FP7 Support Action - European Exascale Software Initiative

DG Information Society and the unit e-Infrastructures

Addressing the Challenge of Exascale

European Exascale Software Initiative EESI

Towards Exascale roadmap implementation

EESI2 – Recommendations

Ultra Scalable Algorithms

S. Requena - WP3 leader

EESI2 recommendations

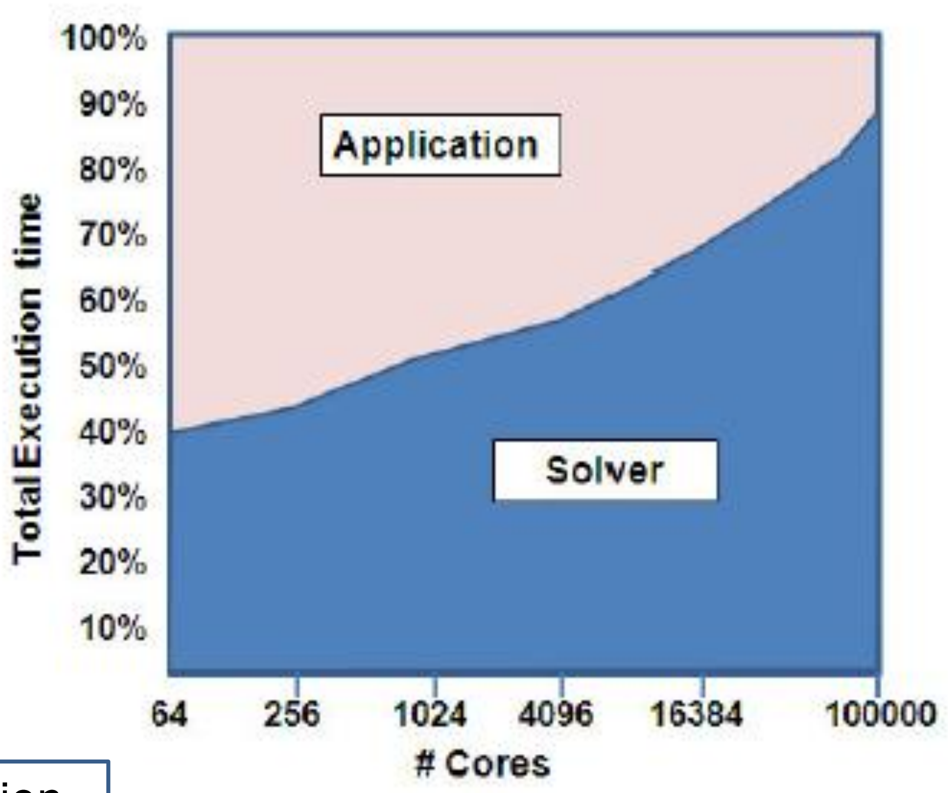

Ultra Scalable Algorithms are one of the 3 pilars of the EESI2 recommendations towards Efficient Exascale Applications.

Impact on a multiphysics application used on a multi-petascale system. Weak scaling on up to 100k cores
(courtesy of D. Lin)

Ultra Scalable Algorithms - Motivations

The Exascale Constraints

Systems	2012 -> 2015	2015 -> 2017	2018 -> 2022	Difference 2012 & 2018
System Peak [PF/s]	25	200	1000	O(50)
Power [MW]	6-20	15-50	20-80	O(5)
System Memory [PB]	0.3-0.5	5	32-64	O(100)
GB RAM/Core	0.5-2	0.2-1	0.1-0.5	5 times less
Node Performance [GF/s]	160-1000	500-7000	1000-10000	O(10)
Cores/Node	16-32	100-1000	1000-10000	O(50)-O(500)
Node memory BW [GB/s]	70	100-1000	400-4000	O(5) – O(50)
Number of nodes	10.000-100.000	5000-50.000	100.000-1.000.000	O(10)
Total concurrency	O(10^6)	O(10^7)	O(10^9)	O(1000)
MTTI	days	O(1 day)	O(1 day)	10 times less

