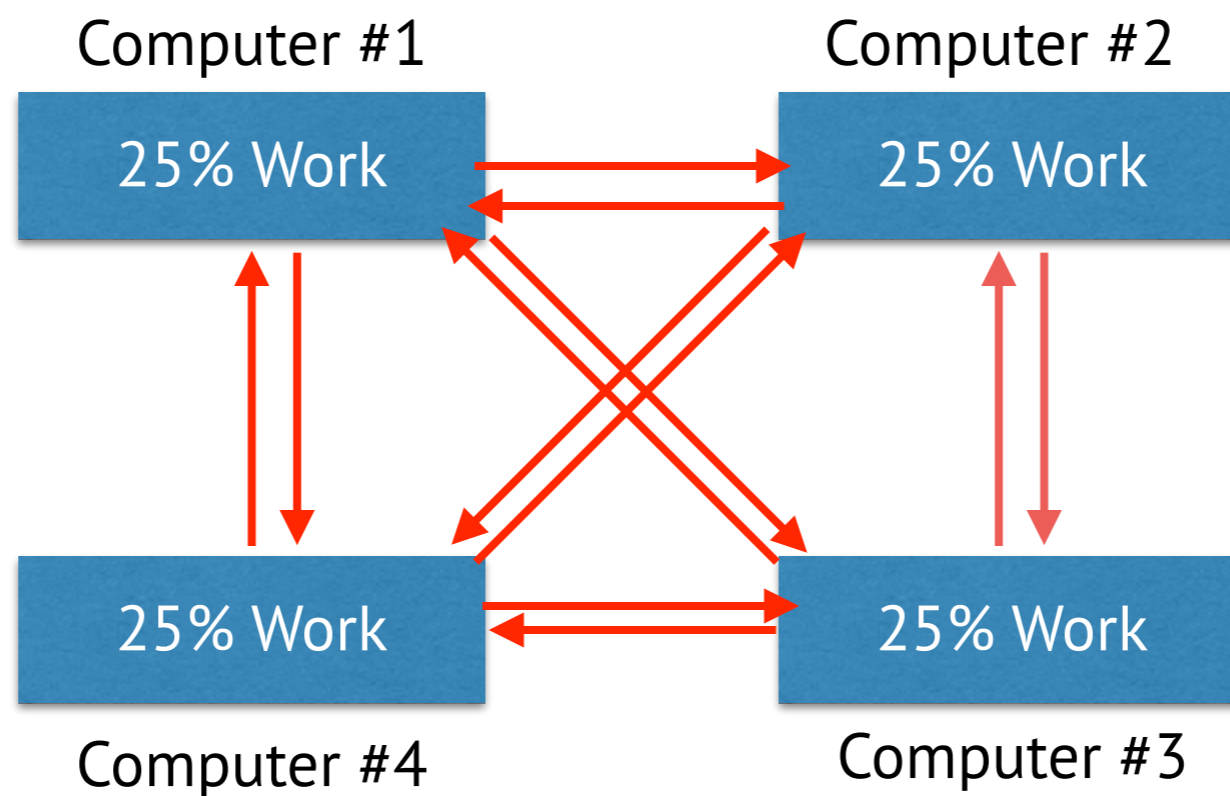




Session1:

**General performance aspects
of electronic structure software**

What happens in a parallel program?

Work (100%)

