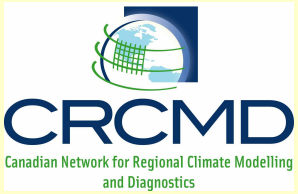

Machine structures

Katja Winger, UQAM
UQAM, 28. Jan. 2010

Marvin

Linux cluster

50 nodes with two dual-core processors each

1 headnode
48 compute nodes => **192 cores**
1 post processing node

Used to run CRCM5

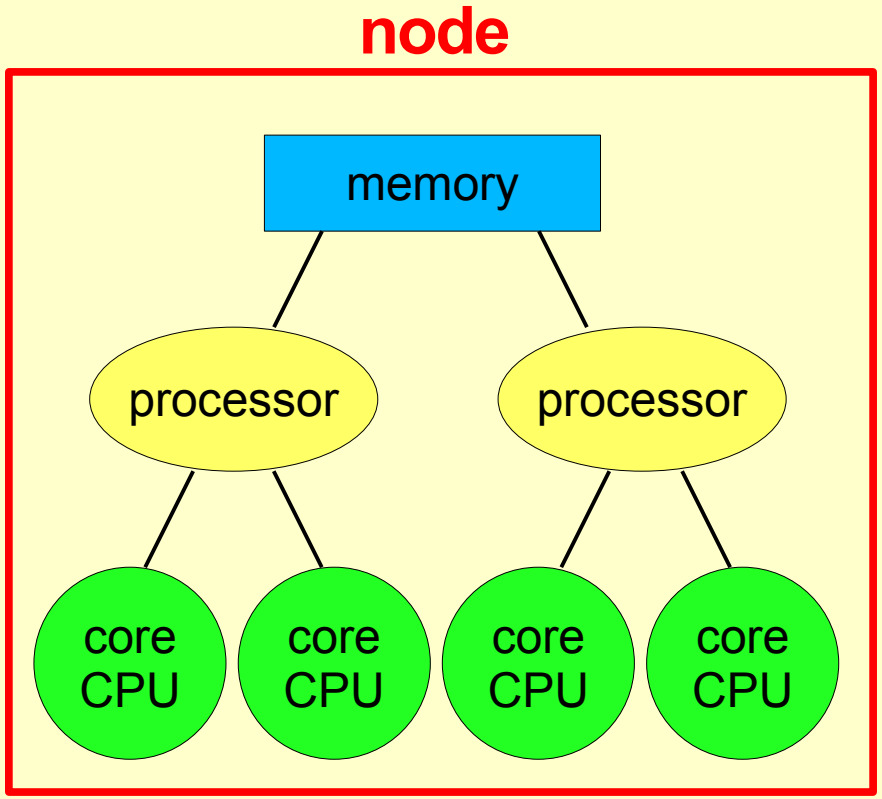

sti (st1/st2/st3)

