How to run VASP

Part 1: Influential settings

Quick summary

- **Processor cores ⩽ atoms**
- **NCORE** = cores per node
  NPAR = compute nodes
- **NSIM** = 1 (Intel), 4 (AMD)
- **KPAR** = compute nodes
- **ALGO** = Fast (or VeryFast)
- **NBANDS** = *N**cores
- **Use the right VASP binary**
  (*NGZhalf*, and *wNGZhalf* for gamma-point calculations)

Q: What's missing?

A: Read the VASP manual!

Efficiency:
Running as many jobs as possible within a given allocation of computer time.

Speed:
The amount of time (in real time, "human time") to run a specific simulation from the time it starts.

Time-to-solution:
The amount of time (in real time, "human time") it takes to get the results = runtime + the time waiting in the queue.

Efficiency

Using as few cores as possible

Speed

As many cores as possible (up to a limit)

Time-to-solution

- *Finding the "right" size which fits the cluster.*
- *Not using up all your allocated time (in order to get higher priority in the queue).*

VASP scaling on Carter

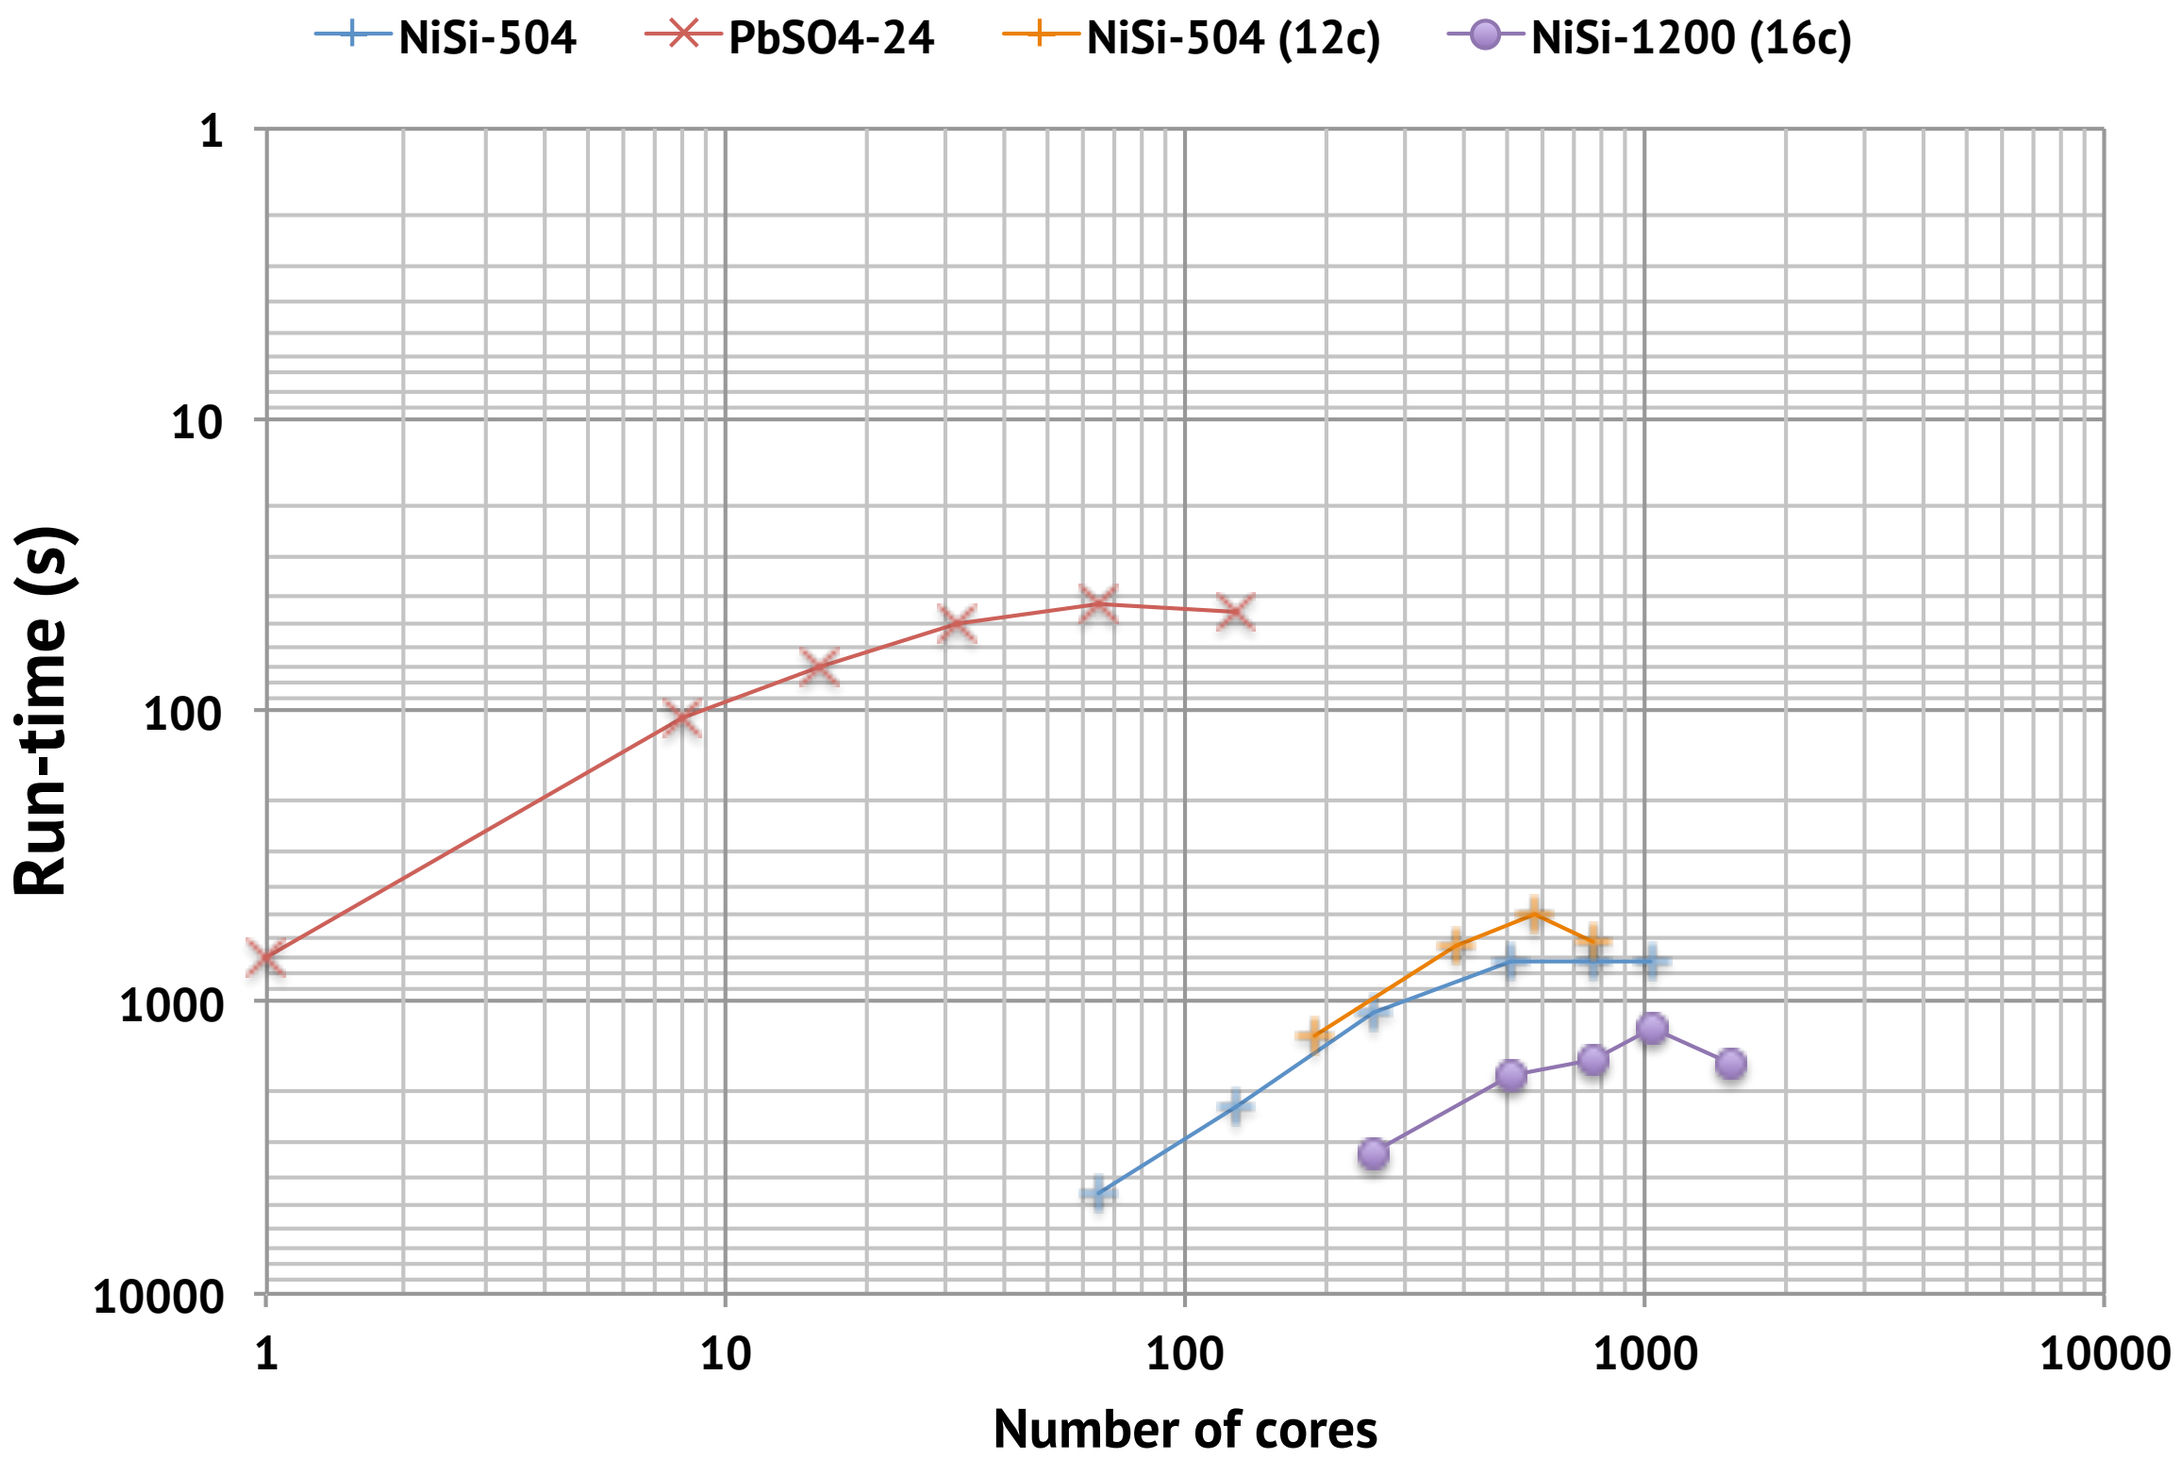

Safe & Efficient:
cores < n(atoms)

Daring & not as efficient:
cores ≈ 2 x n(atoms)

Crash!
cores > n(bands)

Why should I care?

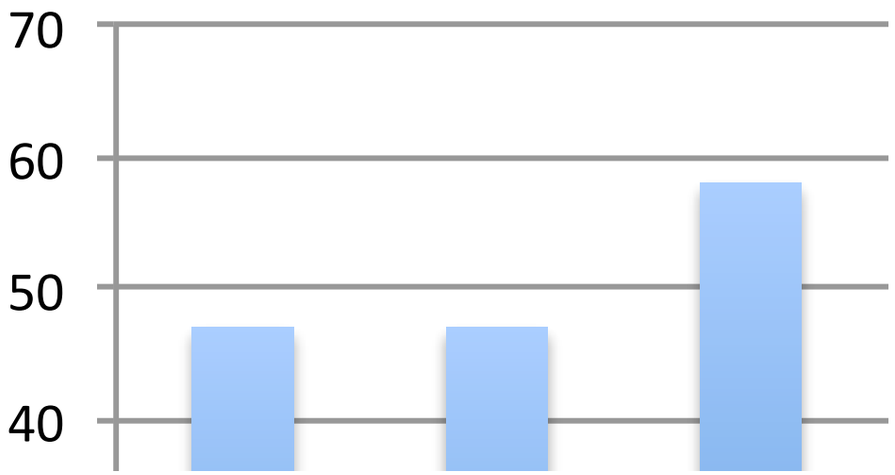

Does 10% matter?

Dear Professor N.N.,

We regret to inform you that we have decided to cut your VSC time allocation from 100,000 core hours/month to 90,000 core hours/month.

Best regards,
VSC Staff

NPAR & NCORE

- Applies to the Davidson/RMM-DIIS algorithms (not hybrid calculations)

