

# Supercomputers and Beskow

*some background for today's visit to PDC*

Stefano Markidis

HPCViz Department

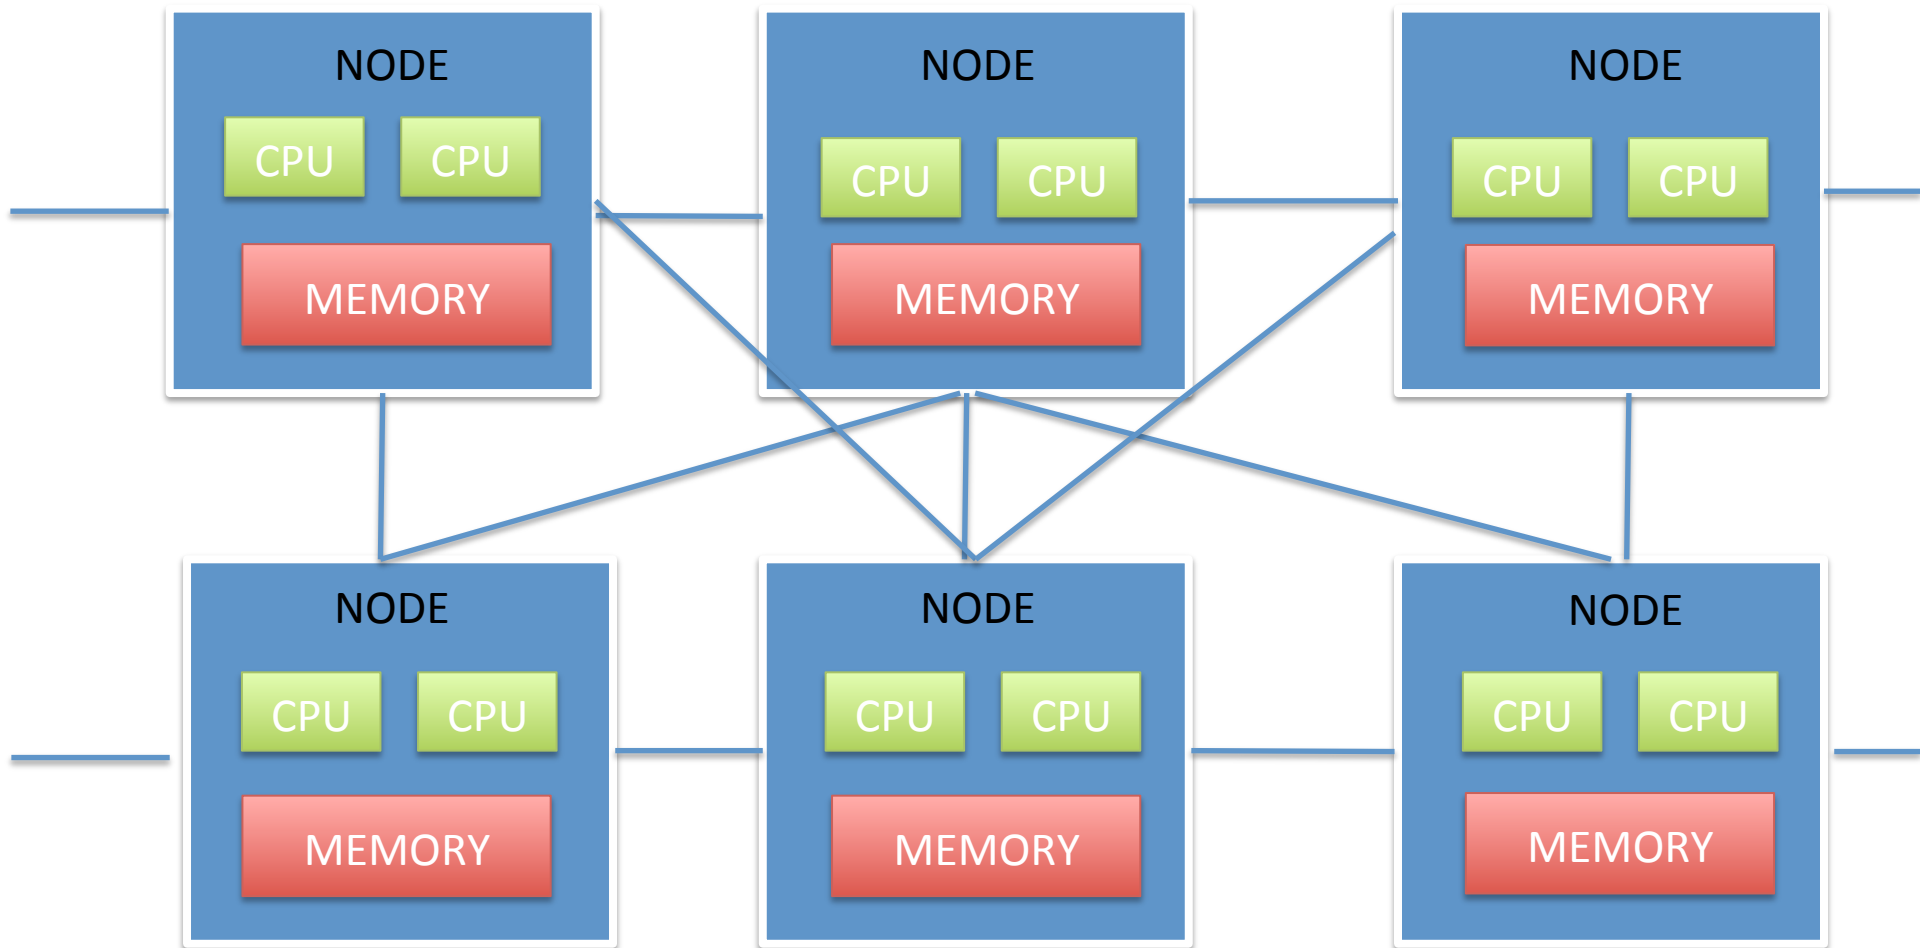
KTH Royal Institute of Technology



# Outline

- What is a Supercomputer ?
- How you rank them ?
- The Power Wall
- The Beskow supercomputer
- Simulations on Beskow

# What is a Supercomputer ?



IT IS A NETWORK OF (SEVERAL) COMPUTING NODES. EACH COMPUTING NODE CONSISTS OF 1 OR MORE CPUS AND ONE SHARED (AMONG NODE CPUS) MEMORY. EACH CPU HAS 1 OR MORE CORES (NOWDAYS YOU CAN HAVE UP TO 64 CORES).

# What Makes a Supercomputer a Supercomputer (instead of your office network) ?

- The use of **high-speed dedicated** interconnection (between nodes) network.
- How come is network is important ? How fast you compute? How fast you can access your memory on node ? How fast is communication ? Communication across network is slowest process!