Linux cluster
3 nodes with two quad-core processors each

Used for storage and post processing

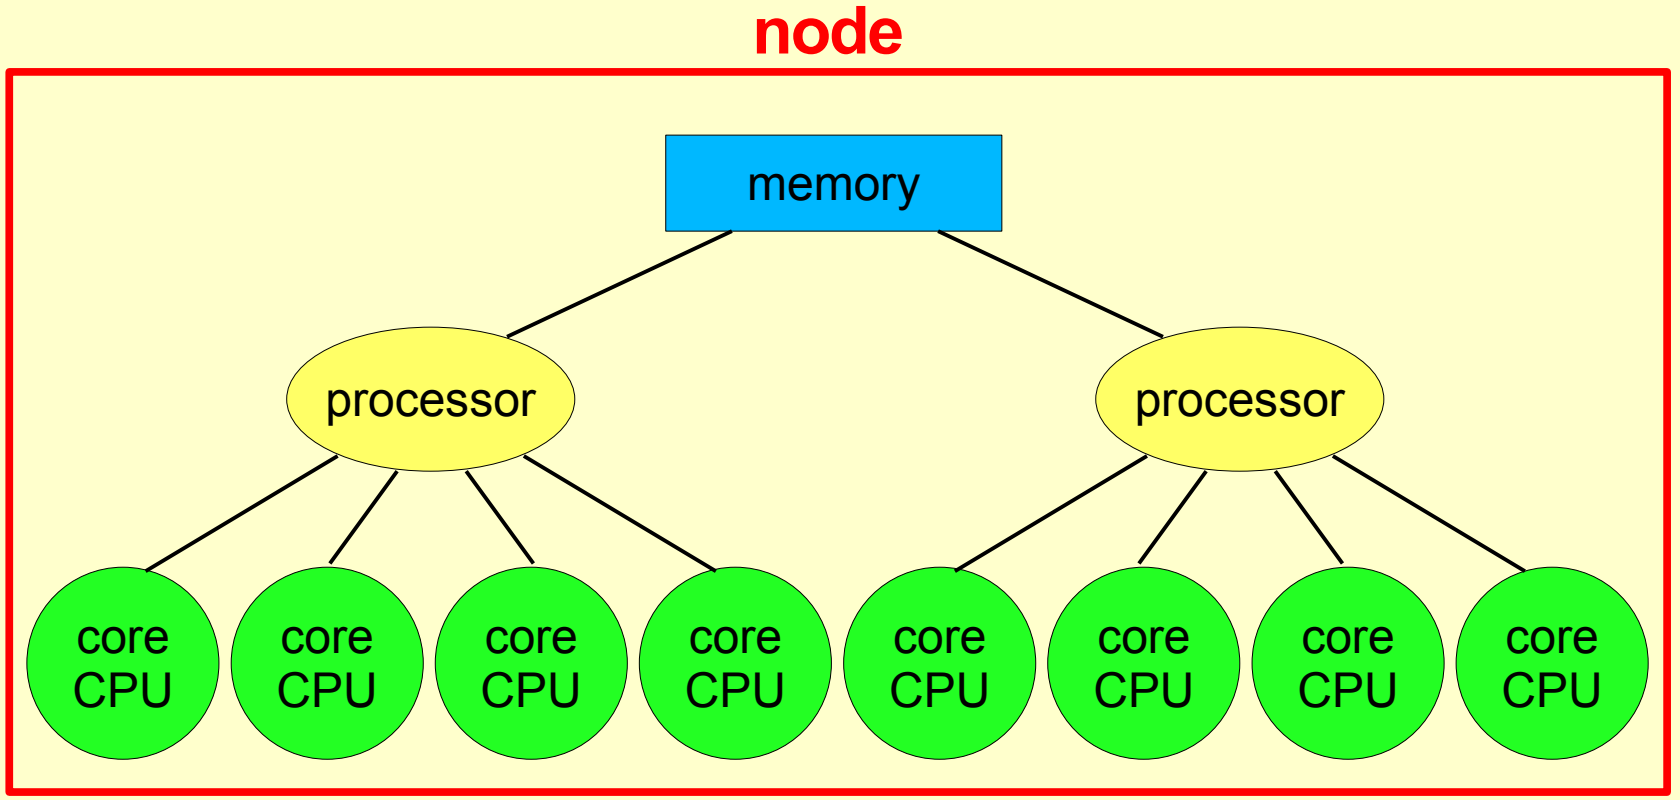

3 compute nodes => **24 cores**

