RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

FRANCE 2030
PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

WP3 - Runtime Systems at Exascale

Samuel Thibault
University of Bordeaux

Runtime systems (WP3)

Exascale Challenges

Inter-node MPI communications

Intra-node PCI communications

- Memory/NIC/GPU

Task granularity

- GPU/CPU

Node loss

User feedback

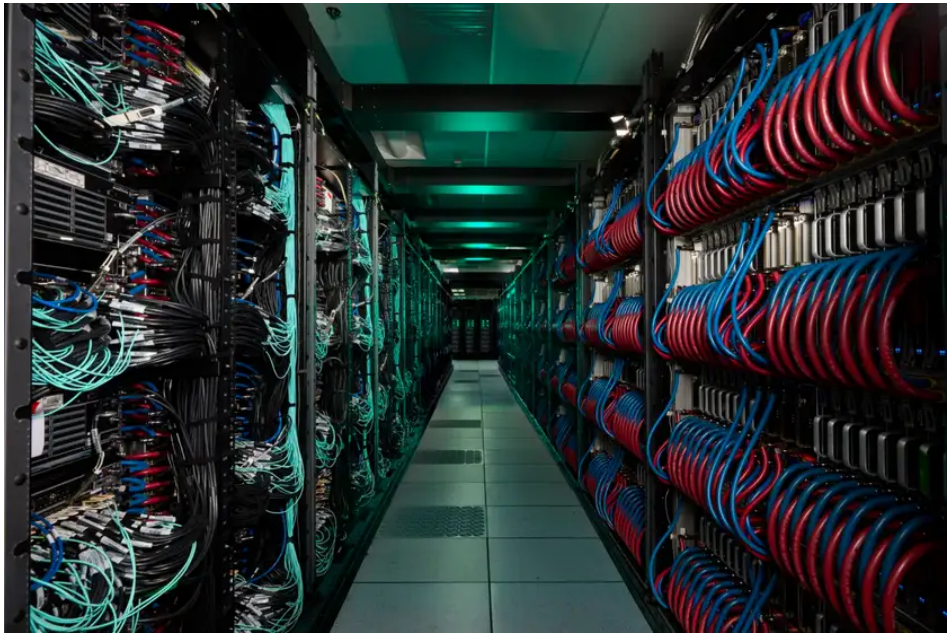

El Capitan

Runtime systems (WP3)

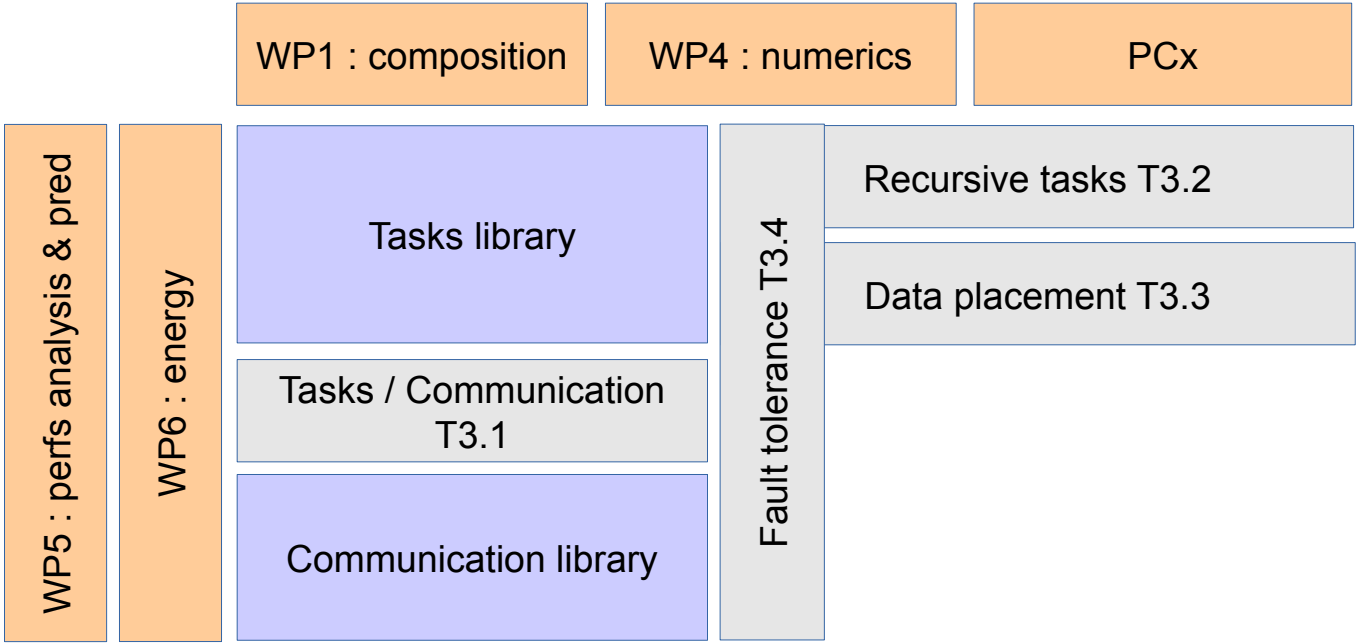

Runtime systems (WP3)

