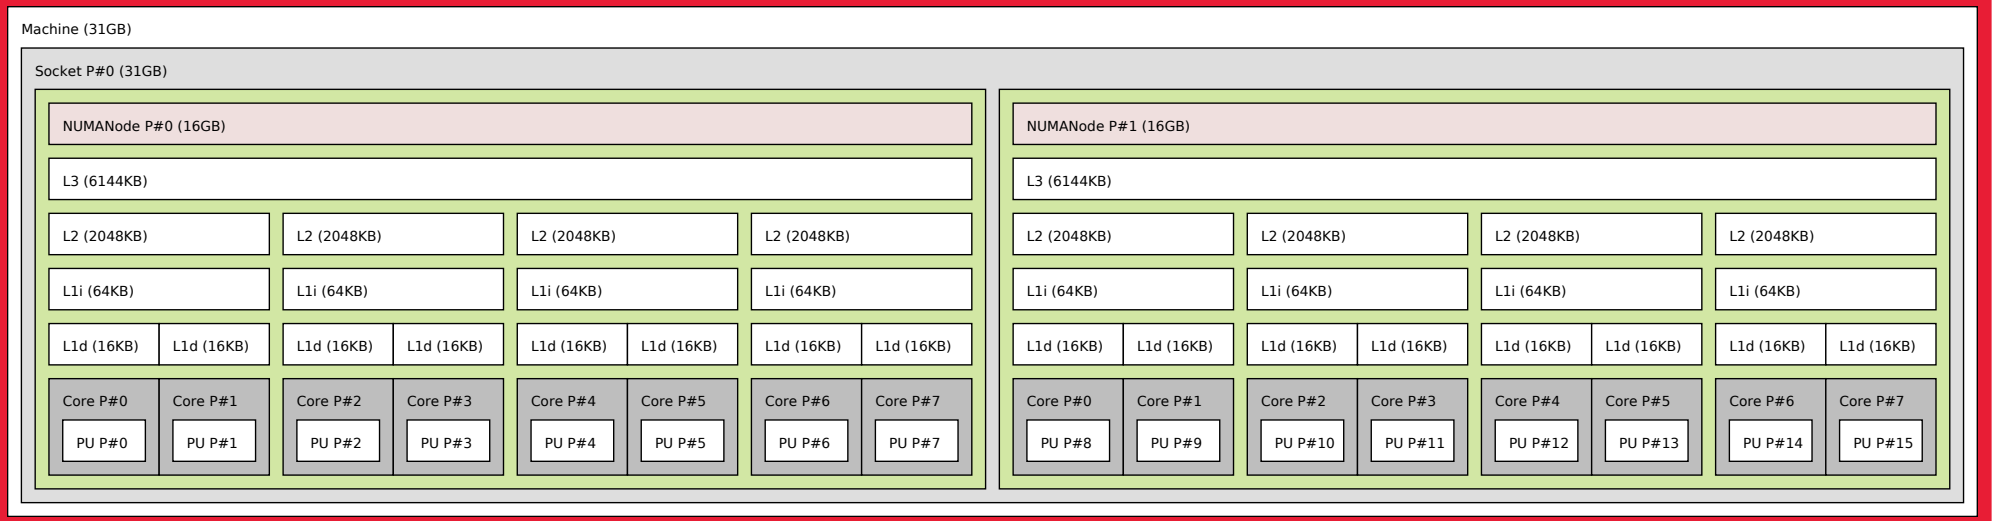

Memory hierarchy is growing deeper and larger.

- caches levels
- NUMA nodes
- Sockets
- 3D RAM
- NVRAM
- IO
- Network

INTRODUCTION

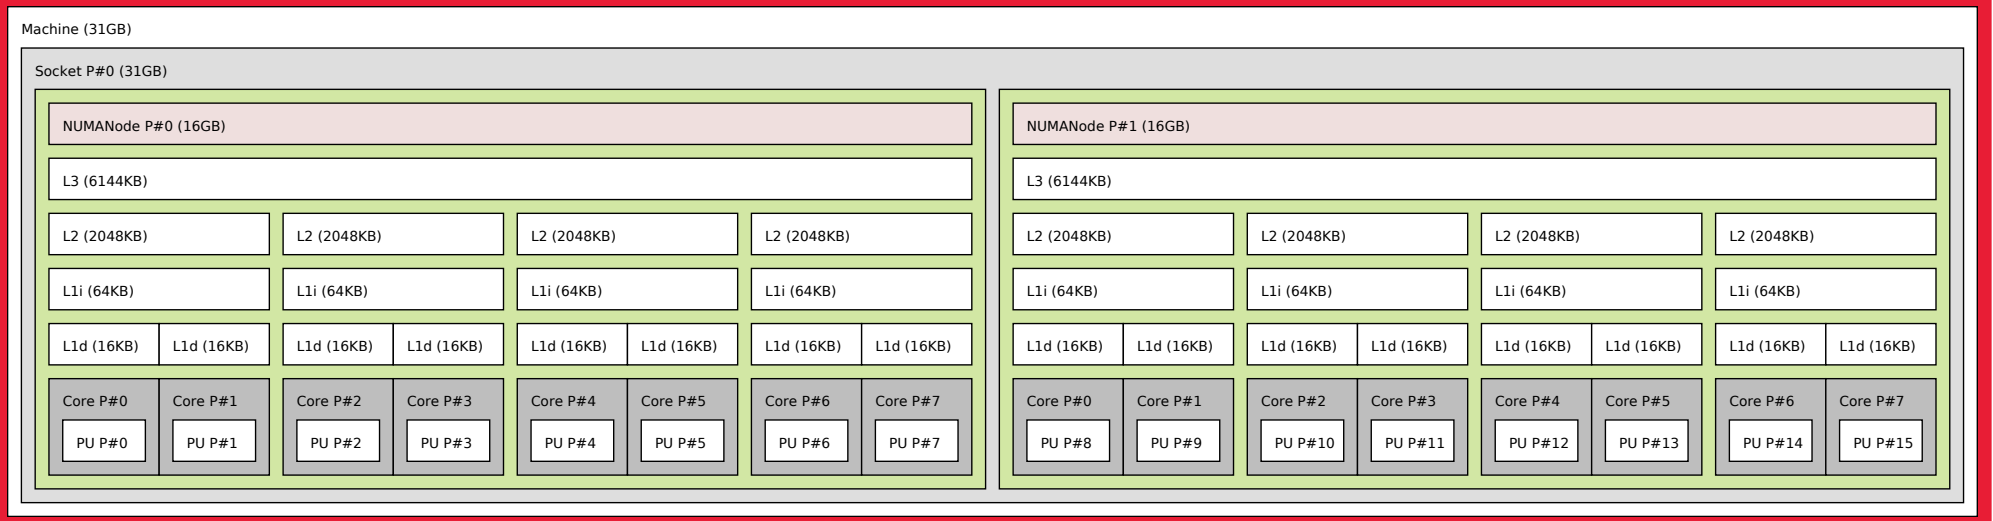

Hence, data locality becomes a concern for performance.
Process mapping enhance data locality.

How to efficiently choose processes placement ?