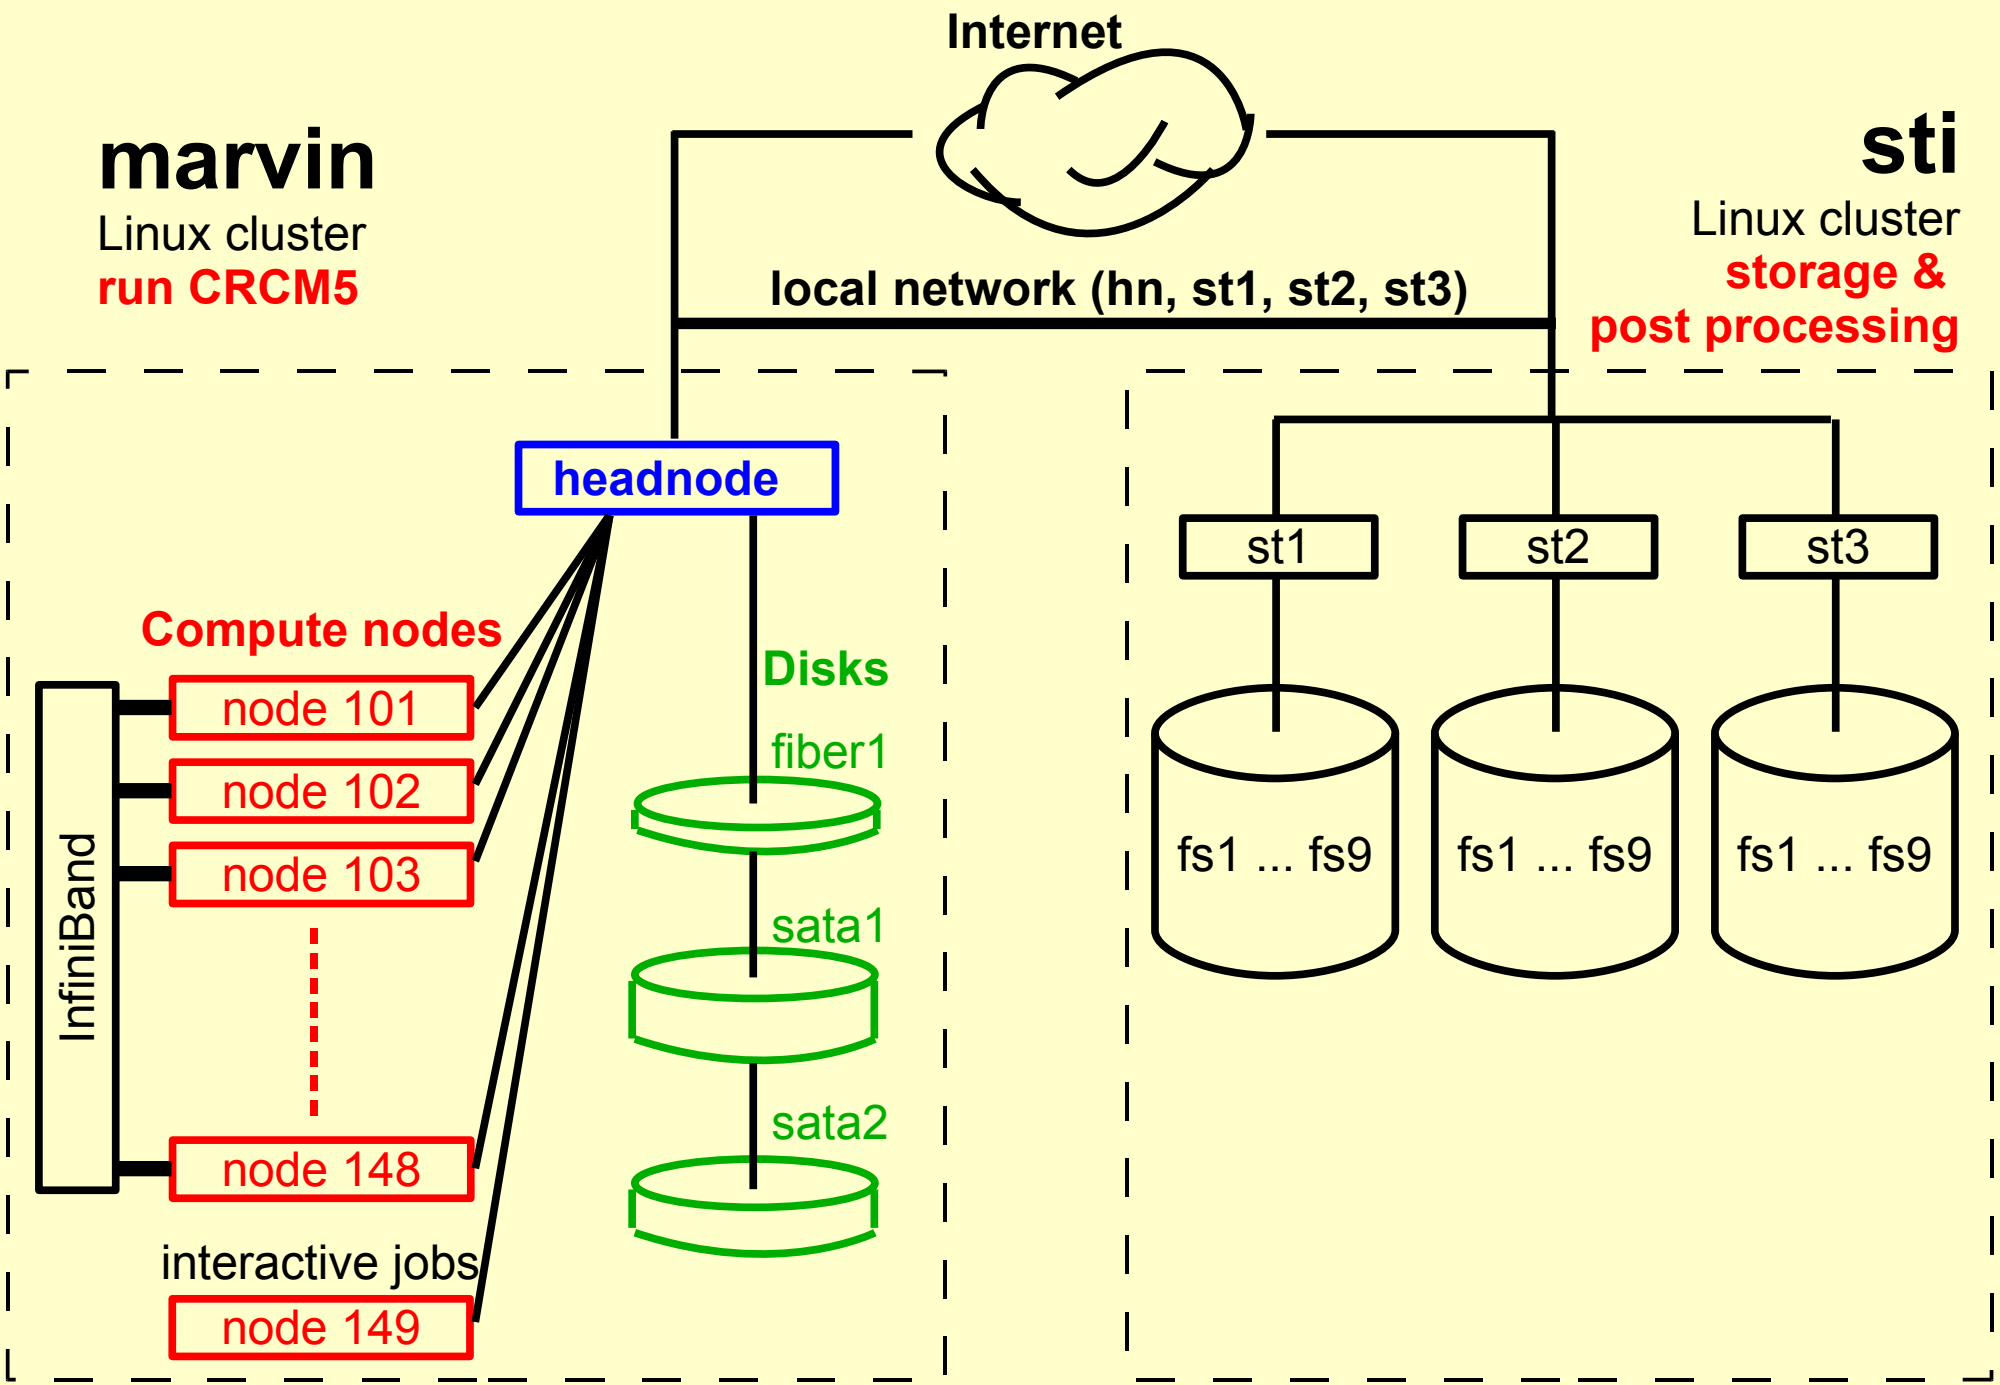
Internet
marvin
Linux cluster
run CRCM5
local network (hn, st1, st2, st3)
sti
Linux cluster
storage &
post processing
headnode
Compute nodes
node 101
node 102
node 103
node 148
interactive jobs
node 149
InfiniBand
Disks
fiber1
sata1
sata2
st1
st2
st3
fs1 ... fs9
fs1 ... fs9
fs1 ... fs9

