



High Performance Computing

# ScaleMP at HPC UGent

customer presentation @ Supercomputing 2012

Nov. 13th 2012

Kenneth Hoste

*kenneth.hoste@ugent.be*





High Performance Computing  
*Empowering researchers*



# HPC UGent

## *Who we are:*

- central contact for high-performance computing at UGent (Ghent University, Belgium)
- part of the central ICT department (DICT)
- member of the VSC (Flemish Supercomputer Centre)

## *What we do:*

- system administration of the HPC infrastructure (STEVIN)
- support for researchers using our infrastructure
- from easy to hard problems, and also giving lectures



High Performance Computing  
*Empowering researchers*



# Our users

Over 600 user accounts

Over 250 different software packages installed

Researchers from various and very diverse fields:

- computer science & engineering
- genomics
- quantum chemistry
- medicine
- ...

## Brief overview of our infrastructure STEVIN:

### *two MPI clusters*

- i) 194x8 cores Intel Harpertown, 16GB/node, IB DDR
- ii) 34x32 cores AMD Magny Cours, 64GB/node, IB QDR

### *three batch clusters*

- iii) 56x8 cores Intel Nehalem, 12GB/node, Gb Ethernet
- iv) 167x8 cores Intel Nehalem, 12GB/node, Gb Ethernet
- v) 64x16 cores Intel Sandy Bridge, 32GB/node, Gb Ethernet

*and a ScaleMP setup...*



High Performance Computing  
Empowering researchers



# ScaleMP @ STEVIN

## *Our ScaleMP setup:*

16x Dell R710, Mellanox IB IS5030, Cisco 1Gb switches

16x2 hexacore Intel Westmere (192 cores), 96GB each

2 single port IB QDR + 1 IB 36-port QDR switch

vSMP (version 4.0.225.53) with load balancing support

partitioning setup (4x4); 4x 48-core systems with ~350GB each

or rather 3x 48-core + 2x 24-core systems right now

(limited to partition to max. 1TB of memory per system due to license)

OS: Scientific Linux 6.2, vSMP 2.6.32 kernel

(stock kernel doesn't work due to SAS controller driver and vSMP)

local disks: 2x 300GB 15k RPM in HW raid0,

vSMP software raid0 over the disks





High Performance Computing  
Empowering researchers



# ScaleMP @ STEVIN

## Tier I system

### *VSC Tier1 setup:*



first VSC supercomputer  
hosted at Ghent University

528 HP SL230 G8 servers, Intel Sandy Bridge 2x8 cores  
64GB memory/node (4GB/core), 8448 cores in total

Infiniband FDR 4x (56 Gbit/sec full duplex)

152.4 TFlops (#118 in Top500 list June 2012)

400TB dedicated scratch storage

We will experiment with ScaleMP on this system as well in the future.

This is possible due to *vSMP floating licenses* and the sufficiently similar (ScaleMP certified) hardware as our current vSMP setup.

Some of our users have a need for:

- (a lot) more memory per node
- more cores per node (non-MPI software)
- both at the same time

We also wanted:

- something flexible (in various aspects) ✓
- good price/performance ratio (i.e. affordable) ✓
- something transparent (for users) ✓

Initial interest in ScaleMP product just before SC10.

Booth visit convinced us to look into it further.

After proof-of-concept on our Harpertown MPI cluster, we decided to purchase dedicated hardware + a vSMP license.

Some initial problems, but *great support* from the ScaleMP team resulting in adequate and timely solutions.

Setup has been running great now for 1.5 years.



High Performance Computing  
Empowering researchers



# Use cases

## *Gaussian: quantum chemistry*

good thread parallelization, poor multi-process parallelization

built with Intel compilers, custom link wrappers for pinning

**excellent scaling**, close to linear (tested up to 48 cores)  
for runs that are dominated by parallel links (e.g. I502, I1002)

serial parts can benefit from more memory and/or better IO

relevant part of .com input file:

```
%Nproc=48
```

```
%MEM=250GB
```

```
#p b3lyp/gen pseudo=read opt(ts,noeigentest,calcfc)  
freq(noraman) scf=tight int(grid=ultrafine) ginput  
iop(2/16=1,6/7=3)
```



High Performance Computing  
*Empowering researchers*



# Use cases

*VASP: ab-initio quantum-mechanical molecular dynamics*

high memory requirements

limited scalability

source patched and built with Intel compilers

vSMP cluster allows for efficiently handling  
experiments that do not scale well using (only) MPI  
or require huge amounts of memory



High Performance Computing  
*Empowering researchers*



# Use cases

## *Trinity: next-generation genome sequencing*

only multi-threading support, no MPI

high memory requirements (for efficient performance)

significant shorter runtime on vSMP cluster:

*inchworm*: from 6.5 days (6 cores) to 2.5 hours

*chrysalis*: from 5 to 3.5 days (~400GB memory req.)

*butterfly*: from about 1 day (master-slave) to +/- 2 hours



High Performance Computing  
*Empowering researchers*



To conclude...

## *Feedback*

users are very happy with our vSMP-based cluster

good performance / scalability

makes experiments possible that would be infeasible otherwise

## *Future plans with vSMP*

experiment on Tier1 system powered by Intel Sandy Bridge

extend use cases to commonly used Tier2 software

R (multi-threaded), MATLAB, Velvet, ...

suggested by ScaleMP team as additional potential applications



High Performance Computing

# ScaleMP at HPC UGent

customer presentation @ Supercomputing 2012

Nov. 13th 2012

Kenneth Hoste

*kenneth.hoste@ugent.be*