Affinity Mapping

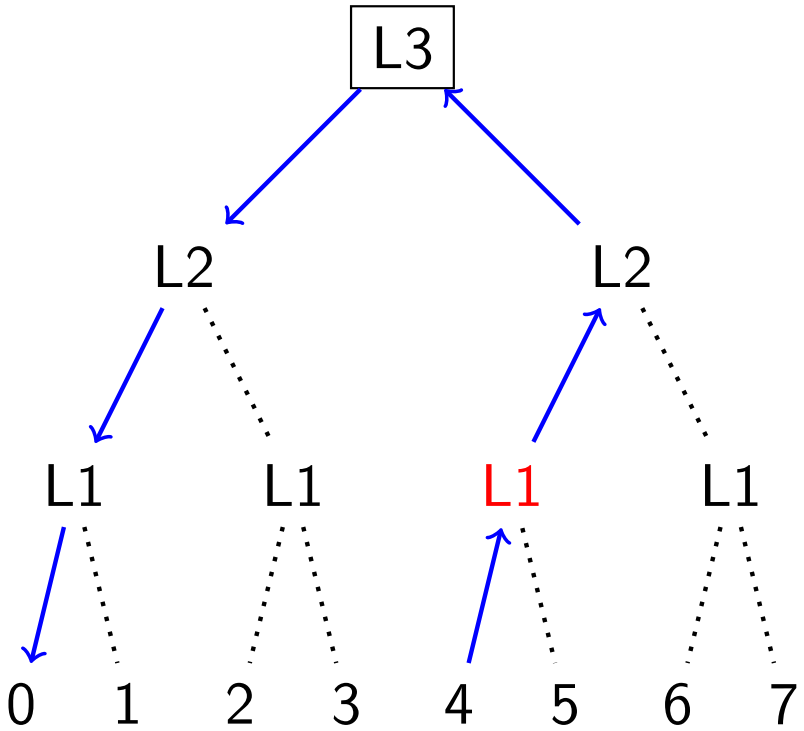

Improves

- Communications
- Data sharing
- Synchronisation

Affinity Mapping

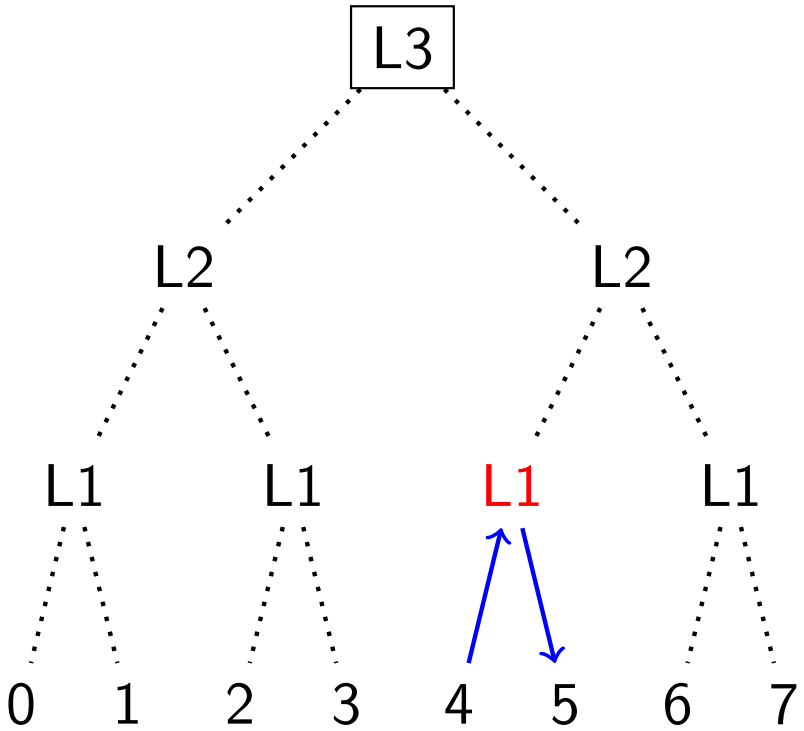

Improves

- Communications
- Data sharing
- Synchronisation

Affinity Mapping

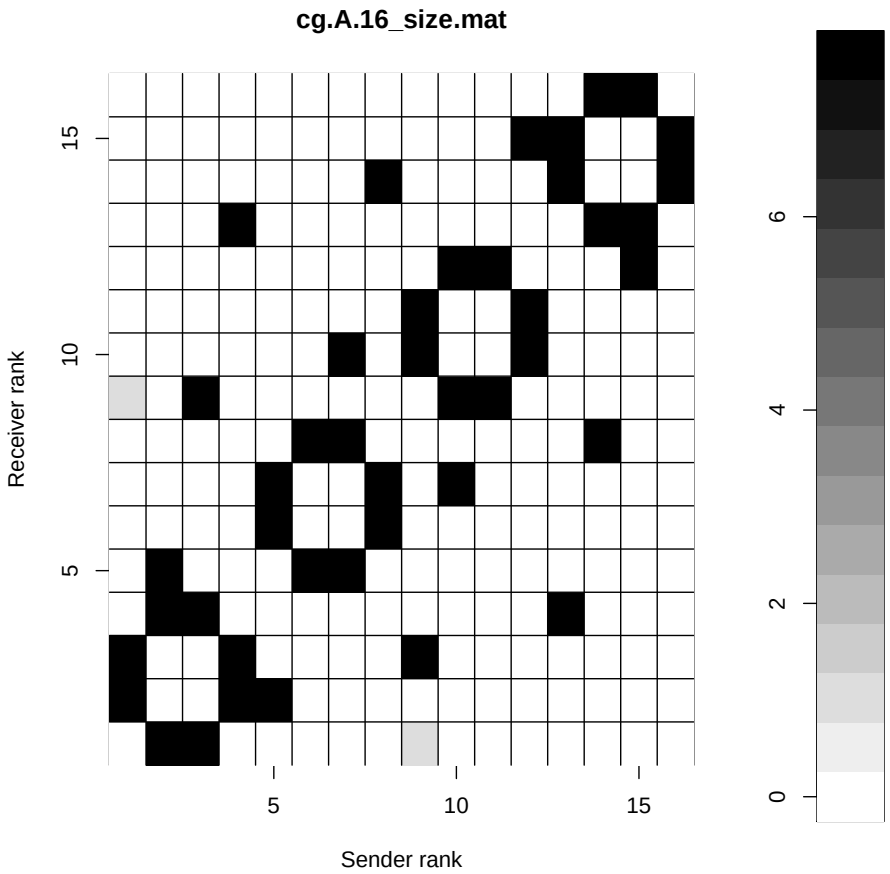

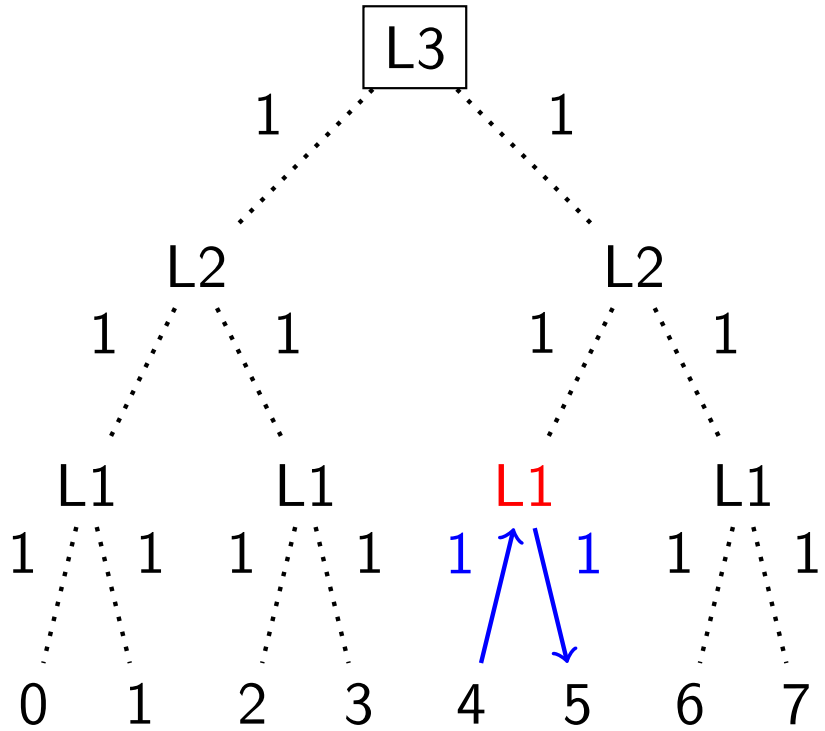

Affinity Mapping

Afinnity metrics

number of messages.

total messages size.