Quotas on marvin

You can always check your quota with:

quota -v

Quotas:

HOME:	200 MB	
fiber1:	50 GB	(no storage, this is where the model is running)
sata2:	100 GB	(this is where the post processing is running and where you can store you executable, analysis files, geophysical fields, ...

MPI

"**M**essage **P**assing **I**nterface"

distributed memory

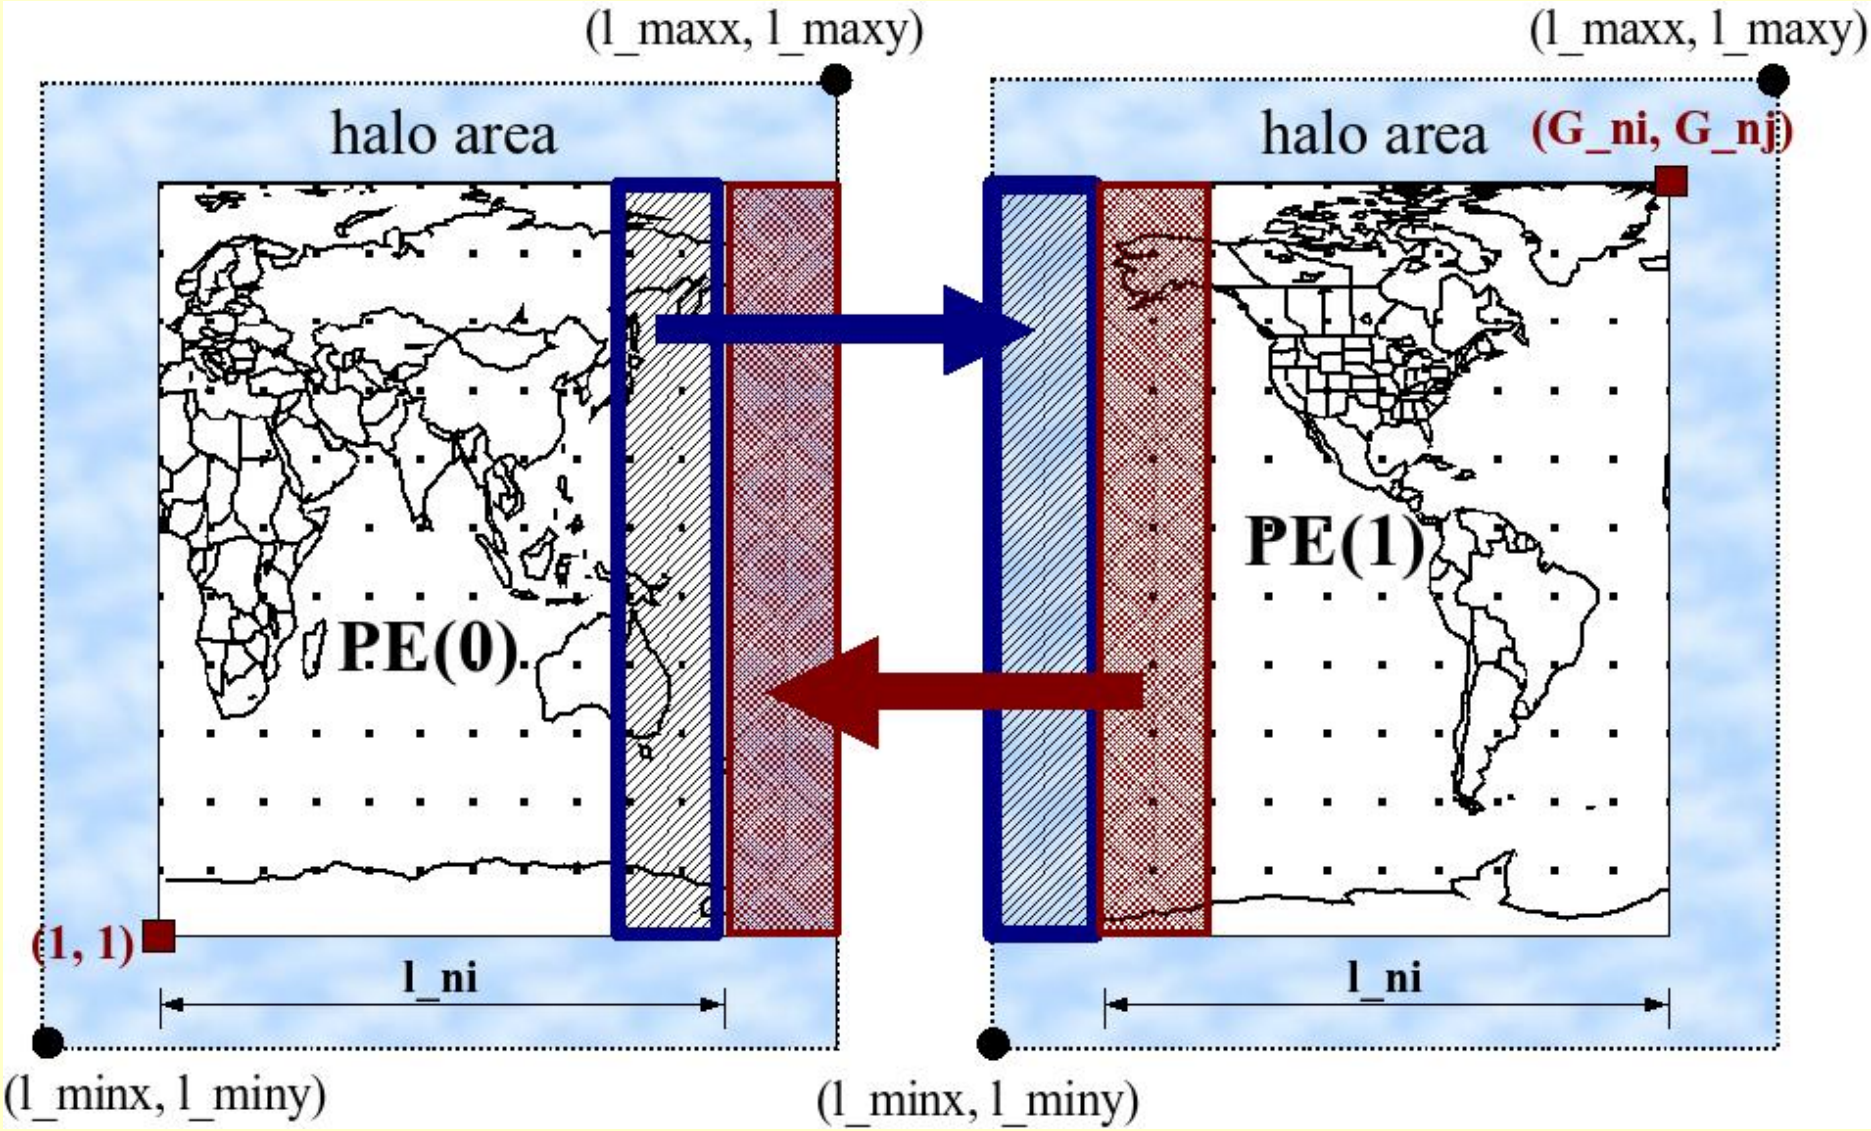

Courtesy of Vivian Lee

MPI

"**M**essage **P**assing **I**nterface"
distributed memory

To run the model in MPI you partition the domain into tiles, xx-yy.
Each cpu will execute the model on one tile.