- Controls/activates band parallelization

- NPAR = 1: no parallelization over bands = BAD

- NPAR = (1-2)x number of compute nodes

- NCORE = cores working on one band

NCORE = `cores per node (or socket)`

Li2FeSiO4 (128 atoms) on Triolith

Standard DFT / 4 compute nodes / 64 cores

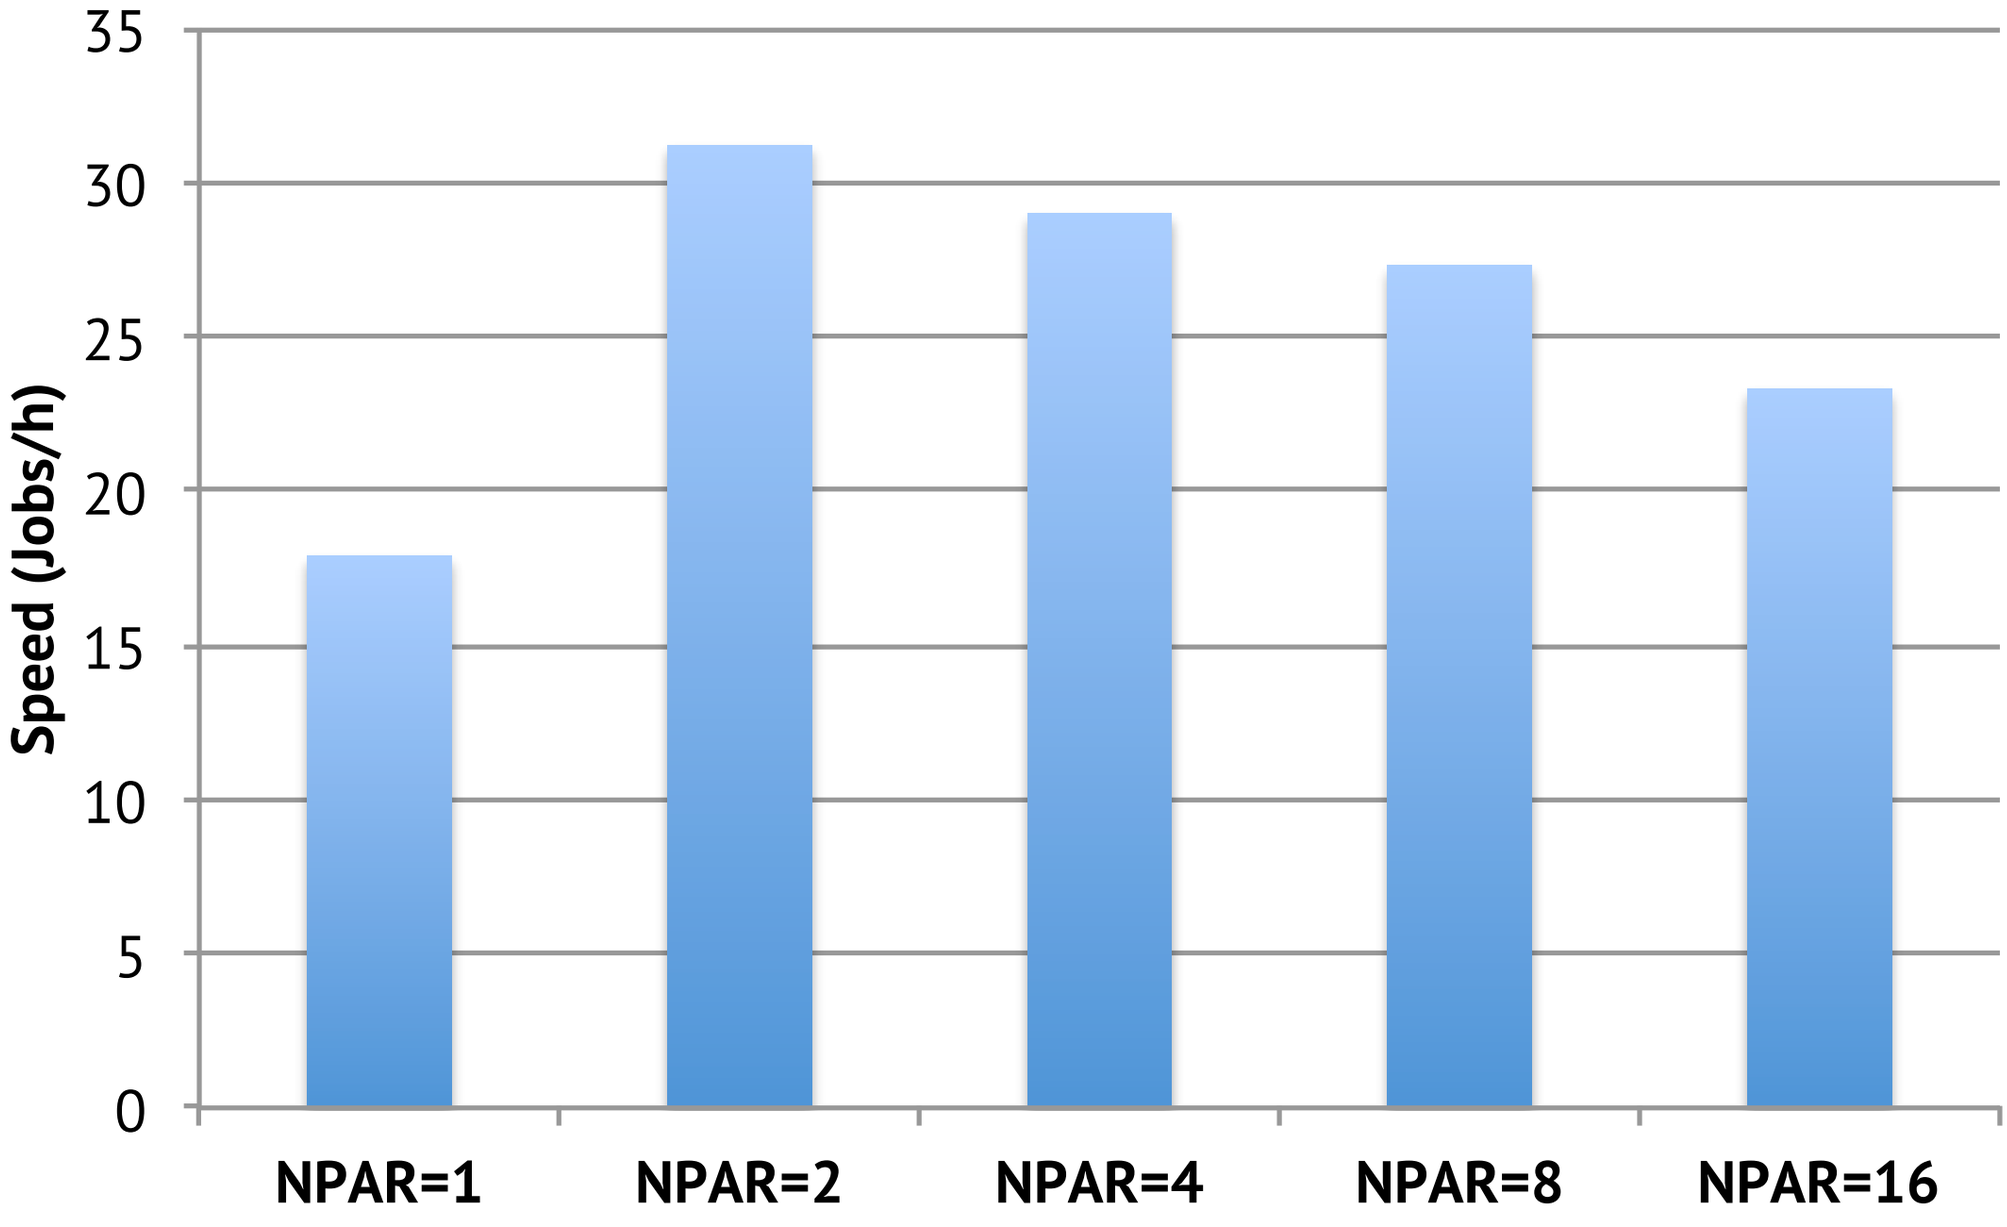

Speed-up with optimal NPAR

FeCo alloy (53 atoms)