Tree weights

Latency.

Bandwidth.

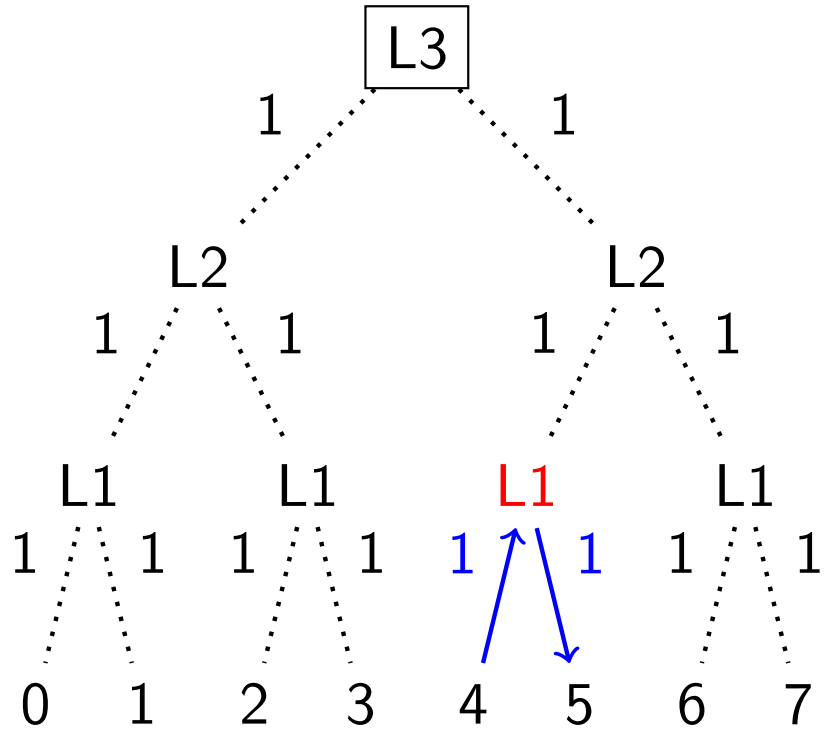

Concurrency Mapping

Reduces

Races on memory levels

Contention on bandwidth.

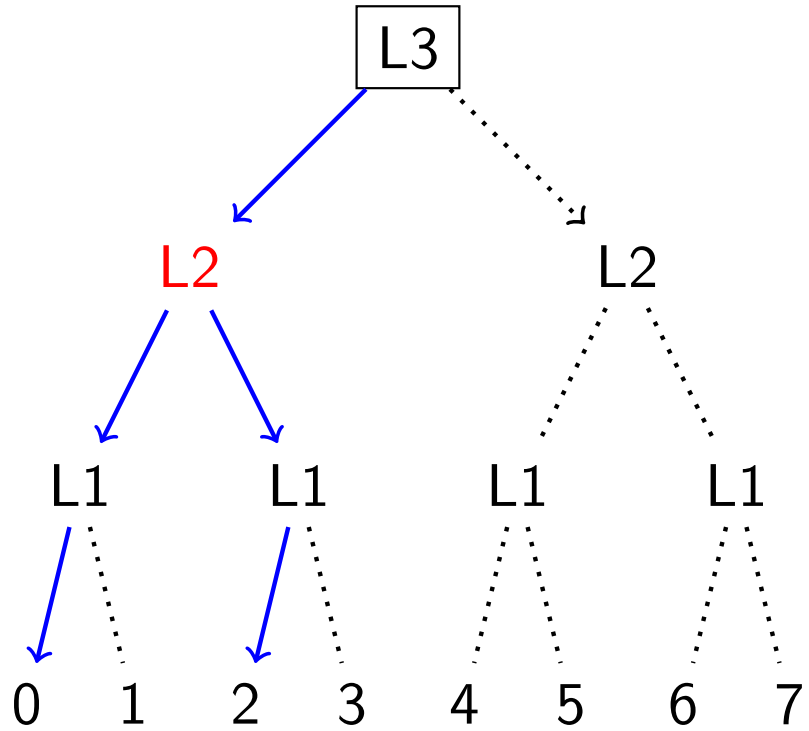

Concurrency Mapping

Reduces

Races on memory levels

Contention on bandwidth.