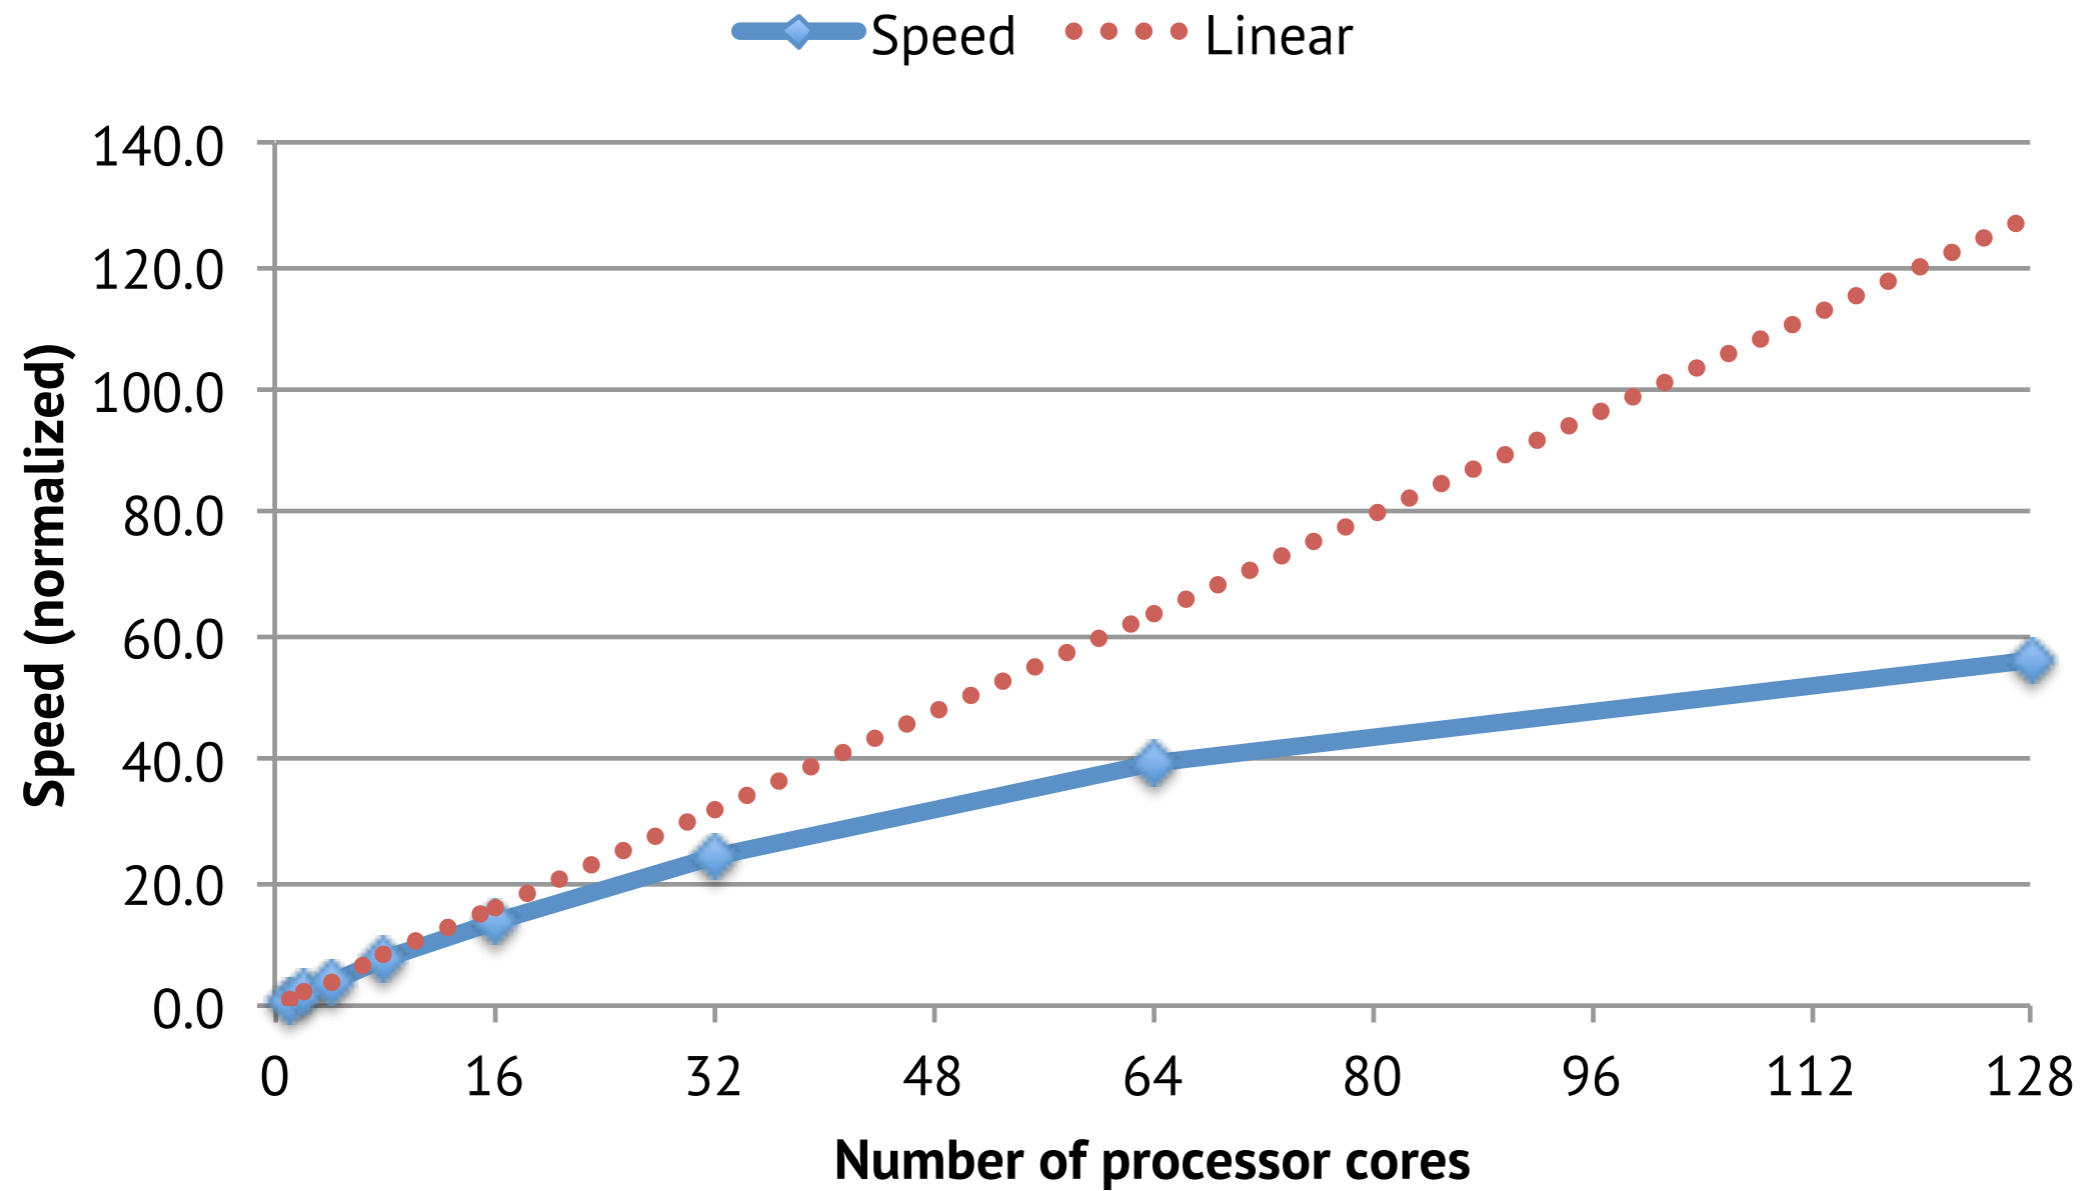




Running in parallel is


Parallel speed-up according to Amdahl's law

(Program is 99% parallel!)



Wasted computer time vs parallel job size

(Program is 99% parallel)



Dear Professor,

We regret to inform you that we have decided to cut your time allocation from 1,000,000 core hours/month to 500,000 core hours/month.

Best regards,
SNAC

Habit of PhD student “A”:

#SBATCH -N 4

#SBATCH -t 04:00:00

Habit of PhD student “B”:

#SBATCH -N 32

#SBATCH -t 01:00:00

“Mechanical Sympathy”

– Hardware and software working together in harmony

Magic numbers

Use 24 or 25 compute nodes?

511 vs 512 bands (or grid points)?

You have 4 compute nodes: 7 or 8 k-points in Brillouin zone?

Magic numbers

Use 24 or 25 compute nodes?

24 servers are connected to 1 Infiniband network switch

511 vs 512 bands (or grid points)?

$$511 = 7 \times 73$$

$$\text{whereas } 512 = 2^9$$

You have 4 compute nodes: 4 or 5 k-points in Brillouin zone?

4 (even parallelization over k-points)

Memory issues

The #1 cause of a crashing DFT calculations is
running out of memory
(Especially for DFT-hybrid or GW)

A simple **non-linear** example:

Increase $2 \times 2 \times 2$ k-point grid to $4 \times 4 \times 4$ grid.

How much more memory is needed?

Memory issues

Example 2:

Increase the plane-wave basis set from
400 eV to 600 eV

How much more memory is needed?