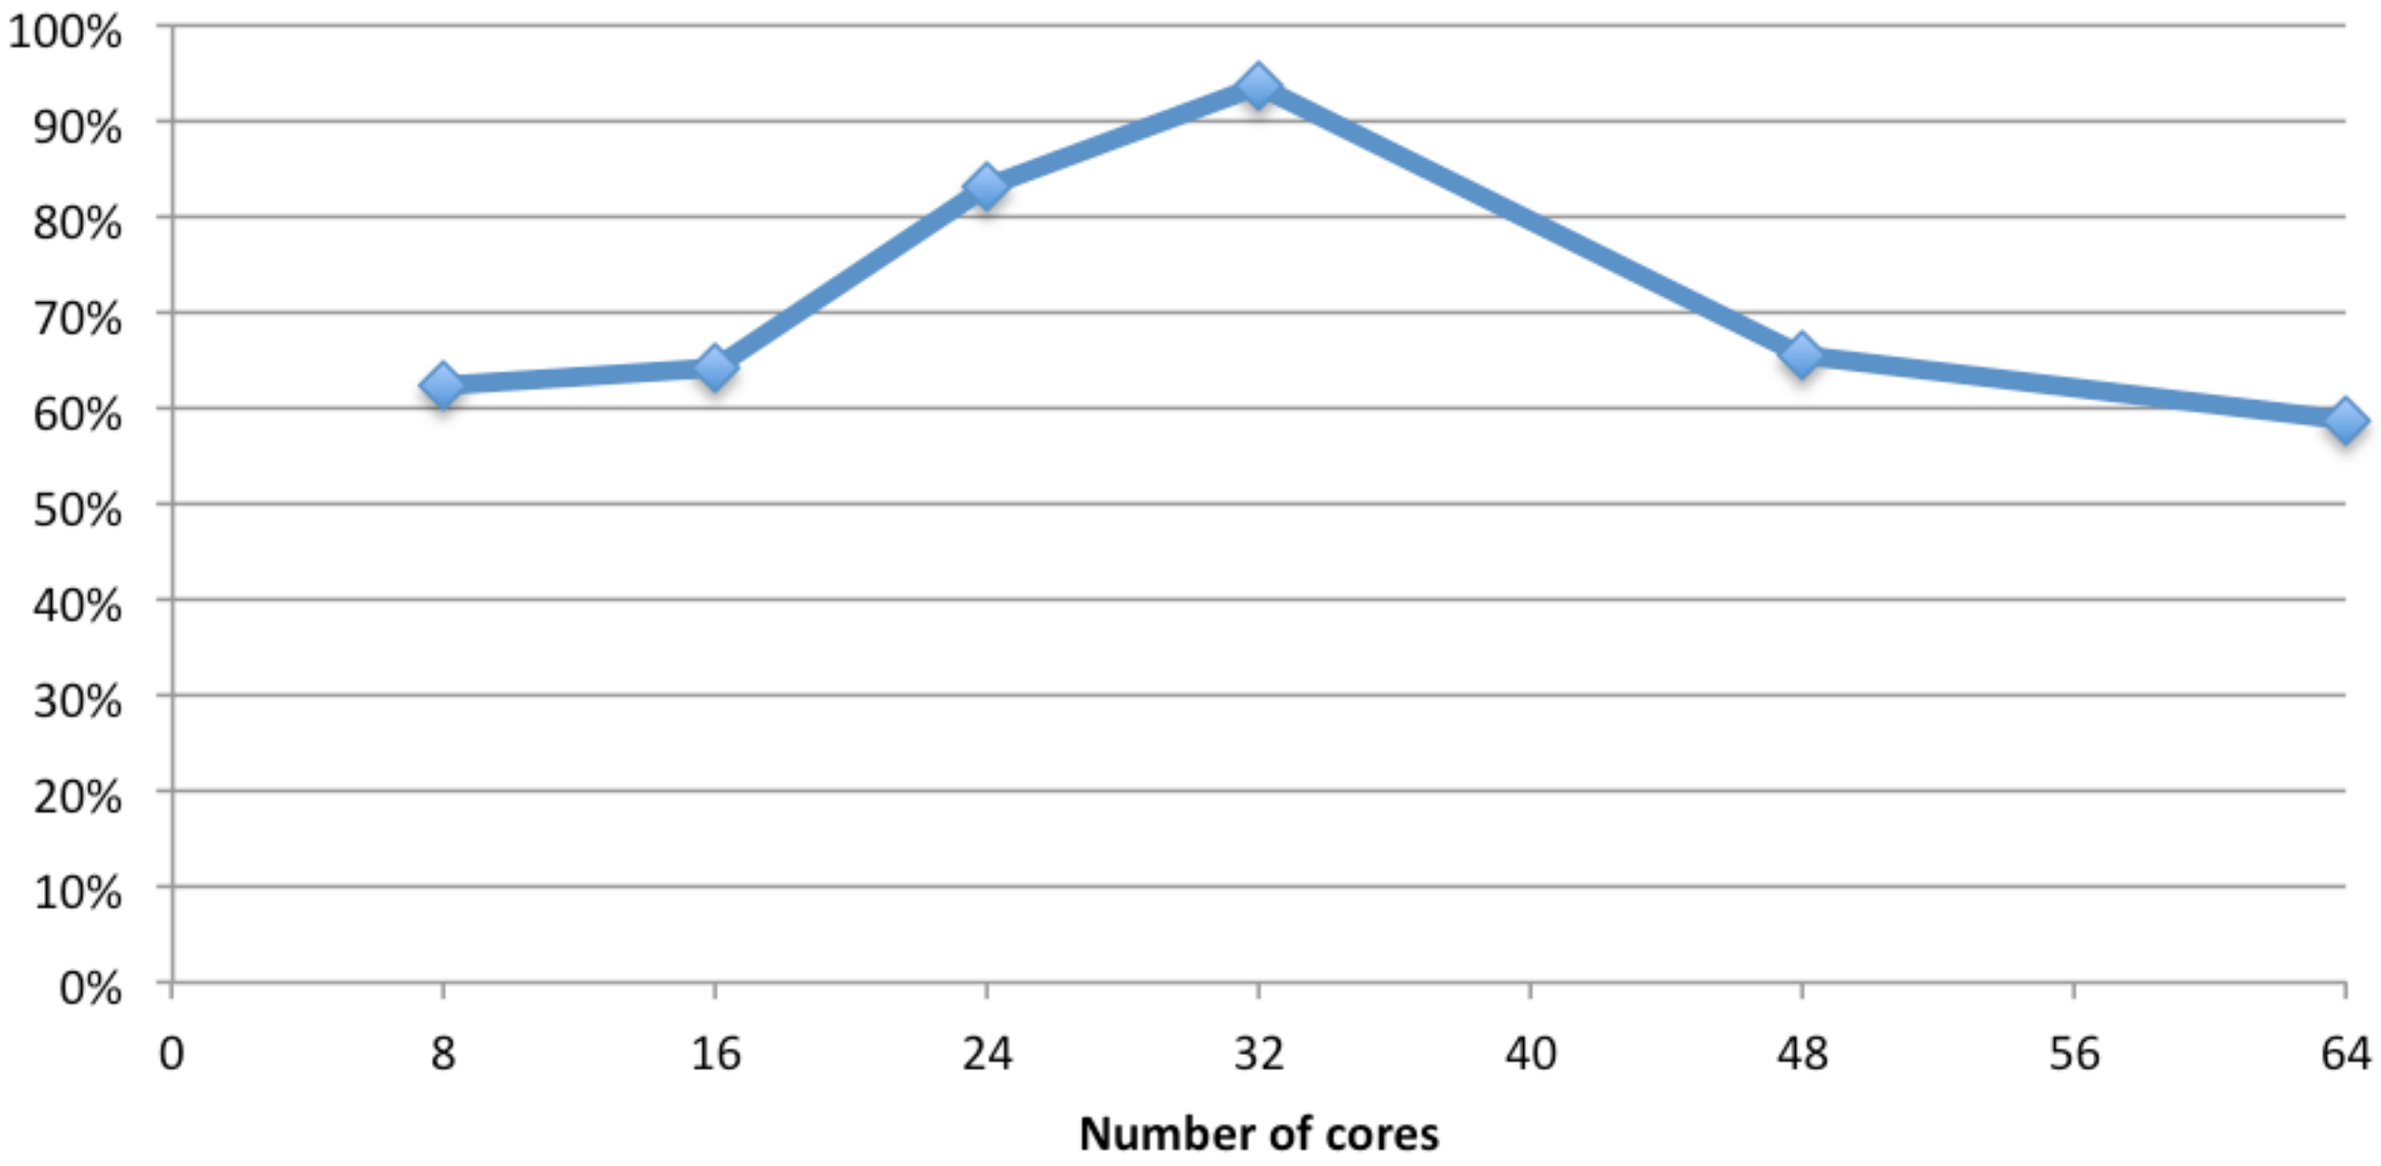

NPAR: caveat

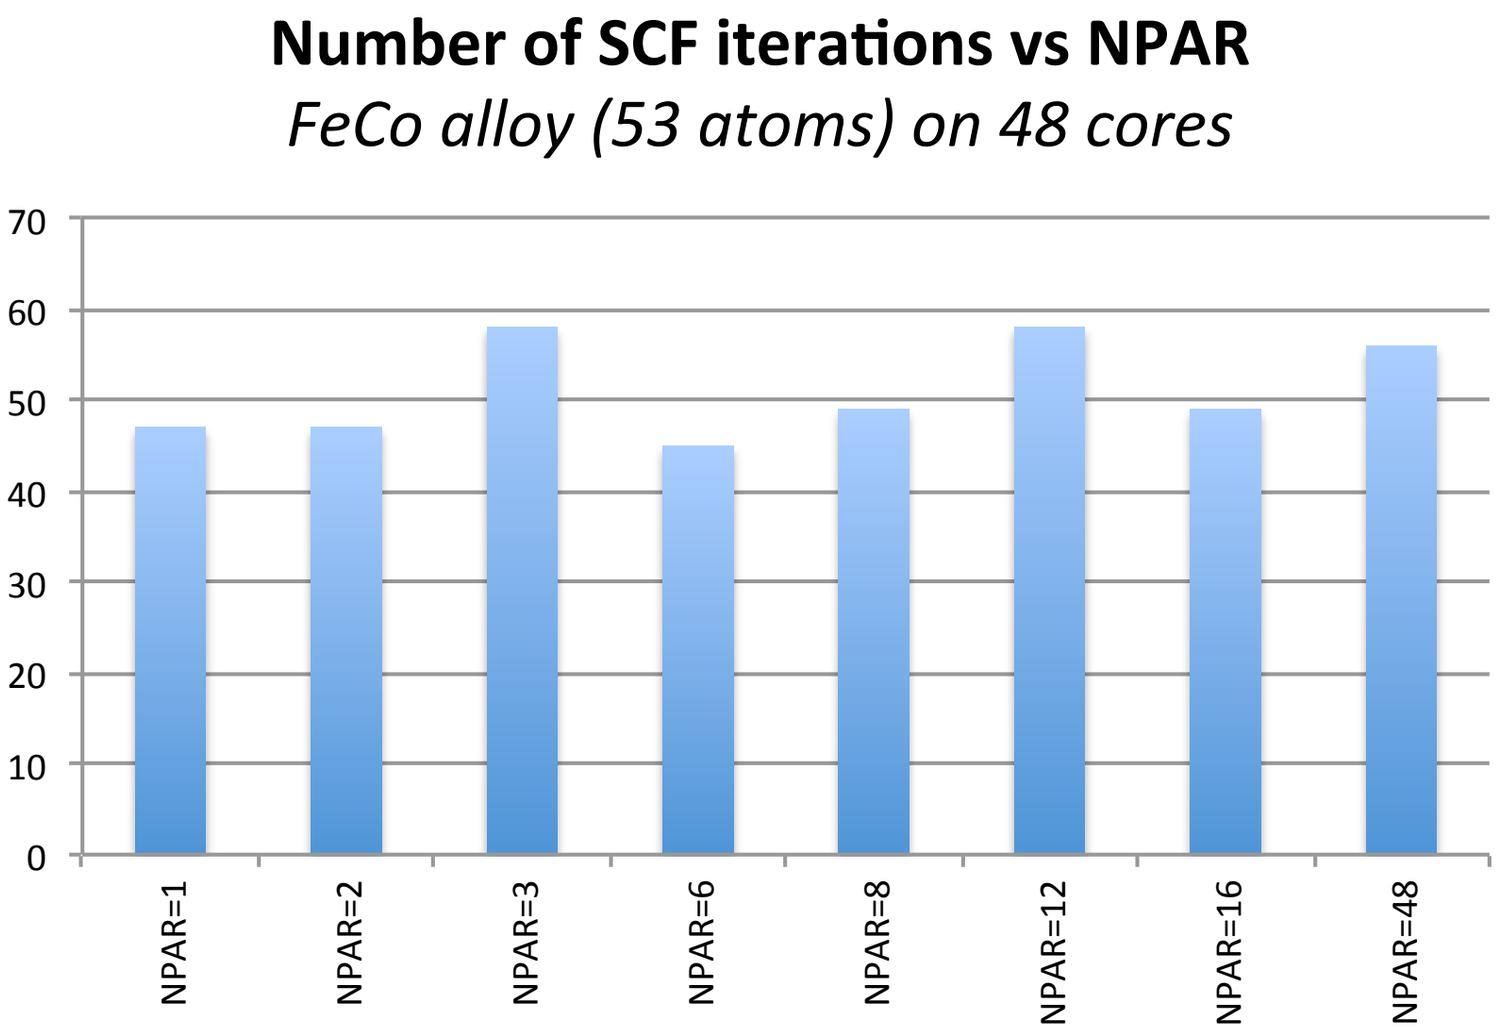
Number of SCF iterations vs NPAR
FeCo alloy (53 atoms) on 48 cores
70
60
50
40
30
20
10
0
NPAR=1
NPAR=2
NPAR=3
NPAR=6
NPAR=8
NPAR=12
NPAR=16
NPAR=48

K-point parallelization ("KPAR")

- Controls parallelization over k-points (version 5.3.2+)
- **Great for hybrid-DFT jobs**
- Not so great for metals with 100-1000s of k-points…

KPAR = *"number of k-points treated in parallel"*
cores / KPAR = *"group size"* = *cores/node?*