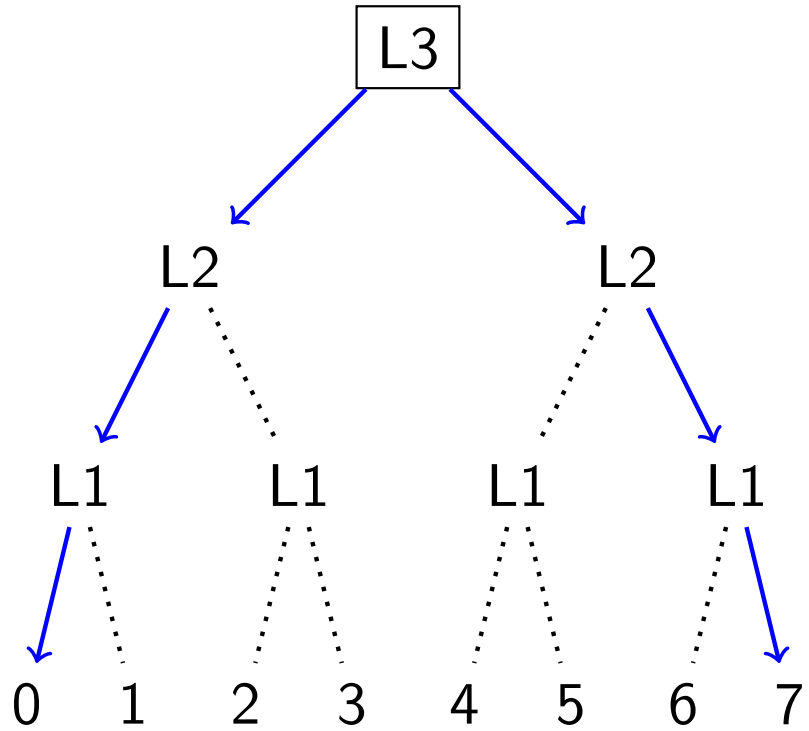

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

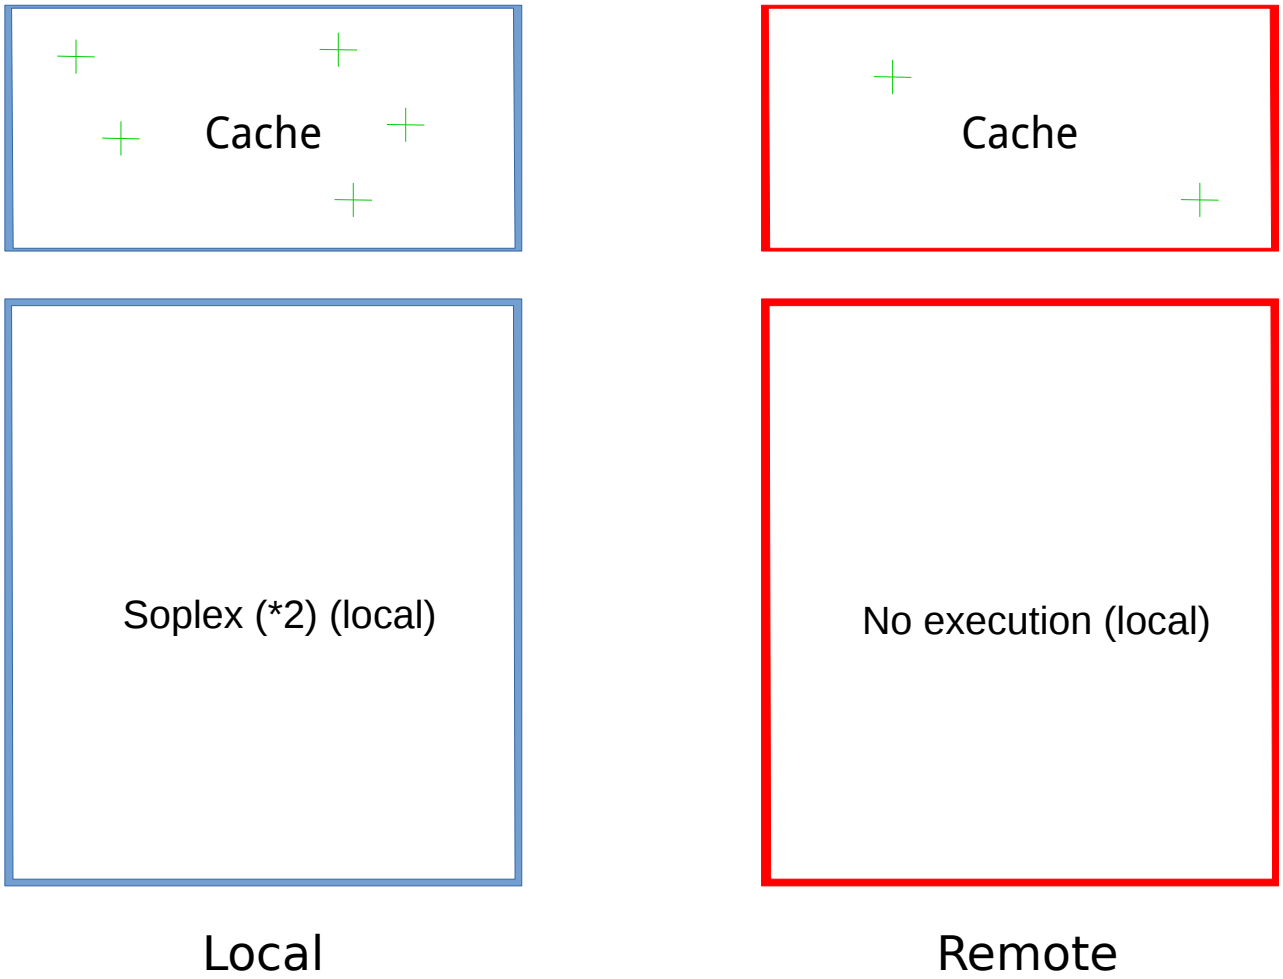

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

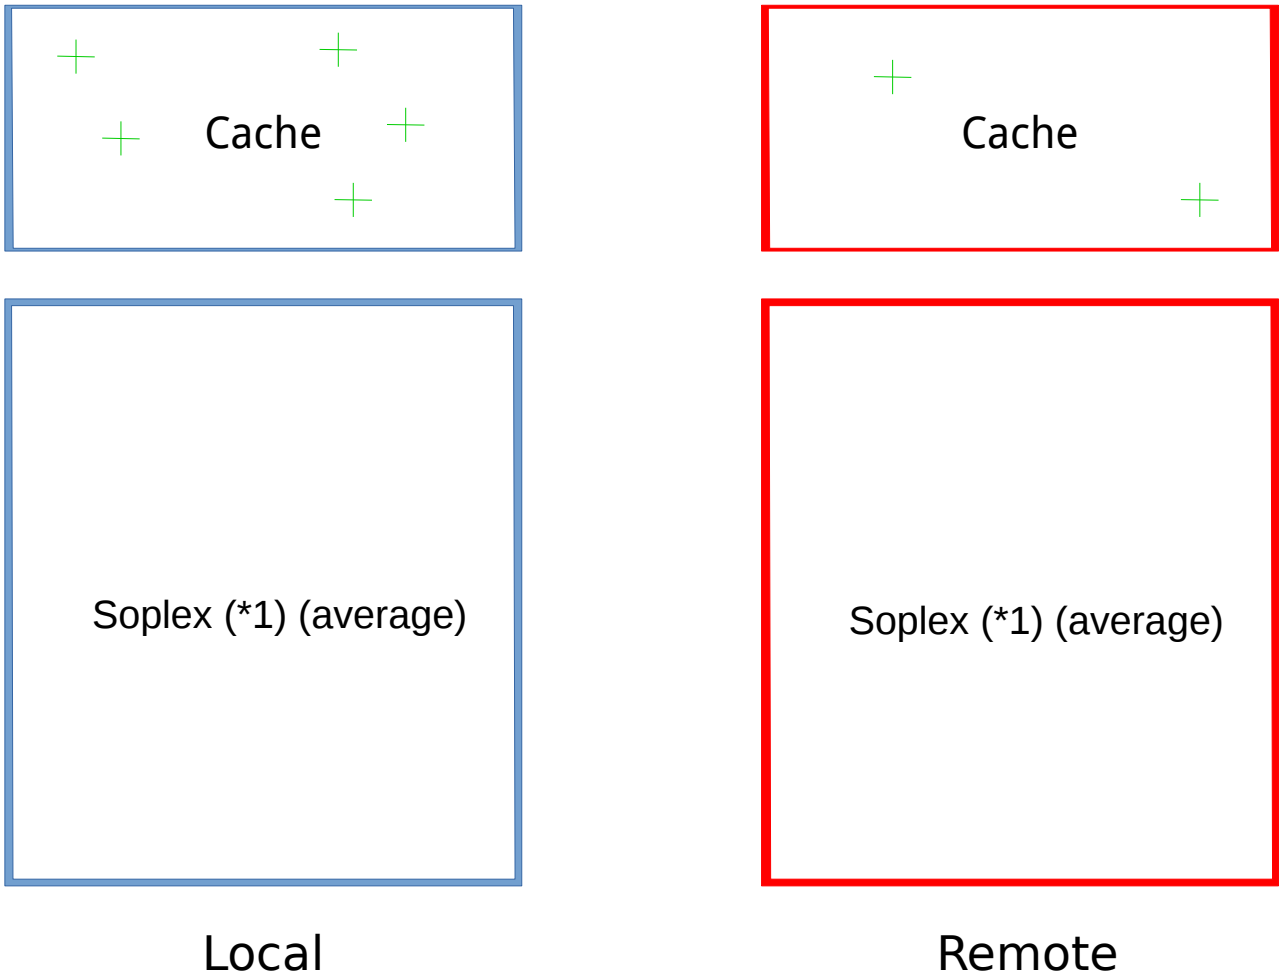

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