MARE NOSTRUM SC

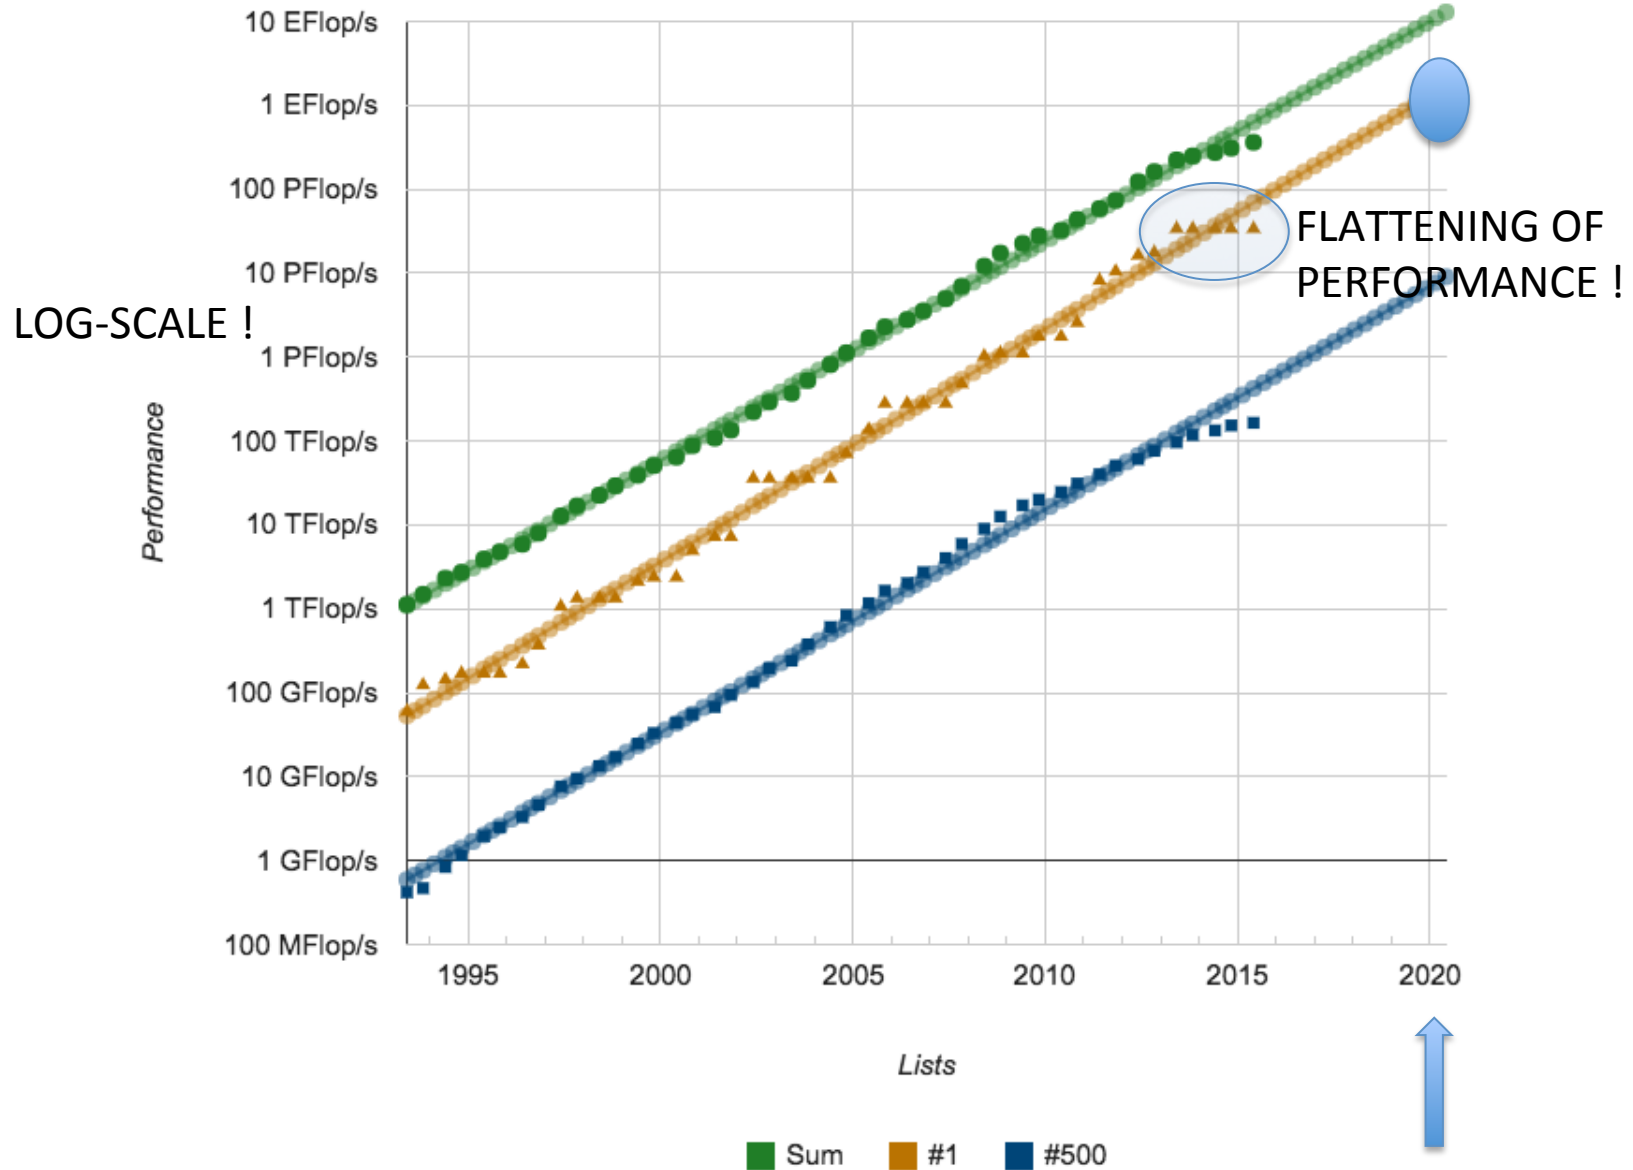
# How You Rank a Supercomputer ?

- The main metric is the number of floating point operations (in double precision) per seconds (FLOPS) a machine can crunch:
  - Theoretical peak performance (RPEAK) : take the FLOPS a single processor X number of processors in the systems (NOT CONSIDERING COMMUNICATION COSTS AND MEMORY ACCESS)
  - Maximal LINPACK Performance (RMAX): Calculate the FLOPS when solving a dense linear system with an optimized library called Linpack. We rank supercomputers using these values.
- The ranking of the supercomputers is published on the <http://www.top500.org/> . The list is updated every twice per year at the 2 main SC conferences.