Inter-node MPI communications

Integrating communication task and communication asynchronism

- Tanguy Chatelain internship: Anticipatory notification of future communication → Almost-last-minute memory allocation+pinning

→ Effective allocation cache (a few GB) instead of permanent allocation (dozens of GB), without added latency

  - [https://inria.hal.science/hal-05147860]

- Will be pursued by Tristan Riehs, **PhD student** starting on oct 25 :
    - Prioritize communications, track lost time, collectives

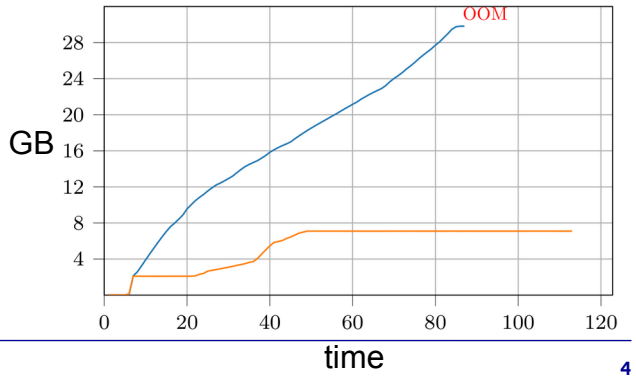

Trace analysis (with WP5)

- **PhD student** to be recruited with WP5 for sep 26

Runtime systems (WP3)

Intra-node PCI communications

Data placement in heterogeneous memory levels

- Optimize placement along DRAM/HBM/NVDIMM/GPU
    - Choose memory type
    - Balance between NUMA nodes
    - Anticipate allocations (e.g. for GPUDirect MPI reception)
- Atoli Huppé, **PhD student**, started in oct 24

Runtime systems (WP3)

Task granularity

Recursive Partitioning Autotuning

- Automatic adaptation of task granularity (GPU/CPU) + CPU-parallel tasks
- Typical 10-30% performance increase over static granularity
  - [https://doi.org/10.1016/j.jpdc.2025.105157]
- Thomas Morin, **PhD student**, will talk about it tomorrow morning
- Pierre Esterie, **eng.**, will help integrate in Chameleon from WP4, and other demos
- Samuel Mendosa, will start **PhD** on oct 25
- Integration with WP1's Comet (WP1 postdoc)
  information from WP2's compiler (WP2 PhD)

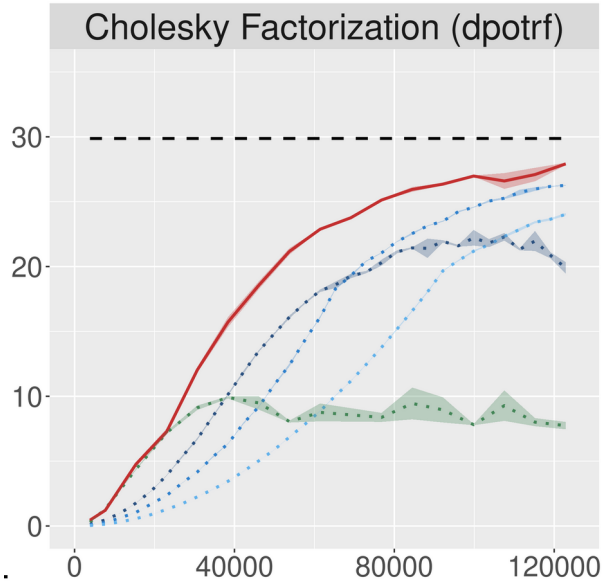

Distributed execution: task mapping & pruning

- Mapping at large tile size, and refine on nodes

Runtime systems (WP3)

Fault tolerance for large-scale systems

Fault-tolerance for large-scale systems with tasks

- Nicolas Ducarton, **eng.**, then **PhD student** started on apr 25, will briefly talk about it tomorrow morning
- Observe checkpoint/restart costs/opportunities
  - Low overhead of fully-asynchronous checkpointing (can be 0% !)
- Leverage state-of-the-art to optimize checkpoint frequency
- Integrate ABFT solutions with fault-tolerant task graph execution
- **Eng.** to be recruited to implement fault-tolerance MPI API in NewMadeleine

