

Specialist course on efficient use of VASP

(and other electronic structure codes
on modern supercomputers)

Peter Larsson, Ph. D
National Supercomputer Centre
Linköping University, Sweden

More info at the web page:
<http://www.nsc.liu.se/~pla>

Electronic structure software packages


General performance aspects

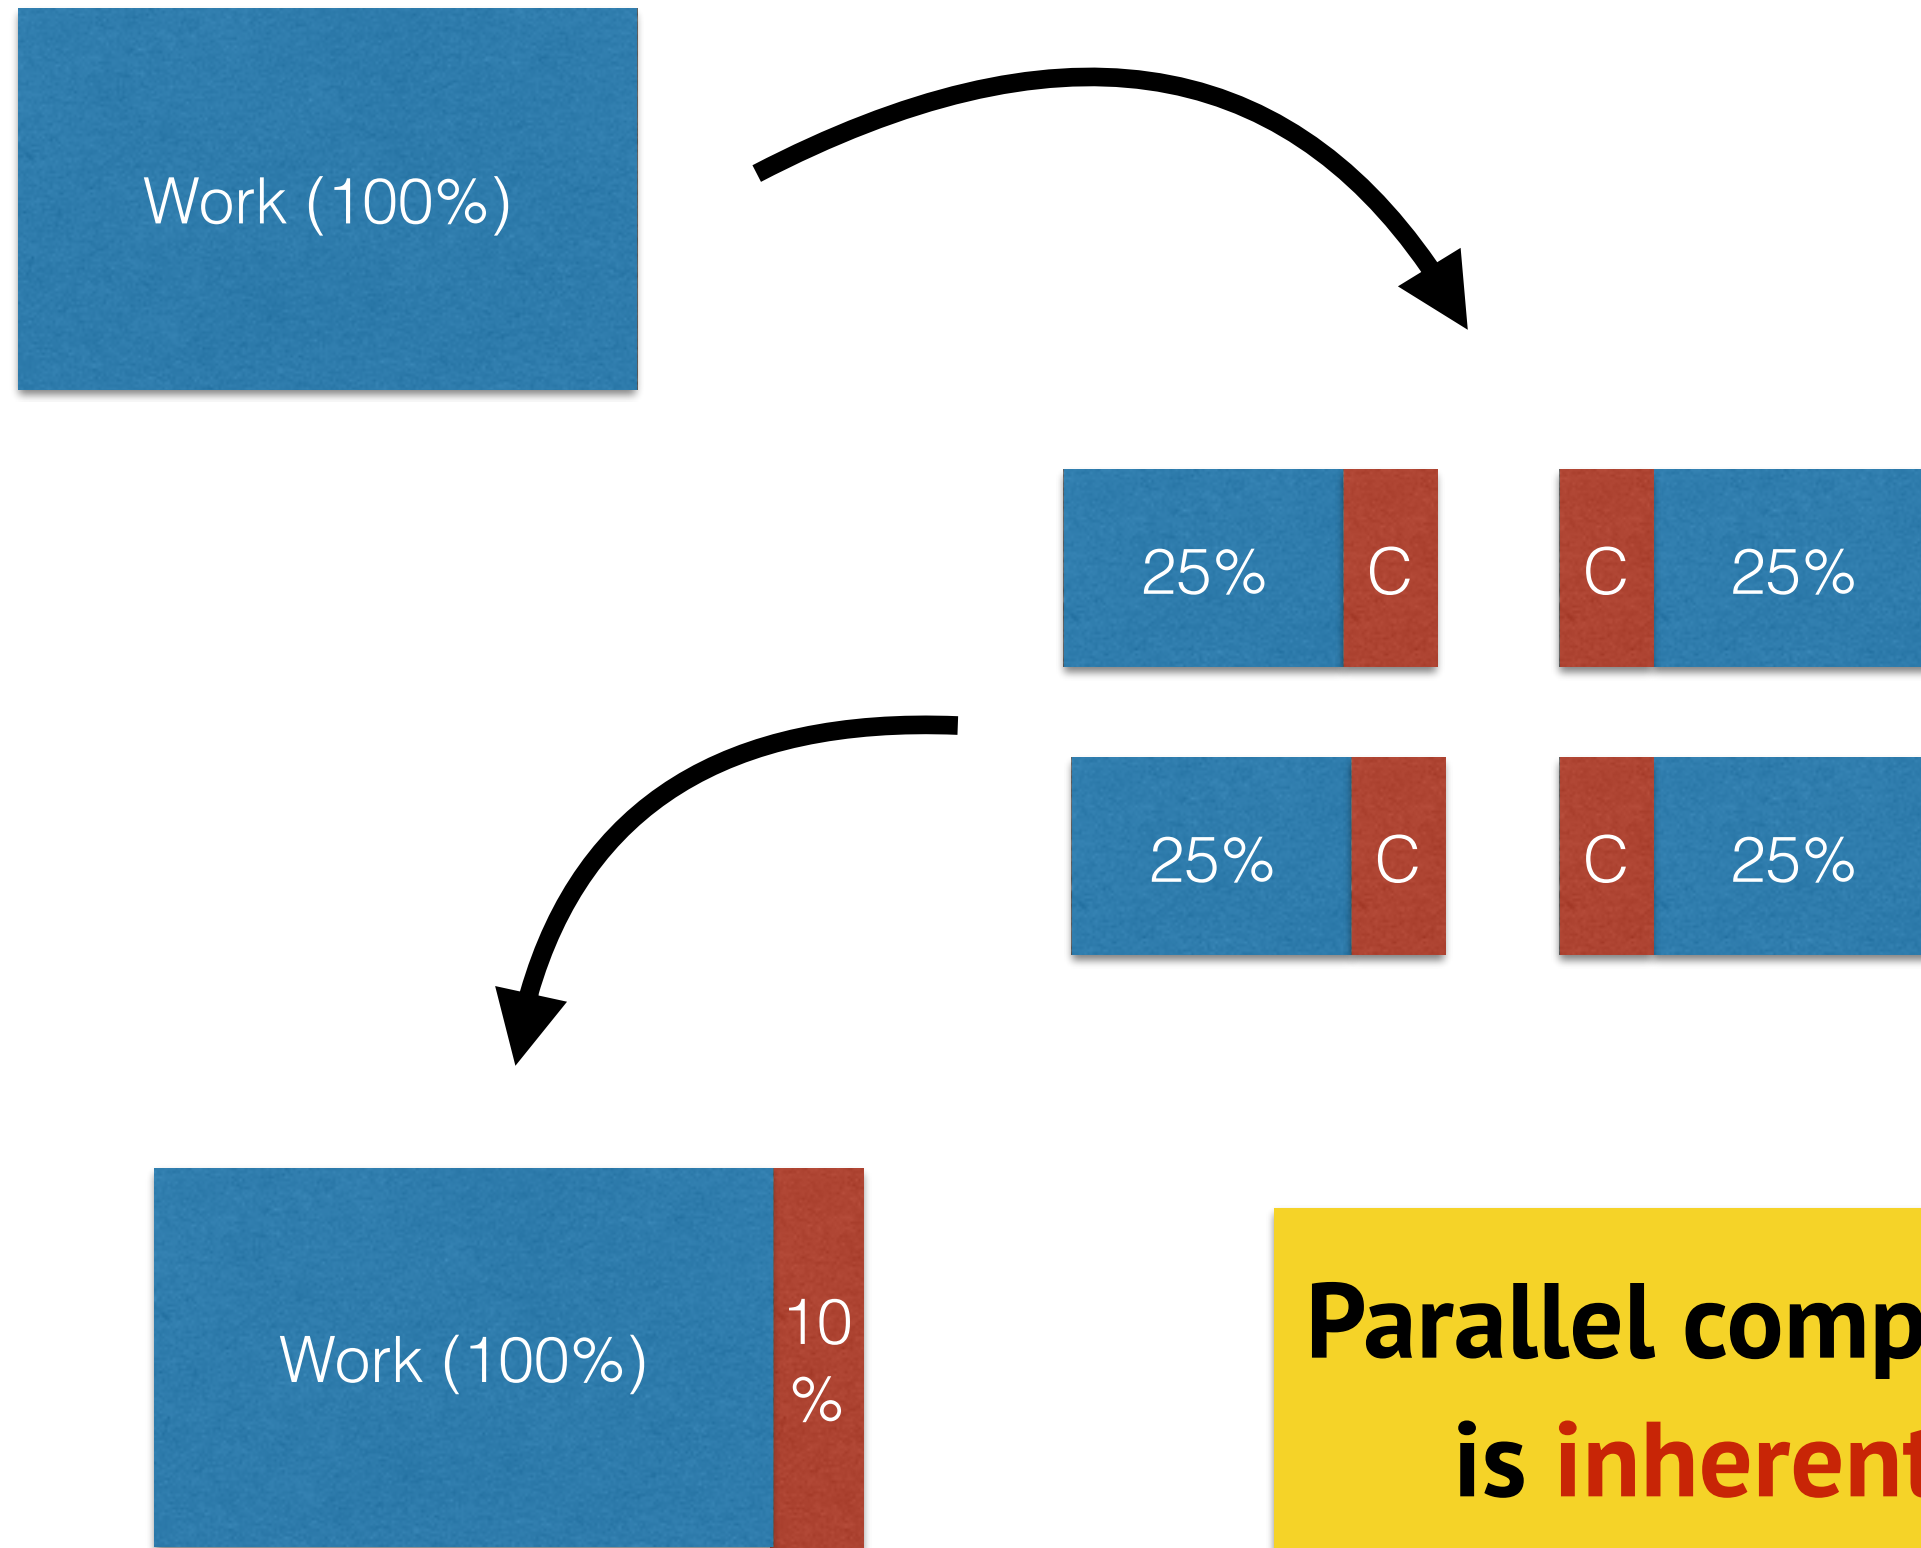


- <http://www.psi-k.org/codes.shtml>

What happens in a parallel program?

Work (100%)






**Parallel computing
is inherently
wasteful!**

$$Speed(n) = \frac{1}{S + \frac{1}{n}(1 - S)}$$


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 N.N.,

*We regret to inform you that we have
decided to cut your time allocation
from 1000 node days to 500 node days.*

*Best regards,
Computer Center Staff*

“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: 4 or 5 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 (with parallelization over k-points)

Post lecture remark regarding power-of-two numbers such as 512. These numbers are nice for load-balancing at a high hardware level (nodes, NUMA zones, processors). At a low level, inside a subroutine, odd numbers such as 513 can actually be better to use because it potentially minimizes memory cache misses.

Memory issues

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

A simple non-linear example:
Increase 2x2x2 k-point grid to 4x4x4 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.*

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