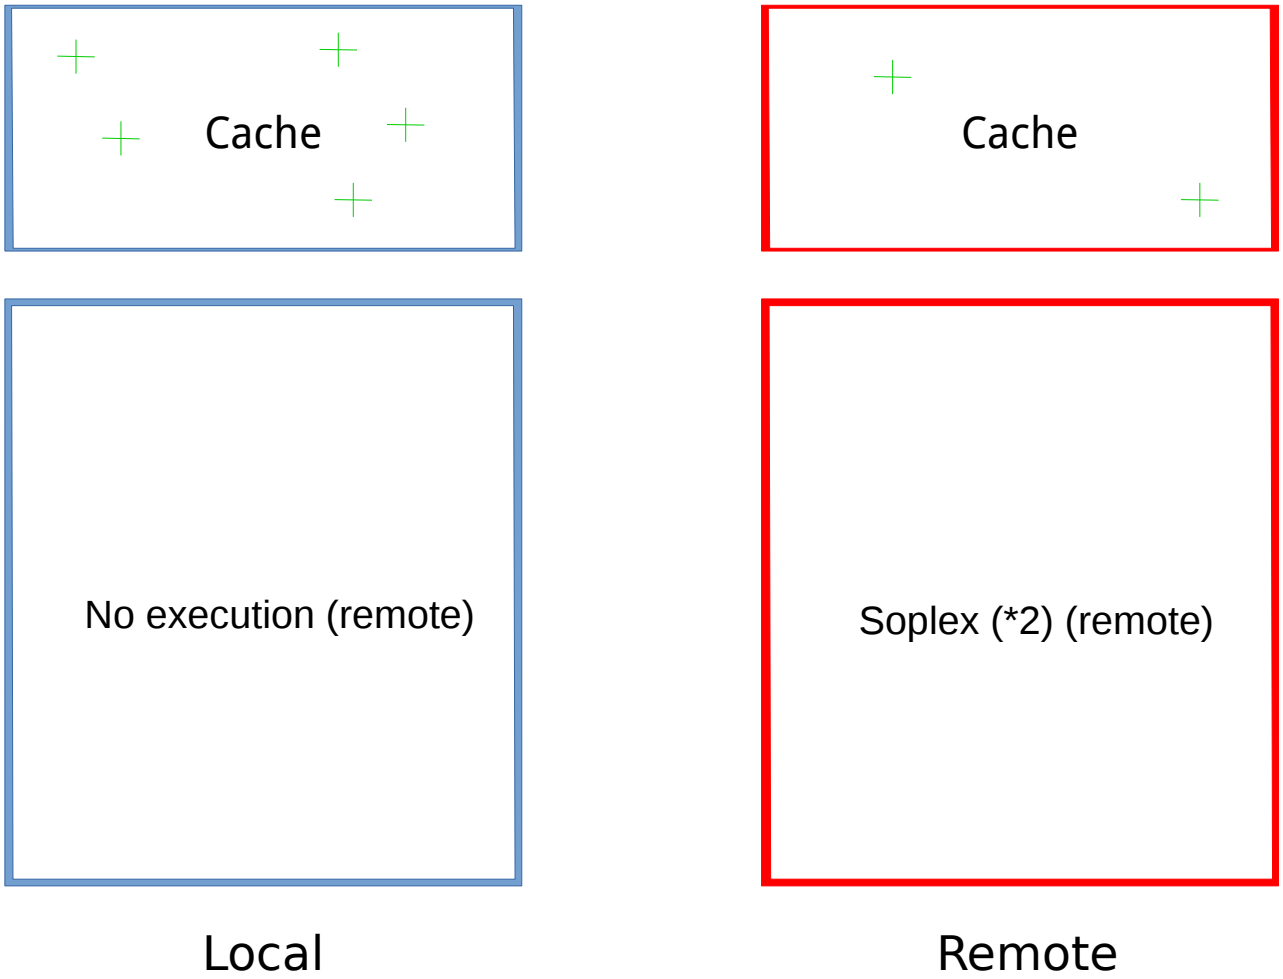

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

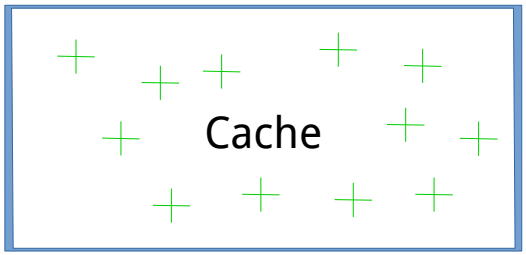

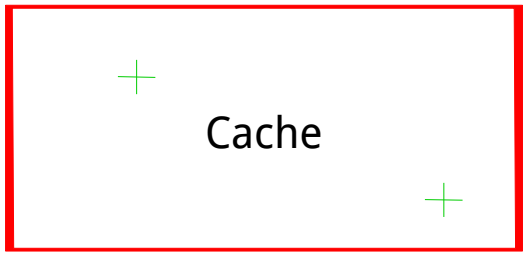

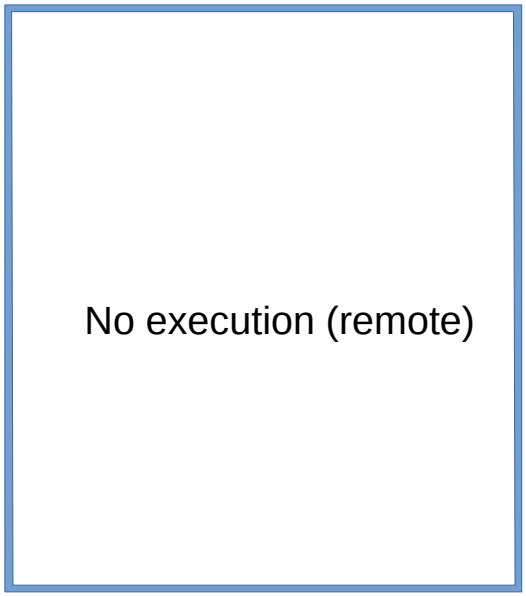

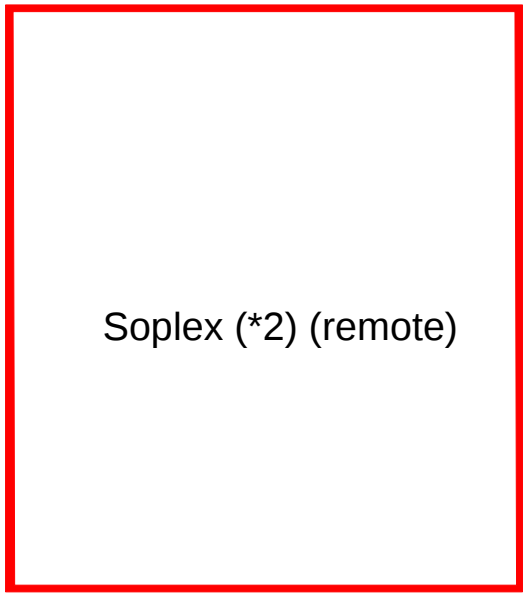

Local

Remote

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