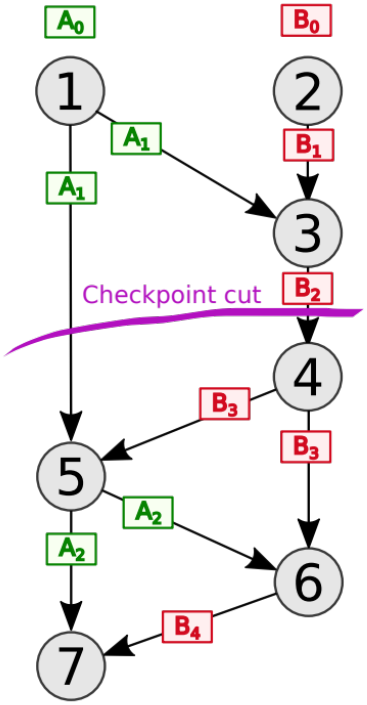

Runtime systems (WP3)

Explore multi-GPU parallelization strategies

- Mini-app from Airbus: noise propagation
  - Finite volumes on unstructured meshes (acoustic waves), C++
  - Explore Partitioning, edge coloring, halos, communication
  - Collaboration with NVIDIA

Interactive visualization and trace analysis tools

- See partitions, task affinities, ghost cells, cpu occupancy
- Training (Bordeaux, Paris Sorbonne, Telecom SudParis, IMT Atlantique)

Bridge the gap with programming models

- Mini-app from Total: mini GEOS (C++)
  - Collaboration with Total and PC1 (Hélène Barucq)
  - Raja, Kokkos, StarPU

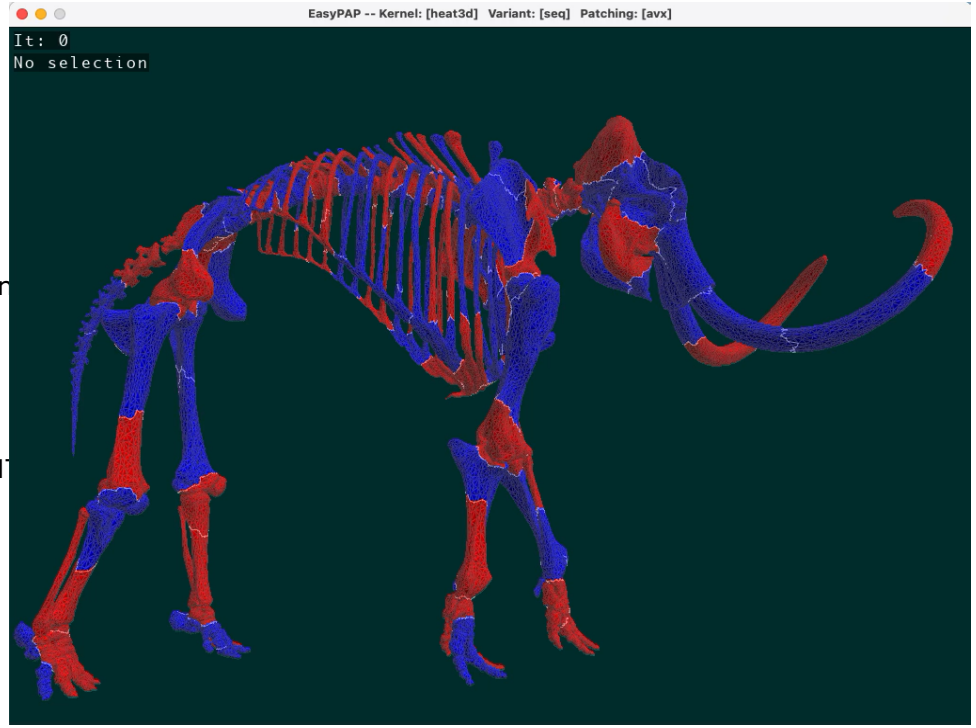

Runtime systems (WP3)

Integration in the Exa-SofT software stack

Software	Status	Integration in the software stack	Usable in other WP/PC ?
StarPU	Production	Tested on Chamelon Up to 4k cores	Yes
NewMadeleine	Production	Tested on Chamelon Up to 4k cores	Yes
Recursive tasks	Prototype	Tested on StarPU, Chamelon On 1 compute node	As prototype
Fault-tolerance	Prototype	Tested on StarPU On 4 compute nodes	No

Runtime systems (WP3)

Conclusion

Towards a scalable runtime-system stack

- Integrate optimizations, from tasks to network
  - Anticipate communication, govern data placement/allocation
  - Control work granularity
- Cope with faults

Next main steps

- Pursue current work
- Integrate with trace analysis/visualization
- Integrate with more demonstrators

RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

FRANCE 2030
PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

cea

cnrs

Inria