KPAR = `min(nodes,#kpts)`

MgO (63 atoms) on Triolith

HSE06 hybrid calculation with 4 k-points

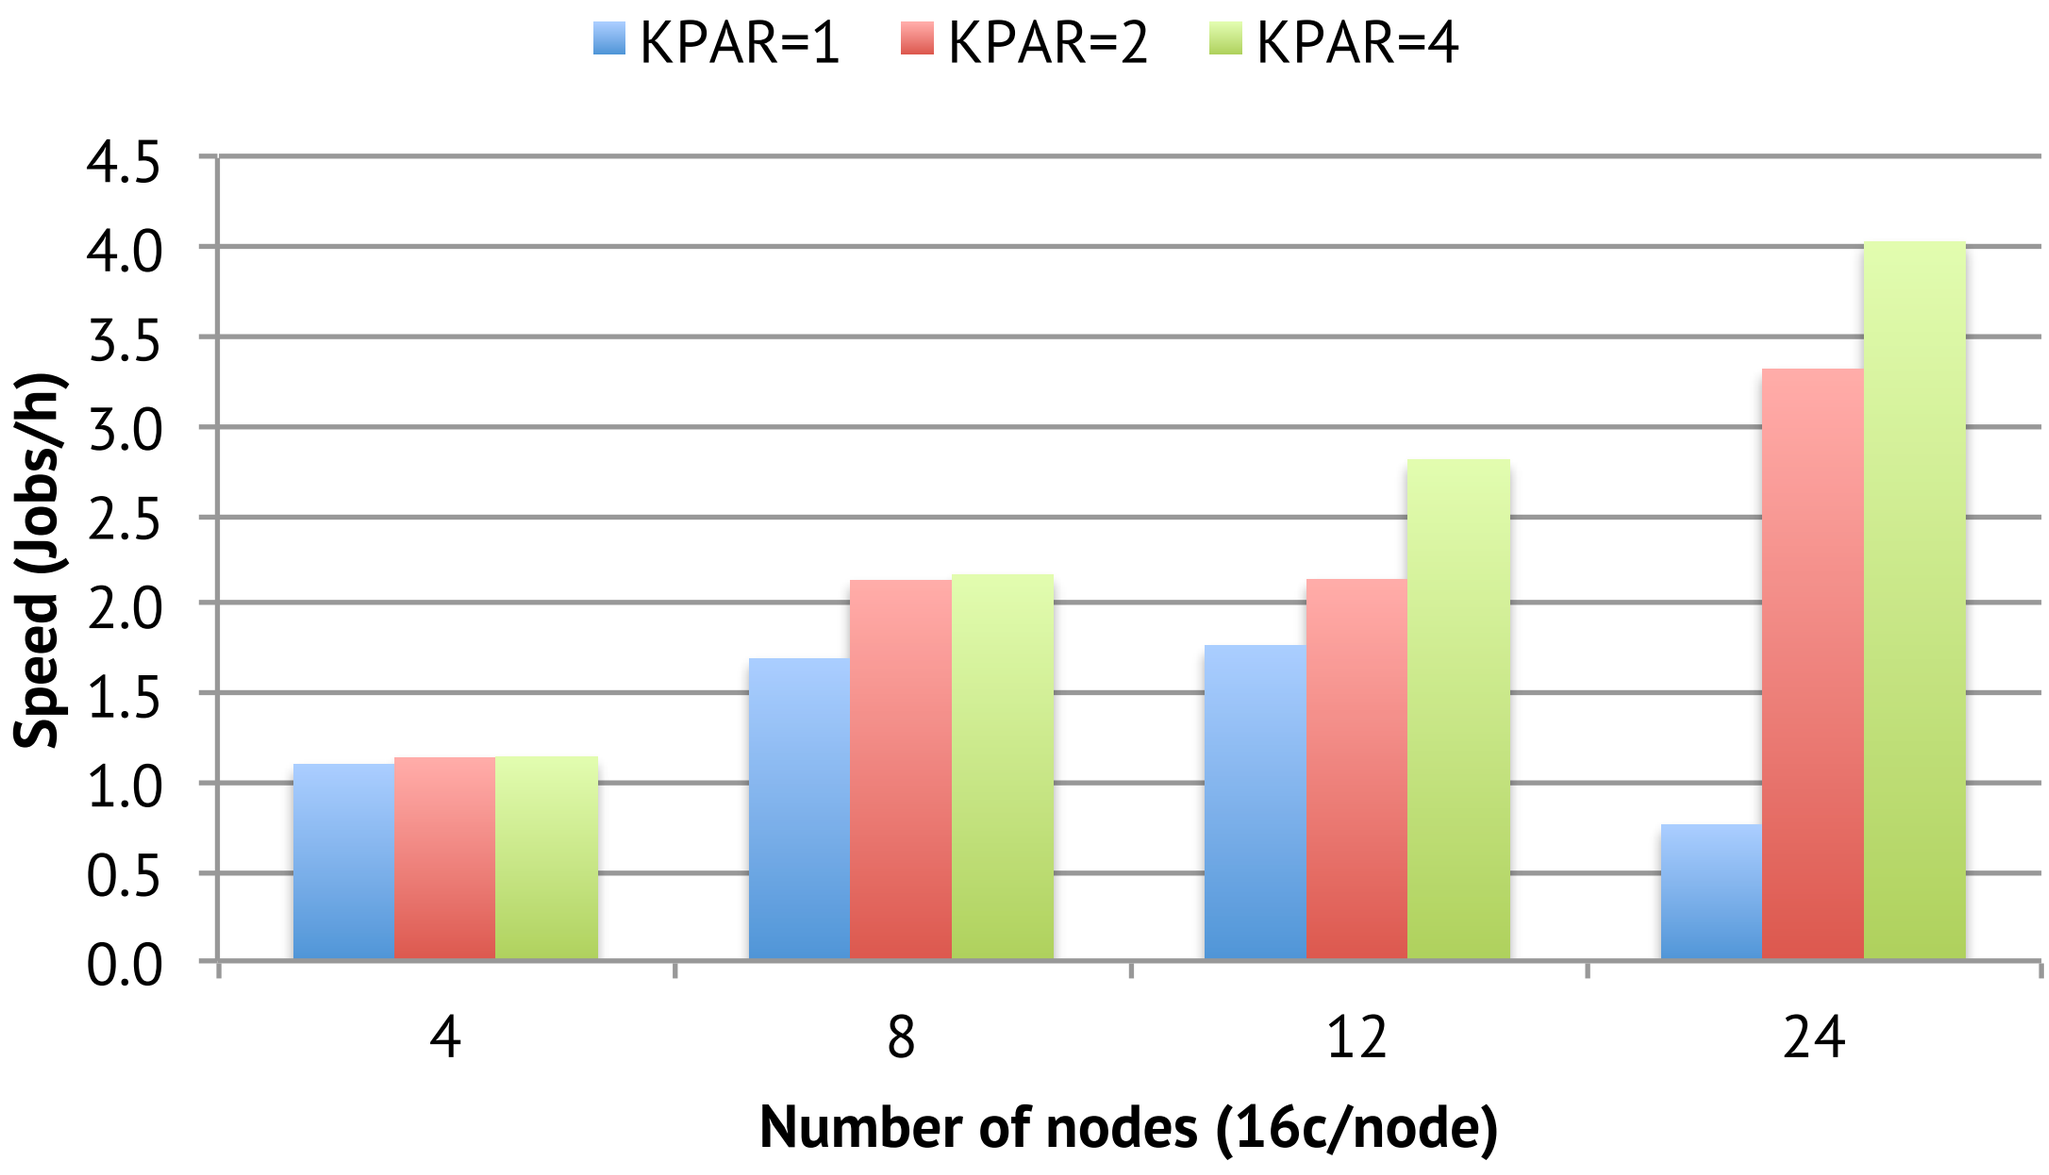

NSIM

- Blocking mode for RMM-DIIS algorithm
  (so only applies for ALGO=Fast/VeryFast)

- Effect depends on BLAS library. Less important nowadays?

- NSIM = 1 for single-node jobs on Intel Sandy Bridge

- NSIM = 4 for older systems (incl. AMD). This is VASP's default.