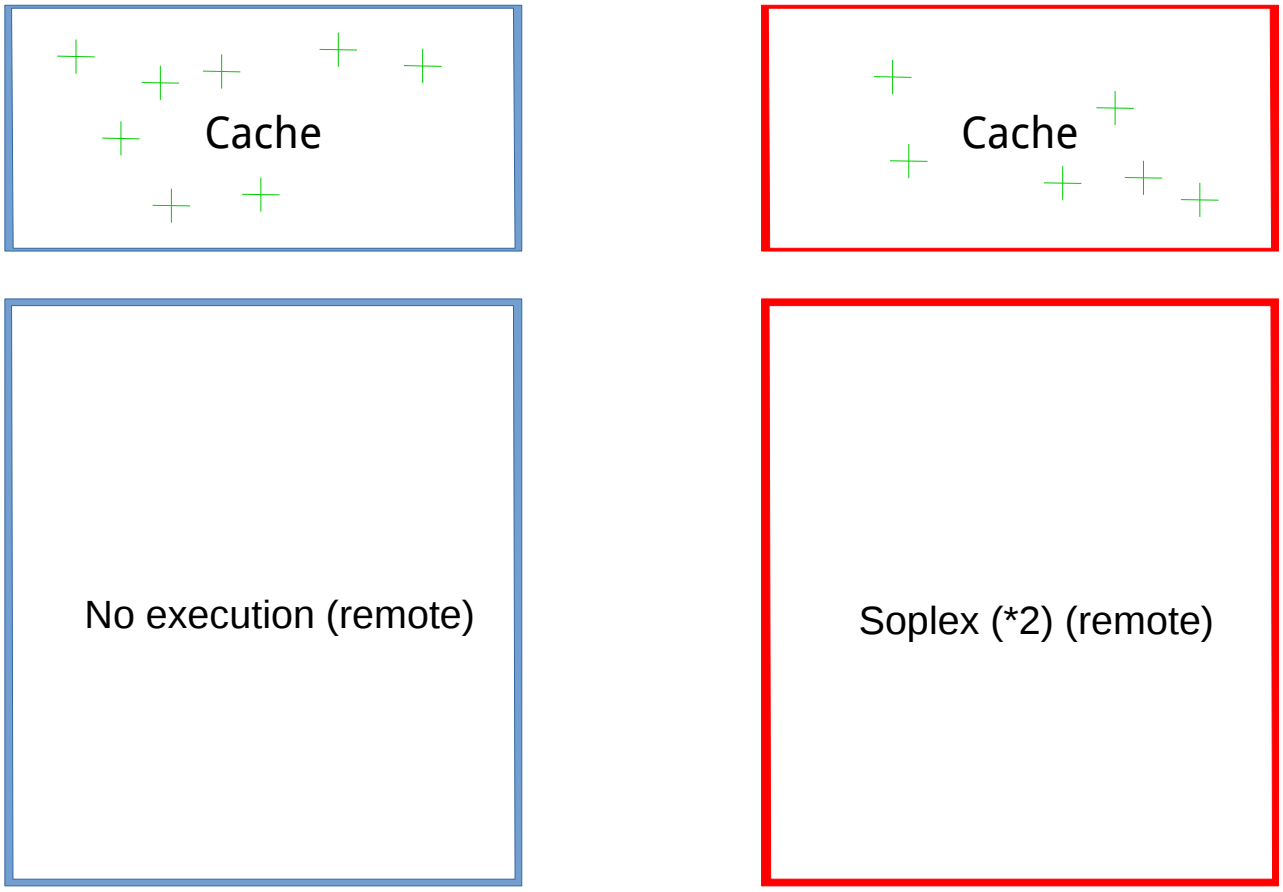

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

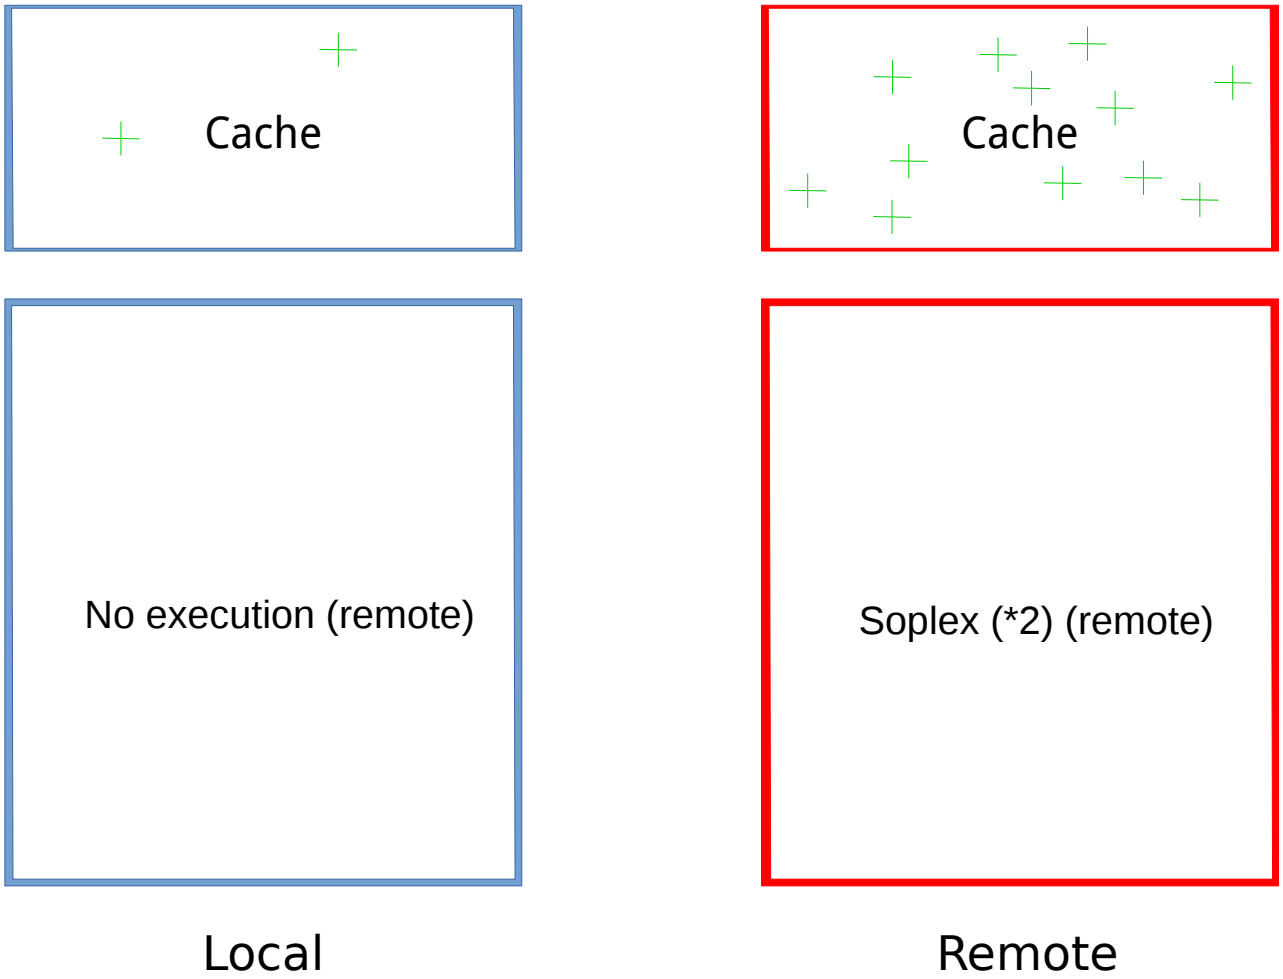

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

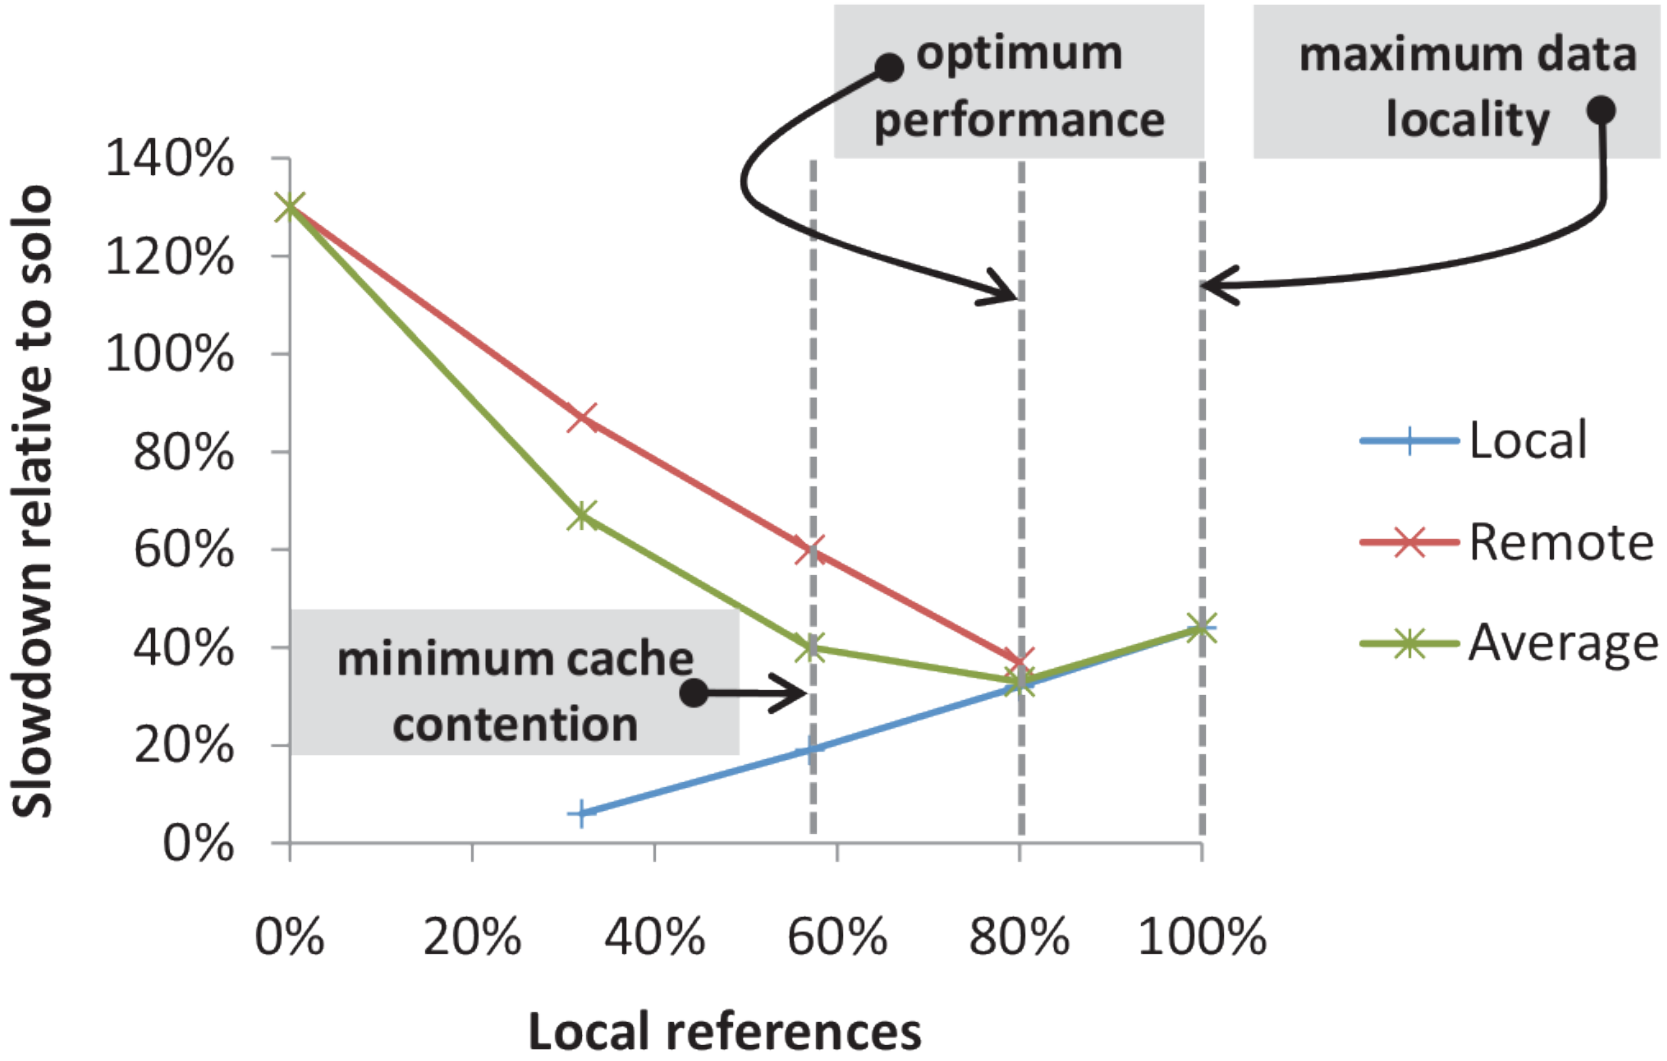

A tradeof case [Zoltan Majo and Thomas R. Gross. 2011.]

We aim to tackle a single part of the problem.

We only consider the local execution.

What are the slowdown causes when cloning the local processus under a shared memory?

Parallel Application Model

Memory continuum $\Longrightarrow$ memory bound application continuum.

Metrics: Performance, data reuse.