Source: Rick Stevens and Andy White, IESP Meeting, Oxford 2010

Ultra Scalable Algorithms - Motivations

Energy Efficiency will require careful management of data locality

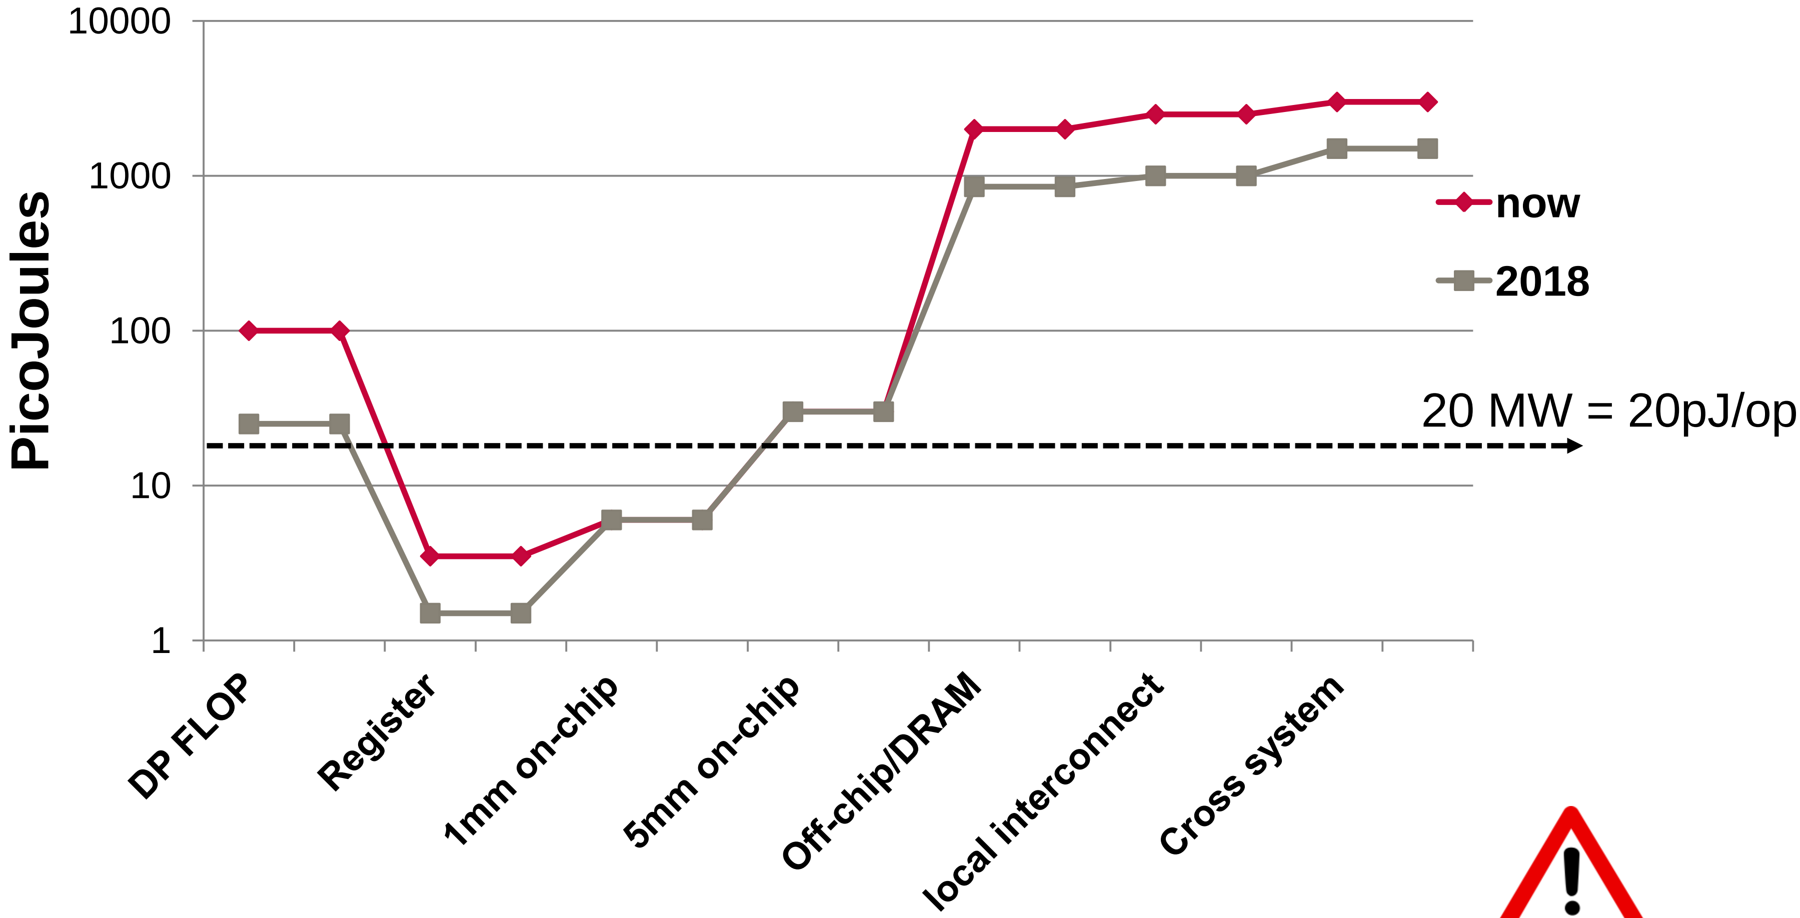

Source : John Shalf/LBNL- SC’12

Mandatory to stay at socket/node level

Ultra Scalable Algorithms - Recommandation

« Algorithms for Communication and Data-Movement Avoidance »

- Rationale : <u>Optimizing data placement and movement </u>will be key to performance, as well as a primary way for solvers to <u>reduce power consumption</u>.

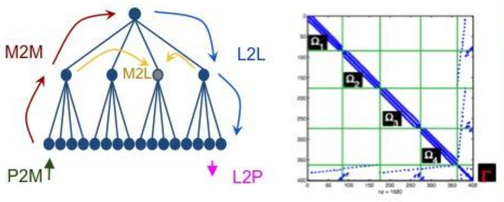

- Goals of the recommandation :
  - Design next gen of dense & sparse hybrid solvers
    - focused on communication reduction/avoiding,
    - use of hierarchical methods (H-Matrices, Fast Multipole Methods, …)
    - efficient paralleliation using directed acyclic graphs **(**DAG) and smart rutime over heterogeneous manycore nodes
  - Focus on operations that are at the intersection with the data mining community (low rank approximation of large matrix)
  - Enable leadership of European researchers in selected areas and allow to reach critical mass

Ultra Scalable Algorithms - Motivations

The Exascale Constraints (again)

Systems	2012 -> 2015	2015 -> 2017	2018 -> 2022	Difference 2012 & 2018
System Peak [PF/s]	25	200	1000	O(50)
Power [MW]	6-20	15-50	20-80	O(5)
System Memory [PB]	0.3-0.5	5	32-64	O(100)
GB RAM/Core	0.5-2	0.2-1	0.1-0.5	5 times less
Node Performance [GF/s]	160-1000	500-7000	1000-10000	O(10)
Cores/Node	16-32	100-1000	1000-10000	O(50)-O(500)