Memory issues

Example 2:

Increase the plane-wave basis set from
400 eV to 600 eV

How much more memory is needed?

$$n_{pw} \propto \text{cut-off}^{\frac{3}{2}}$$




So memory usage increases ca 1.8x

Distributing memory

1 node



4 compute nodes



Memory demands can decrease (or increase!)
by the number of processors used.

Out of memory on Beskow

```
entering main loop
  N      E      dE      d eps      ncg      rms      ort
DAV:  1    0.137475550122E+05    0.13748E+05    -0.44040E+05    6144    0.754E+02
DAV:  2    0.343134701388E+04    -0.10316E+05    -0.10076E+05    6144    0.185E+02
DAV:  3    0.125016272295E+04    -0.21812E+04    -0.21264E+04    6144    0.827E+01
DAV:  4    0.515291496736E+03    -0.73487E+03    -0.72475E+03    6144    0.497E+01
DAV:  5    0.348309184886E+03    -0.16698E+03    -0.16575E+03    6144    0.225E+01
[NID 01184] 2015-01-02 11:50:12 Apid 265643: initiated application termination
[NID 01192] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01164] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01172] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01250] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01349] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01166] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01246] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01215] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01360] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01530] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01184] 2015-01-02 11:50:13 Apid 265643: OOM killer terminated this process.
[NID 01194] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01211] 2015-01-02 11:50:13 Apid 265643: OOM killer terminated this process.
[NID 01350] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
[NID 01171] 2015-01-02 11:50:13 Apid 265643: OOM killer terminated this process.
[NID 01348] 2015-01-02 11:50:13 Apid 265643: Error detected during page fault processing. Process terminated
Application 265643 exit signals: Killed
Application 265643 resources: utime ~384s, stime ~0s, Rss ~4936, inblocks ~30753408, outblocks ~71807232
```

“OOM killer terminated this process”

OOM = out of memory

OOM on Beskow: 2nd variation

```

entering main loop
      N      E      dE      d eps      ncg      rms      ort
DAV:  1      0.464516240572E+05      0.46452E+05      -0.27873E+06      6144      0.108E+03
DAV:  2      0.129044910820E+05      -0.33547E+05      -0.31515E+05      6144      0.322E+02
DAV:  3      0.352220084587E+04      -0.93823E+04      -0.87191E+04      6144      0.178E+02
DAV:  4      0.115885589713E+04      -0.23633E+04      -0.22517E+04      6144      0.988E+01
DAV:  5      0.526902167817E+03      -0.63195E+03      -0.61444E+03      6144      0.552E+01

      N      E      dE      rms(c)
ROT:  1      -0.229870123359E+04      -0.22987E+04      0.106E+02
ROT:  2      -0.226649298470E+04      0.32208E+02      0.641E+01
ROT:  3      -0.224391290515E+04      0.22580E+02      0.166E+00
gam= 0.000 g(H,U,f)= 0.551E+03 0.263E+02 0.799E-20 ort(H,U,f) = 0.000E+00 0.000E+00 0.000E+00
SDA:  6      -0.224374120608E+04      -0.27706E+04      -0.23104E+03      6144      0.578E+03 0.000E+00
forrtl: severe (41): insufficient virtual memory
Image              PC              Routine              Line              Source
vasp-build04-half  000000000AAE1C5  Unknown              Unknown              Unknown
vasp-build04-half  000000000F2FE8D  Unknown              Unknown              Unknown
vasp-build04-half  000000000F8666E  Unknown              Unknown              Unknown
vasp-build04-half  000000000425FA9  Unknown              Unknown              Unknown
vasp-build04-half  000000000401056  Unknown              Unknown              Unknown
vasp-build04-half  0000000005014191  Unknown              Unknown              Unknown
vasp-build04-half  000000000400F1D  Unknown              Unknown              Unknown
forrtl: severe (41): insufficient virtual memory
Image              PC
vasp-build04-half  000000000AAE

```

“forrtl” = Fortran Runtime Library
 “insufficient virtual memory” =
 Couldn’t allocate enough memory


Bottlenecks

1. Total memory bandwidth
2. Network latency
3. Network bandwidth
4. Number of cores
5. Clock frequency

What is faster?

**16 compute nodes using 1
processor (2.2 Ghz)
connected with Infiniband**

**1 node using 16 processors
(2.6 Ghz)**



VASP scaling: Shared-memory vs Infiniband