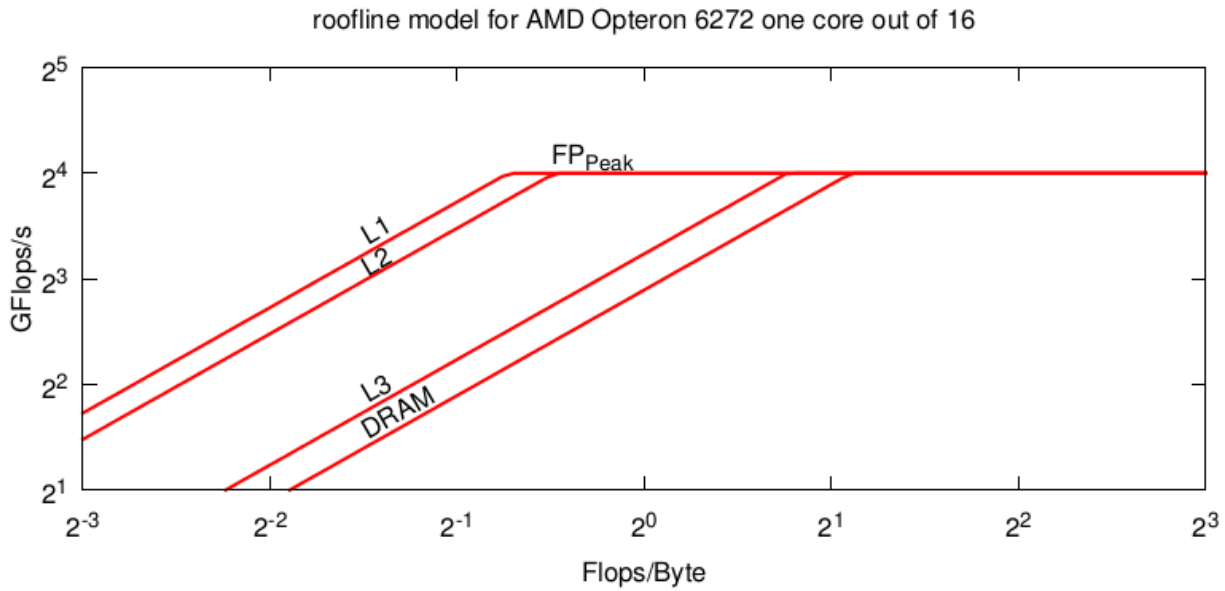

Parallel Application Model

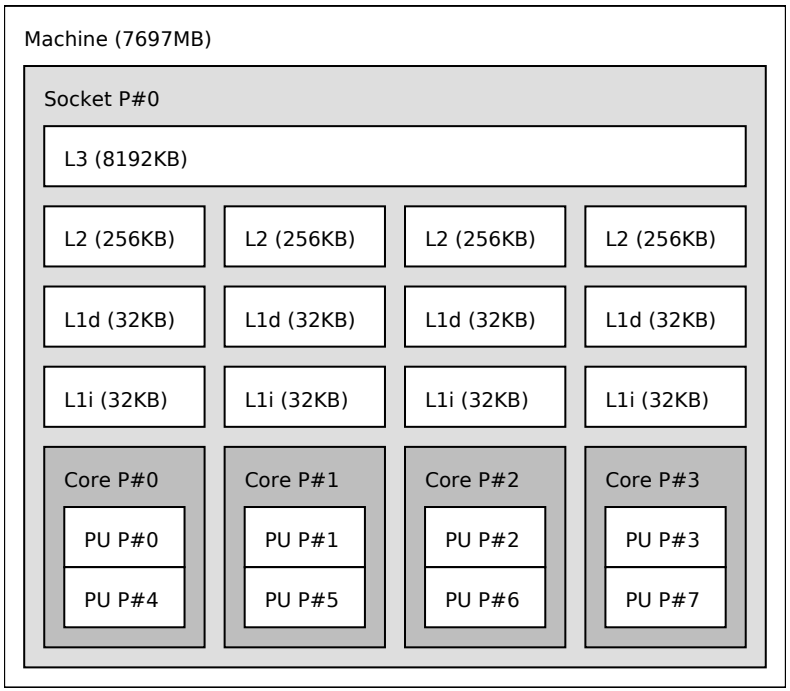

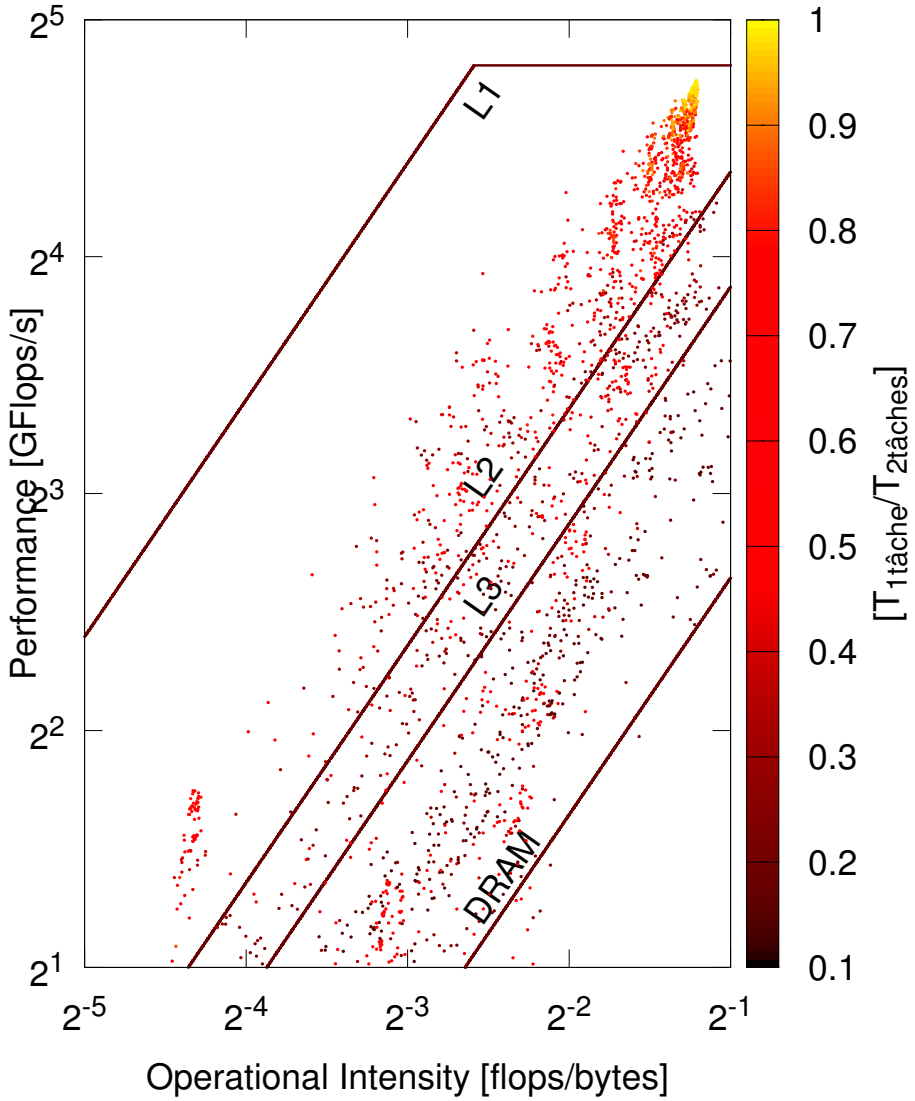

Parallel Application Model

Understanding the slow down causes is not so easy.

We need a lower level of application model.

The roofline representation hides the application dynamic memory usage.

A Visual tool using hwloc & PAPI to track imbalances

We need to examine a lot of benchmarks cases. . .

A lot of low levels metrics. . .

And detect imbalances such as reaching memory limits,
as shown in the Cache Aware Roolfine model.

A Visual tool using hwloc & PAPI to track imbalances

```
FOPS{NUMANode, PAPI_FP_OPS}
L3TCM{L3, PAPI_L3_TCM}
L2ICM{L2, PAPI_L2_ICM}
L1DCM{L1d, PAPI_L1_DCM}
L1ICM{L1i, PAPI_L1_ICM}
DOPS{Core, PAPI_DP_OPS}
INS_per_CYC{
  PU,
  PAPI_TOT_INS/PAPI_REF_CYC
}
```

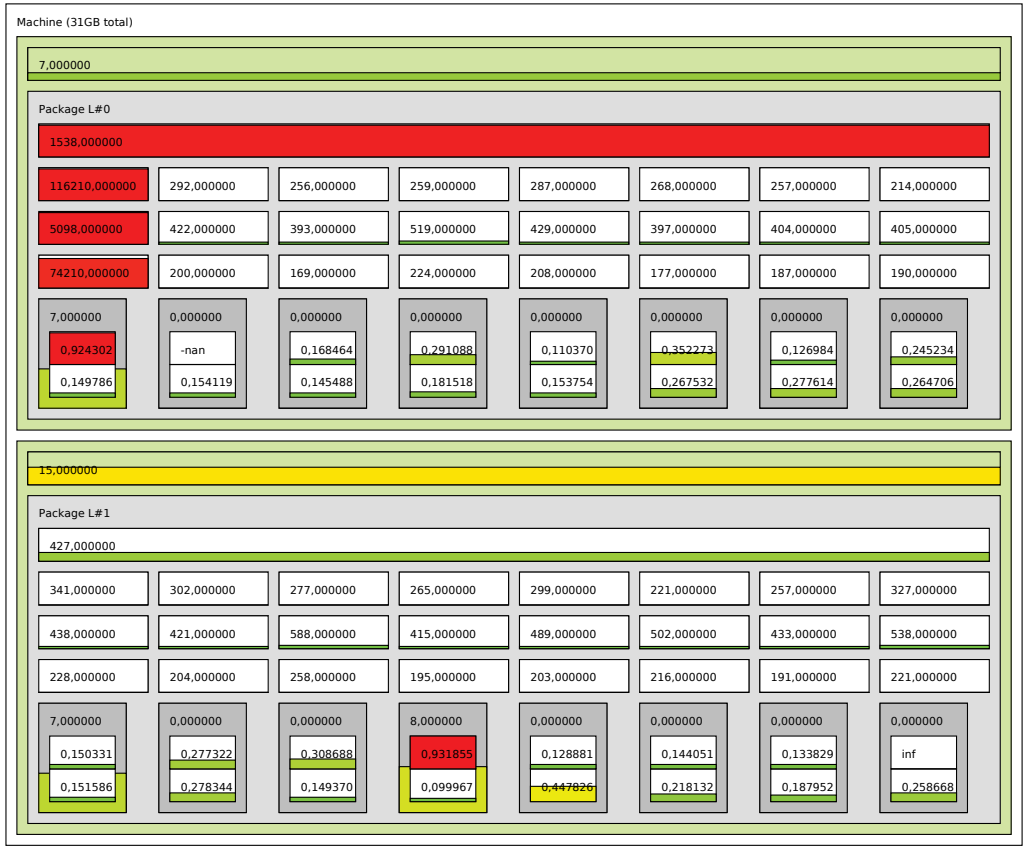

THANK YOU

Coloc

Grenoble

people.bordeaux.inria.fr/ndenoyel