Li2FeSiO4 (128 atoms) on Triolith

Standard DFT / 4 compute nodes / 64 cores

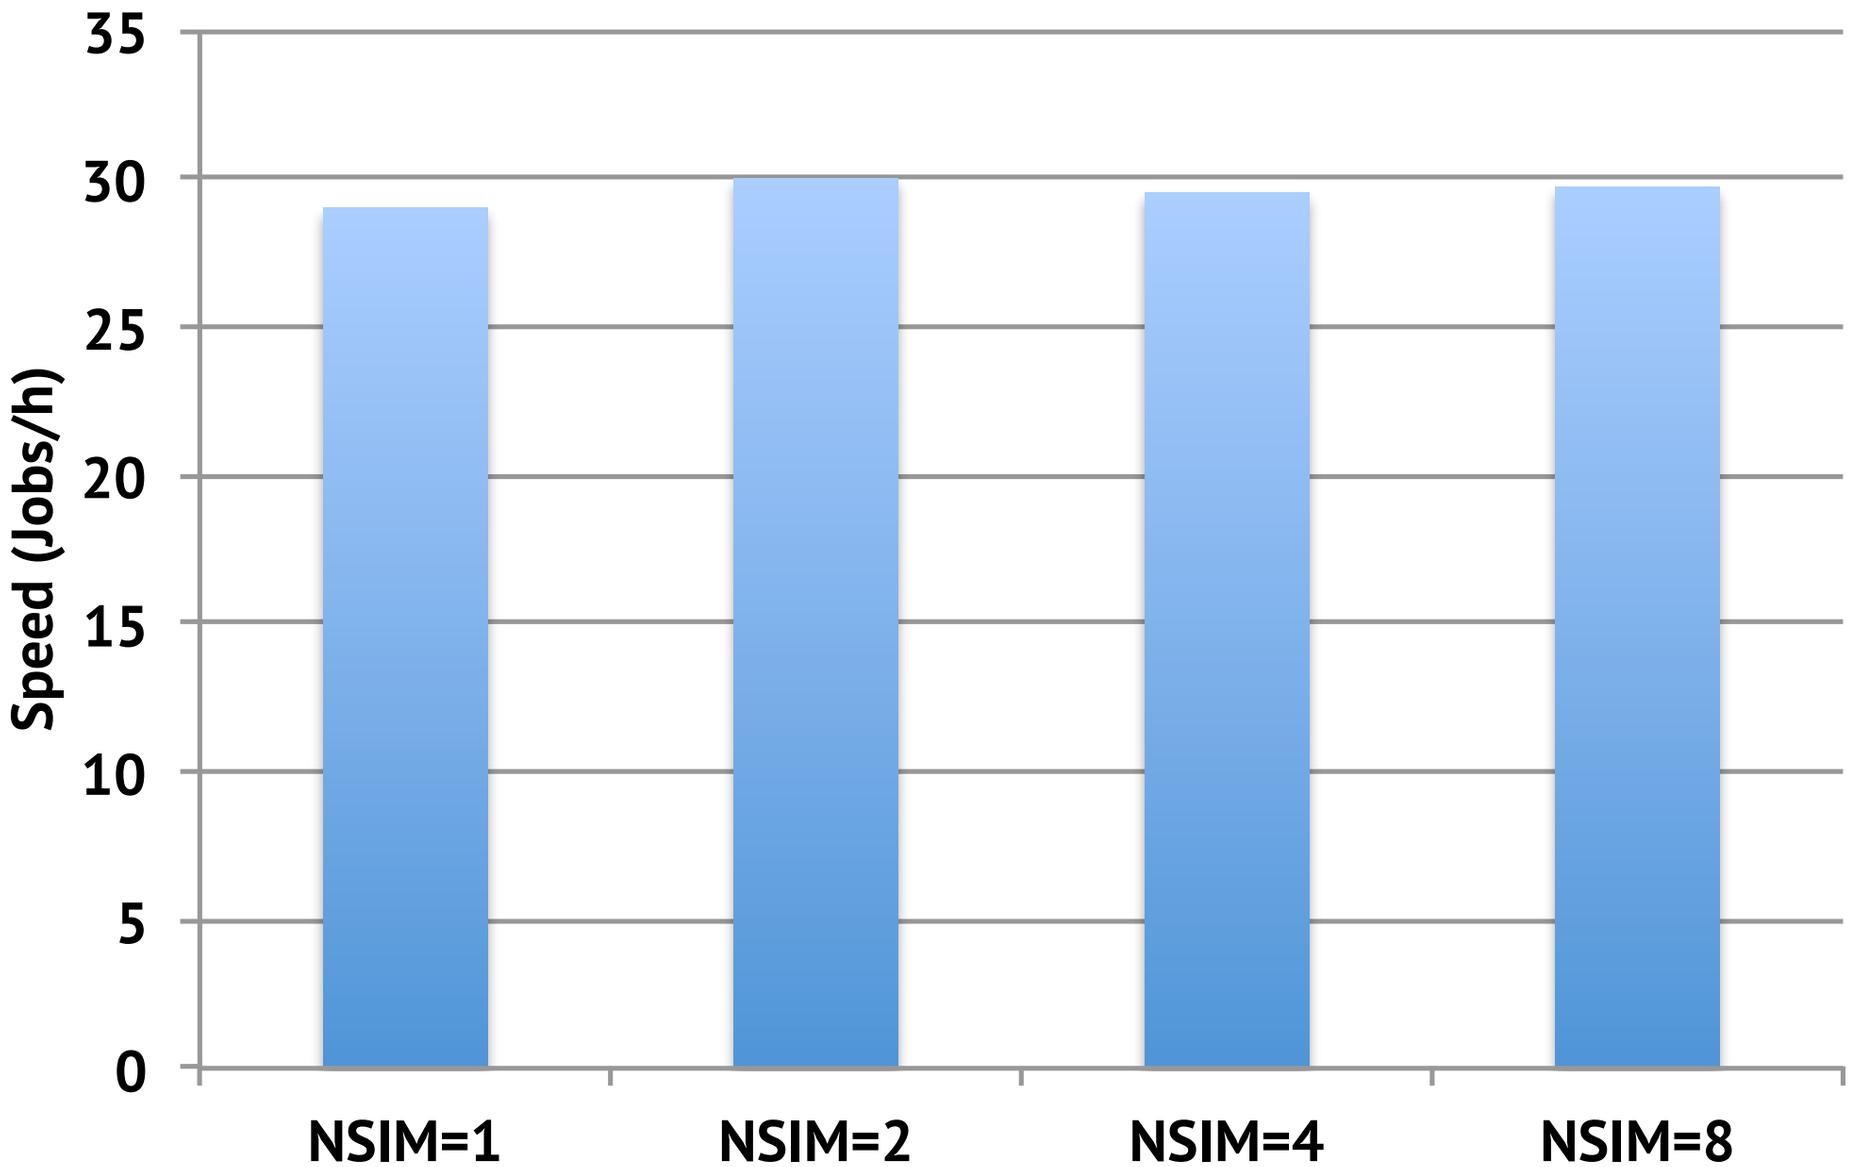

PbSO4 (24 atoms) on Triolith

Standard DFT / 1 compute node / 16 cores

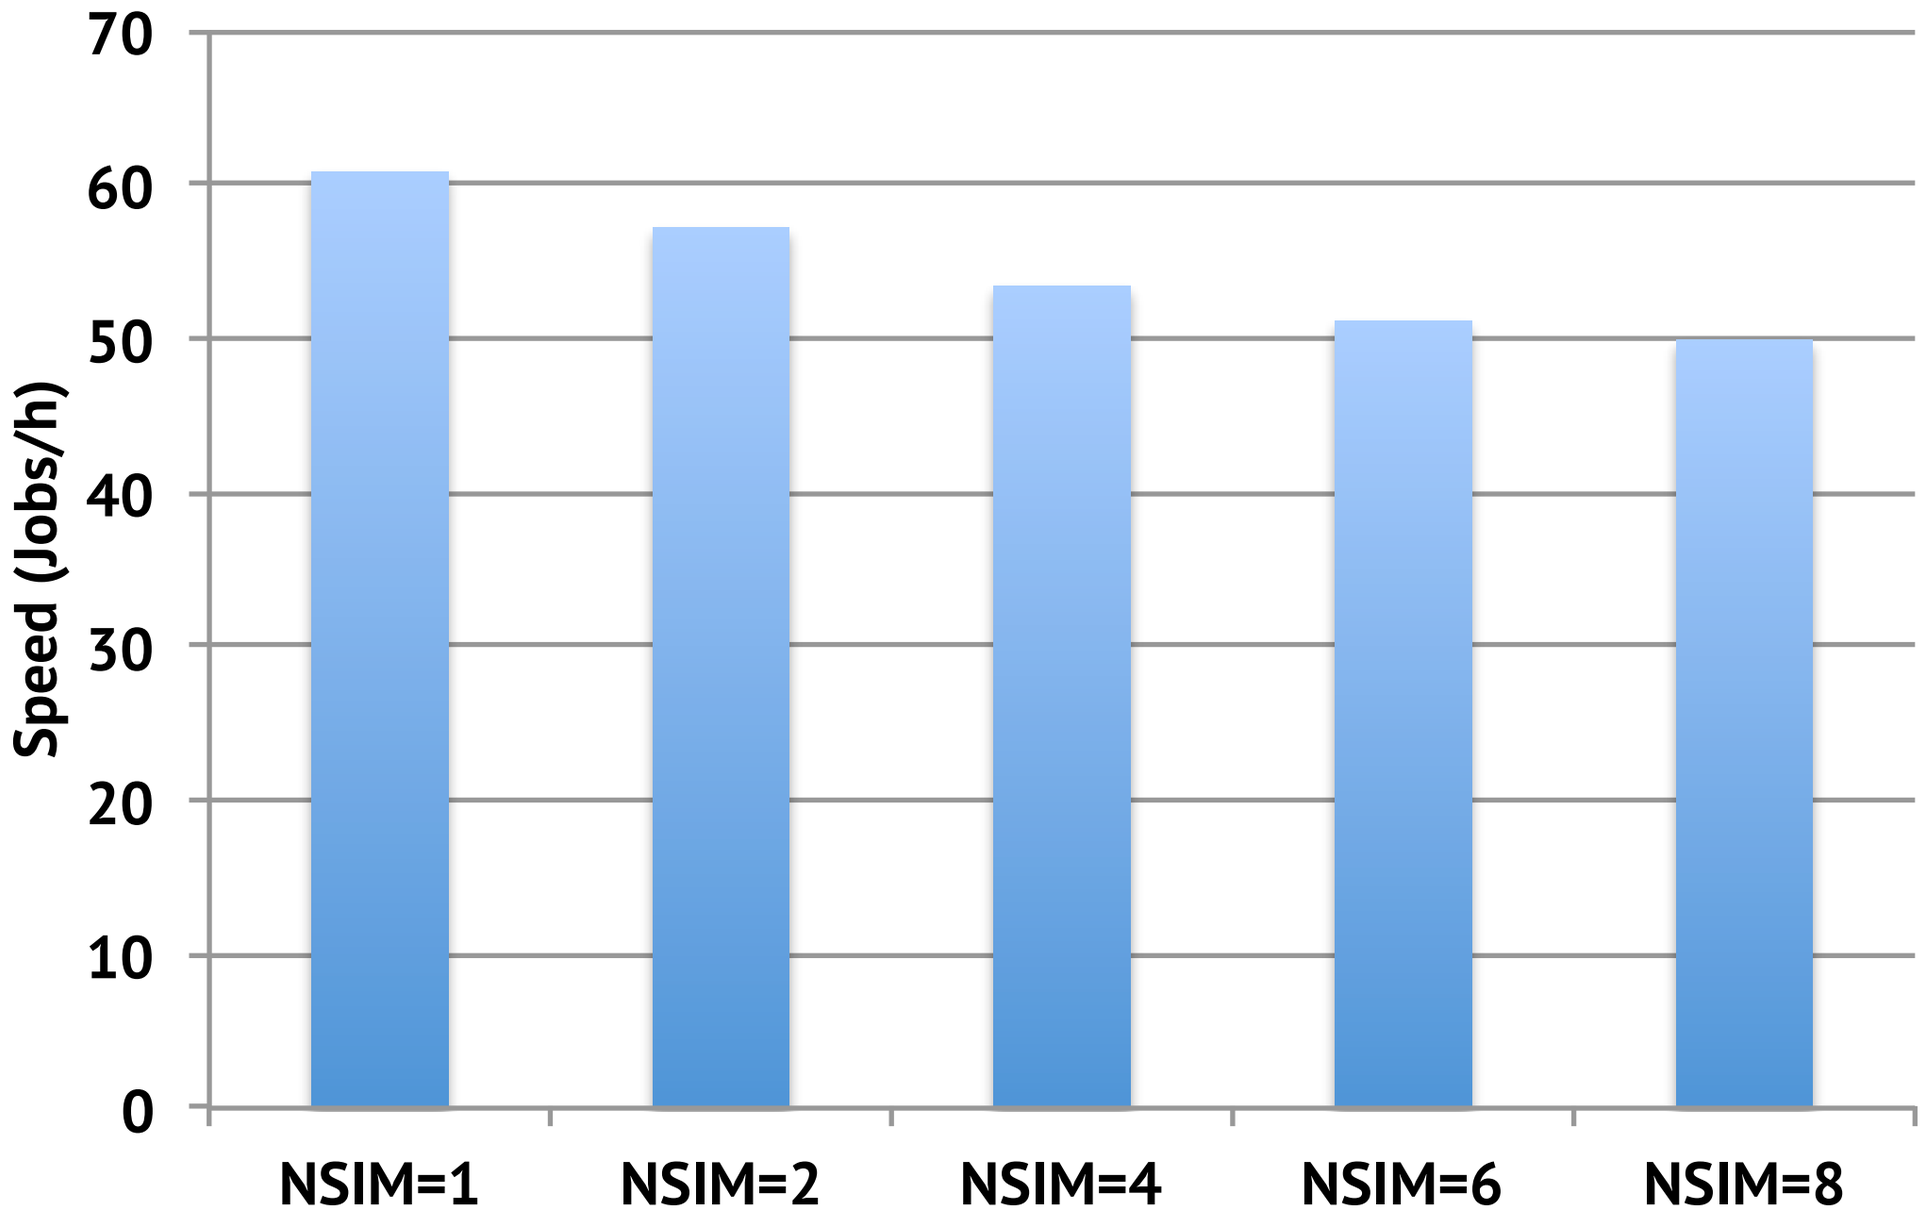

Core binding / pinning

- It is important to lock an MPI process to a specified processor core. Otherwise they can jump around…
- Intel MPI binds by default.
- Older versions of OpenMPI does not. Try e.g.:

mpirun -np 32 --display-map —bind-to-core …

- Core binding on AMD Opteron 6300-series “Interlagos” processors is not intuitive. [Cf. one of my blog posts]