LAM grid partitioned into 4x4 tiles:

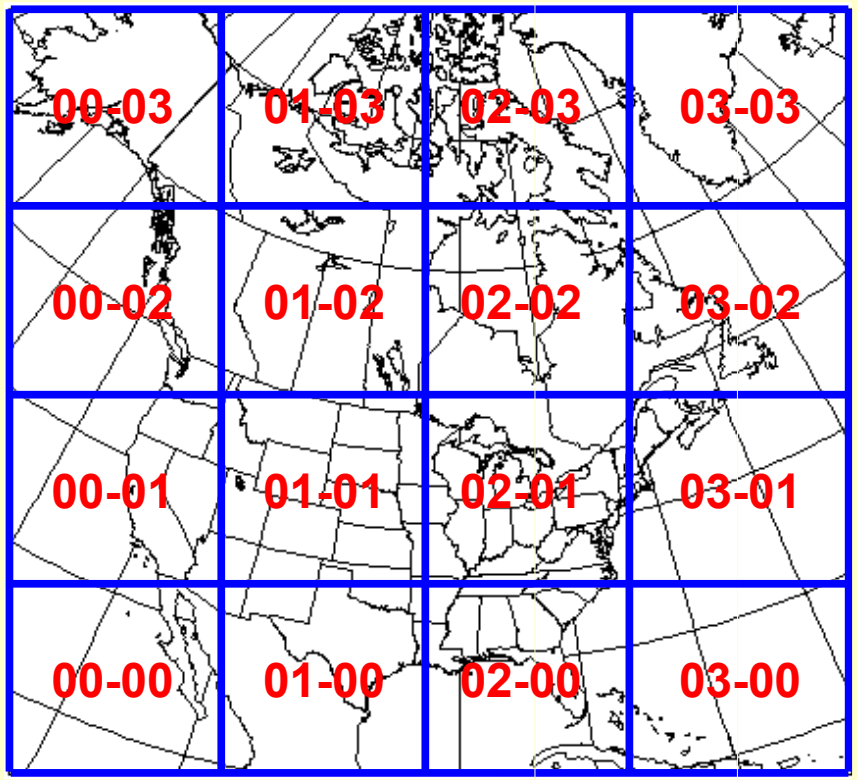

Distribution of tiles over cores

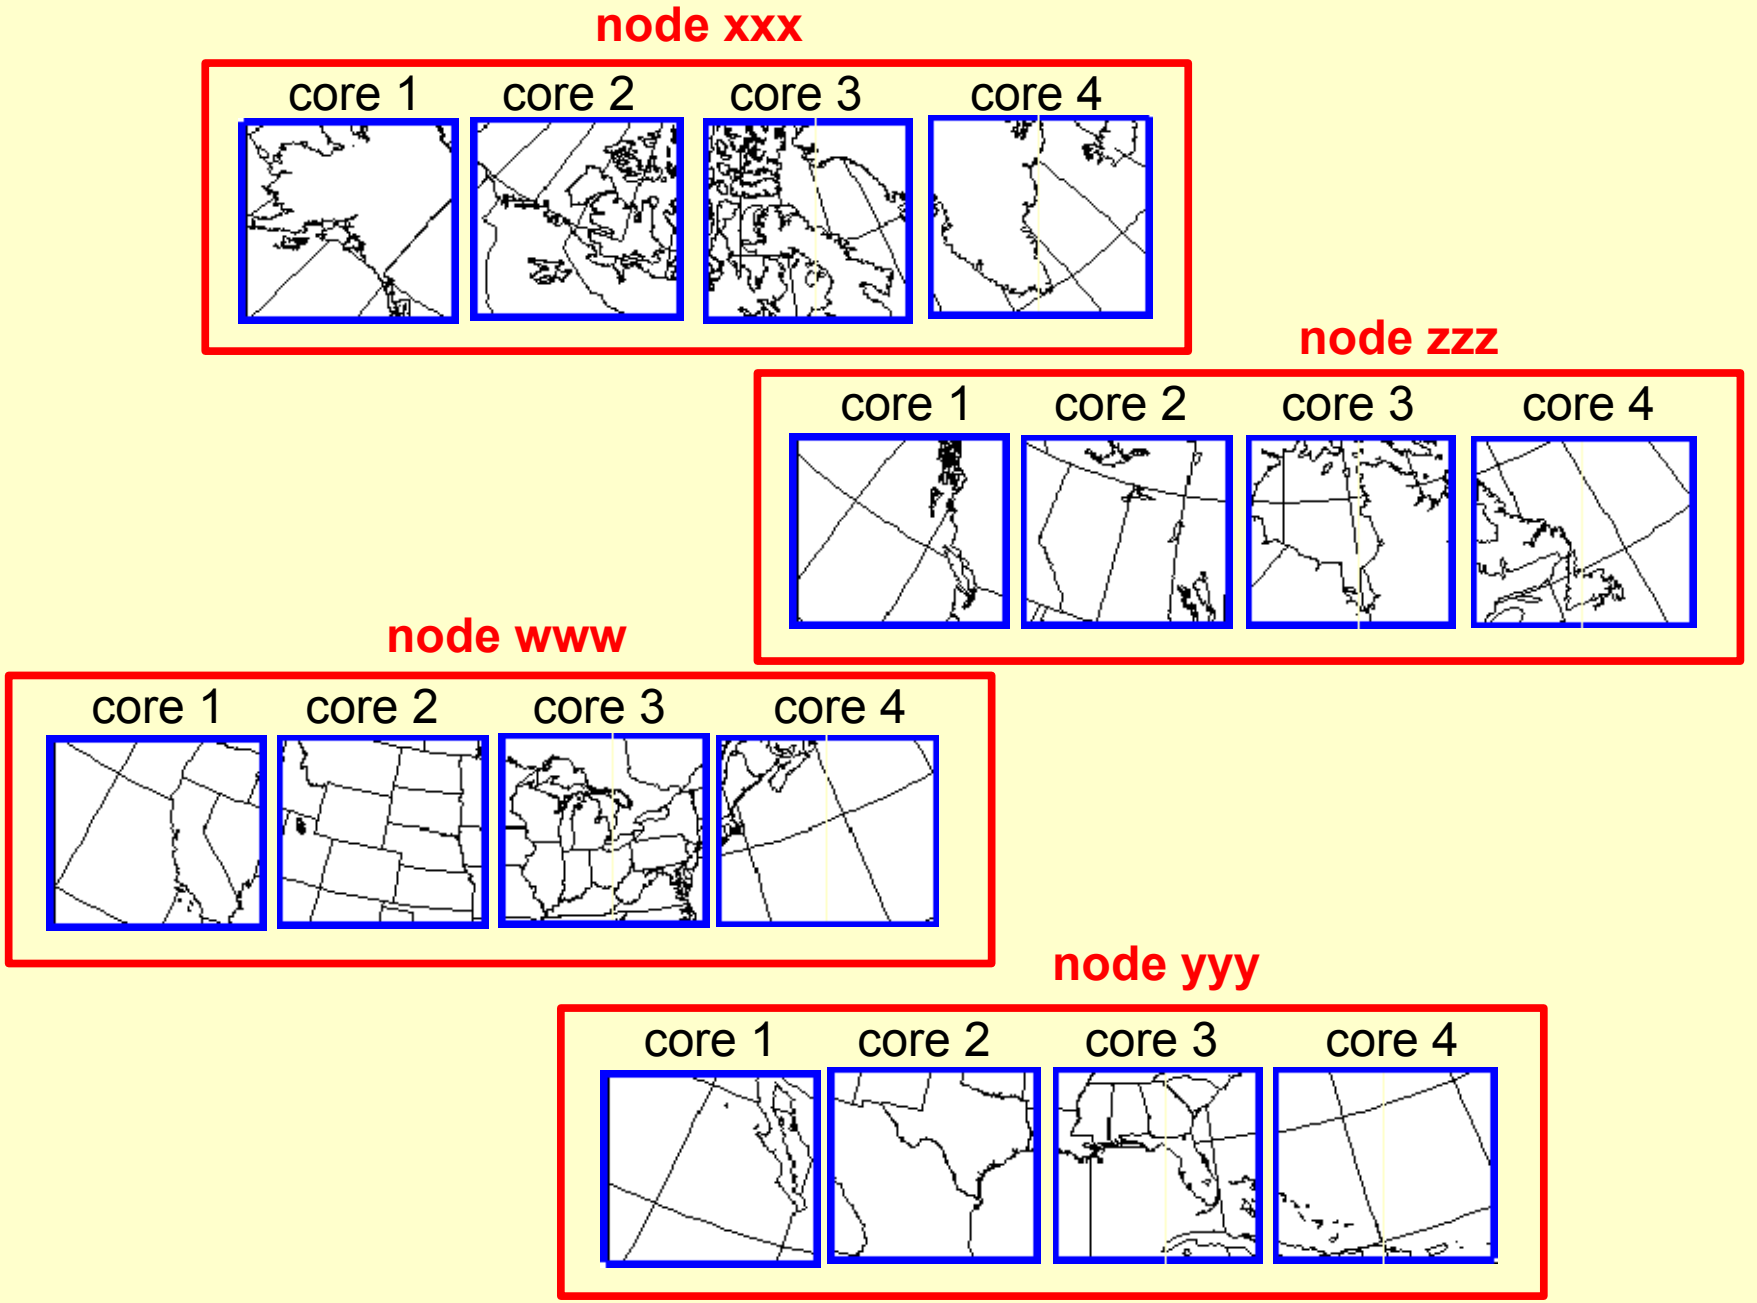

OpenMP

using shared memory

OpenMP in the physics:

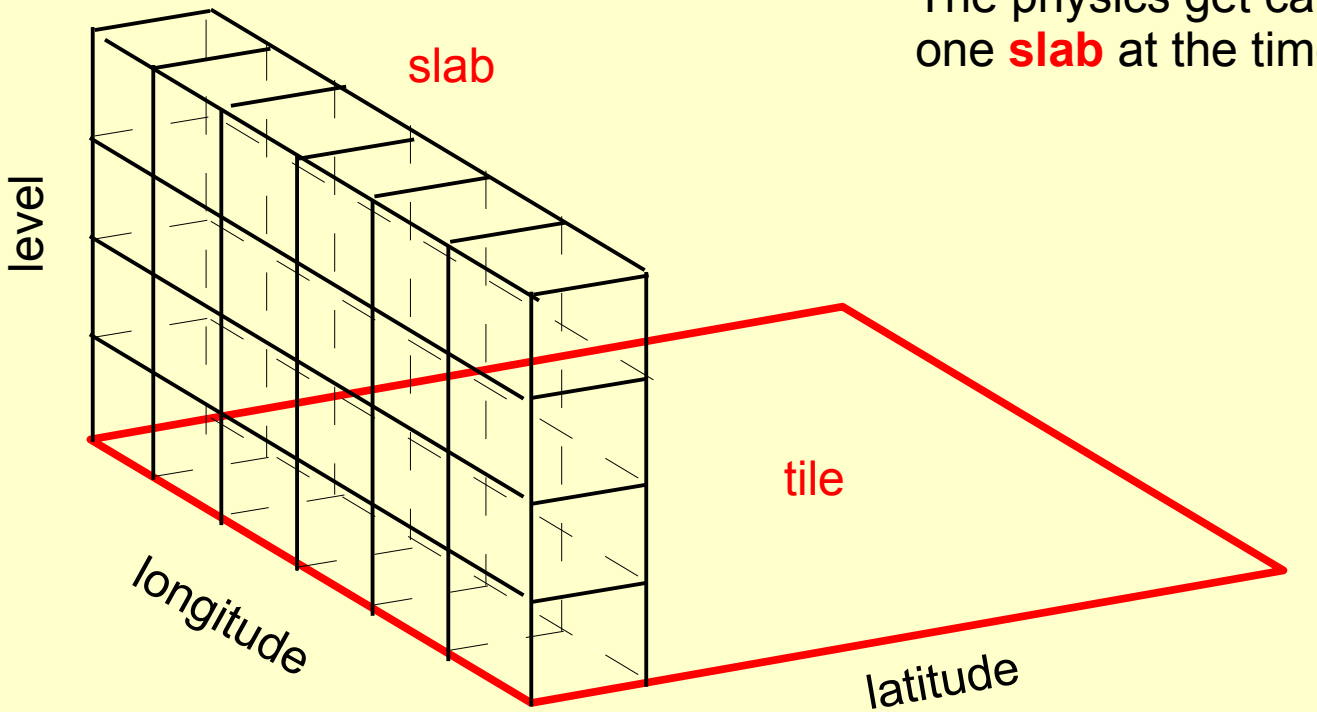

The physics get called one **slab** at the time

Batch jobs

Submit a job via Sun GridEngine

On marvin always use 'soumet'!!!
soumet *jobname* **-t** *time* **-listing** *listing-directory* **-jn** *listing_name* [**-cpus** #*cpus*]

Check on jobs

qstat
~winger/ovbin/qs

Kill a job

qdel *job-ID*

Interactive jobs on marvin

Interactive jobs should only be submitted from node cn149, never from the headnode!
You can logon to the node with:

ssh -X cn149