Node memory BW [GB/s]	70	100-1000	400-4000	O(5) – O(50)
Number of nodes	10.000-100.000	5000-50.000	100.000-1.000.000	O(10)
Total concurrency	$O(10^6)$	$O(10^7)$	$O(10^9)$	O(1000)
MTTI	days	O(1 day)	O(1 day)	10 times less

Source: Rick Stevens and Andy White, IESP Meeting, Oxford 2010

Ultra Scalable Algorithms - Recommandation

« Parallel-in-Time: a fundamental step forward in Exascale Simulations »

- Rationale
  - Spatial decomposition strategies are not enough for exploiting all the massive amount of concurrency of Exascale systems
    - Use of the parallelision accros the time dimension
    - Potential application areas include: climate research, CFD, life sciences, materials science, nuclear engineering, etc
    - European researchers are leading Parallel-in-Time developments

- Goals
  - Establish of multi-disciplinary consortia to co design the deployment of Parallel-in-Time methods, encapsulated in reusable scalable libraries
  - Establish a series of benchmarks for Pro/Cons of Parallel-in-Time methods
  - Fund 2 to 4 internatonal projects between €2M and €4M

Ultra Scalable Algorithms – The Session

After the apetizer, lets go for the main course

- *« Algorithms for Communication and data-movement Avoidance »* by L. Grigori (Inria, France)

- *« Parallelisation in Time with examples on applications »* by M. Bolten (Univ. of Wuppertal, Germany)

- *« User presentation on Exascale Communication hiding/avoiding with examples on applications »* by W. Vanroose (Univ. of Antwerp, Belgium)