# TERA – PETA – EXA FLOPS MACHINE

- The first supercomputer to be capable of doing 1 Tera ( $10^{12}$ ) FLOPS was ASCI RED in 1996 at Sandia. Now we have a TFLOPS on a chip!
- The first supercomputer to be capable of doing 1 Peta ( $10^{15}$ ) FLOPS was Roadrunner in 2008 at LANL.
- Fastest supercomputer now is TIANHE2 in China. It delivers 33 PFLOPS. Leading the ranking for almost four years now (US doesn't like this, see Obama's executive order on July 2015 for exaflop machine by 2025). Beskow is 1.4 PF (#43).
- Who is going to be the first to reach the ExaFLOPS? → Race to exascale

# The Race to Exascale



# The Power Wall

- TIANHE-2 is # 1 with 33 PFLOPS and it uses 18 MW!
- How much power an exascale machine (33 x TIANHE-2) will use ?
- Exascale machine can't consume more than 20 MW. With THIAN-2 almost hit the wall.
  - Power scales with cube of CPU frequency (slower CPU are better but put more!)
  - RAM Memory operation consumes memory (new memory type)





# KTH Beskow Supercomputer

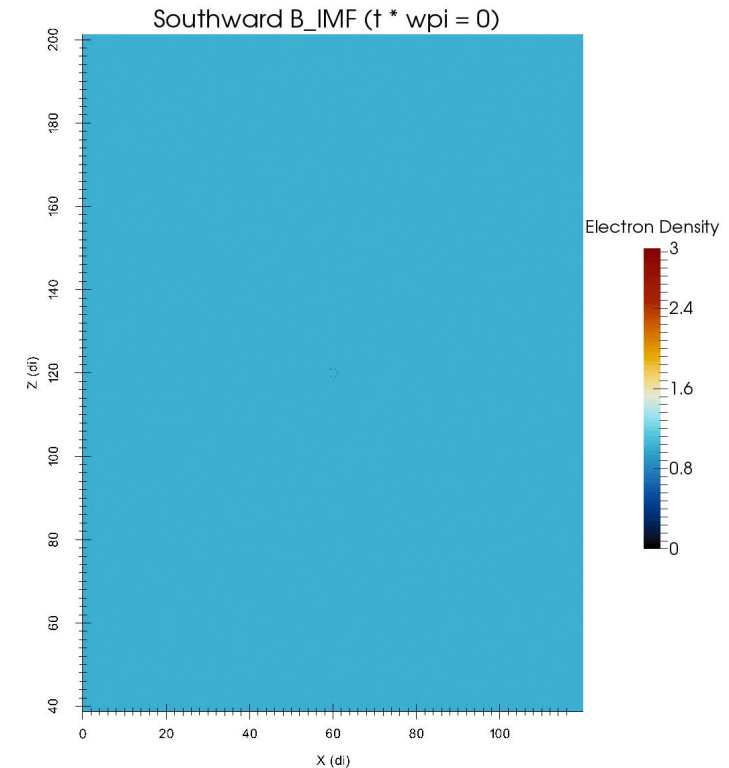
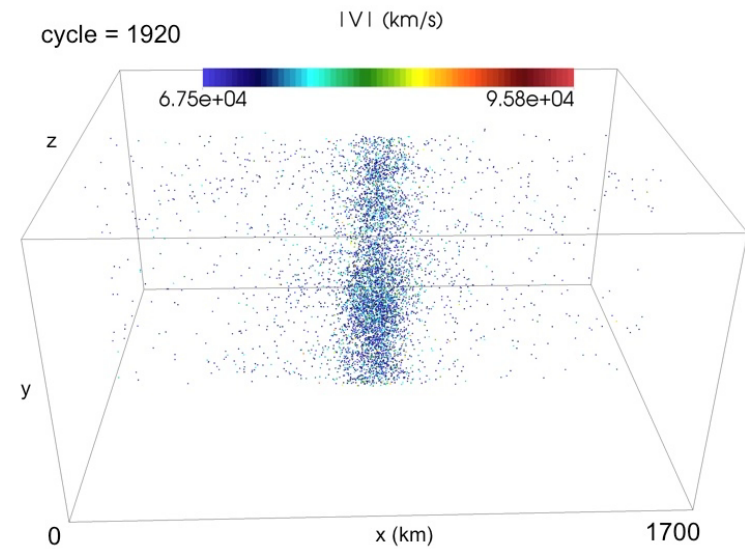
- Fastest supercomputer in Nordic Countries
- Named after “Elsa Beskow”, painter and writer.

42	Air Force Research Laboratory United States	<b>Spirit</b> - SGI ICE X, Xeon E5-2670 8C 2.600GHz, Infiniband FDR SGI	73,584	1,415.5	1,530.5	1,606
43	KTH - Royal Institute of Technology Sweden	<b>Beskow</b> - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	53,632	1,397.0	1,973.7	
44	CEA/TGCC-GENCI France	<b>Curie thin nodes</b> - Bullx B510, Xeon E5-2680 8C 2.5GHz, Infiniband FDR Bull	77,184	1,359.0	1,667.2	2,251



# Simulations on Beskow

- The developer team of the iPIC3D code focusing on algorithms and parallel implementations is at KTH.
- iPIC3D code is used to simulate magnetic reconnection and magnetosphere physics. KU Leuven and University of Colorado are the main users of the iPIC3D code.
- iPIC3D runs routinely on 2K cores of Beskow supercomputer.



# Conclusions

Ready to meet Beskow? It is 15-20 minutes from here

