



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



**EXCELENCIA
SEVERO
OCHOA**

Assignments

Prof. Jesús Labarta
BSC & UPC

ACM Summer school
Barcelona, July 22nd 2019

Agenda

- **Introduction:** before start working with the system
 - Software requirements, login information, configuration, system overview, etc.
- **Paraver tutorial:** getting familiar with paraver
 - A Paraver tutorial to introduce the fundamentals of this tool
- **The matrix multiply example:** playing with this code
 - The matrix multiplication code has appeared during the seminar's slides. A good opportunity to try it!
- **Tutorial exercises:** applying learned lessons (guided examples)
 - Heat and N-Body kernels using TAMPI
 - LULESH using DLB
- **The IFSKer Practice:** let's do it!
 - Submit your report about your experiences applying TAMPI or DLB

Documents & slides (available on-line)

<https://pm.bsc.es/ftp/training/acm-ess-2019>

Table Of Contents

[Introduction](#)

[01-Paraver](#)

[02-Matrix multiplication](#)

[03-TAMPI](#)

[04-DLB](#)

[05-IFSKER](#)

[Software Liability](#)


File Downloads

[HybridProgramming.pdf](#)

[Assignments.pdf](#)

Quick search

Enter search terms or a module, class or function name.




Software requirements

- SSH: Secure SHell (to connect the HPC system)
 - Linux: has native support of secure shell “ssh user@host”
 - Windows: need to install a ssh program
 - PuTTY <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>
 - MobaXterm <http://mobaxterm.mobatek.net/download.html>
- X Server / X forwarding (for wxparaver or .pdf readers)
 - Linux: has native support (remember to connect with **“ssh -X user@host”**)
 - Windows: need to install a X server program
 - Xming <https://sourceforge.net/projects/xming/>
 - MobaXterm already includes a X server within the package

Login information

- Connecting to Nord III (login)
 - nord1.bsc.es
 - nord2.bsc.es
- User provided in your **badge**



Practical information for the ACM Summer School

XXXXX XXXX XX XXXXX XXXX X XXXXXXXXX XX XXXXX XXXX X XXXX XXXXX XXXX XXXXXXXX
XXXX XXXXXXX X XX XXXXXXX XXXX XX XXXXX XXXXXXXXX XX XXXXX XXXX X XXXX

XXXXX XXXX XXXXXXX X XX XXXXXXX X XXXX XXXXX XXXX XXXXXXX X XX XXXXXXX XXXX
X XXXX XXXXX X XXXX XXXXX XXXX XXXXXXX X XX XXXXXXX XXXX XXXXXXX XXXXXXXXXX
X XX XXXXXXX

XX XXXXX XXXX XX XXXXX XXXX X XXXX XXXXX XXXX XXXXXXX X XX XXXXXXX XXXX
X XXXX XXXXX XXXX XXXXXXX X XX XXXXXXX

WIFI in Sala Àgora

XXXXXX
XXXXXX XXXXXXXX


Access to FTP Server for lecture material

XXXXX XXXX XXXXXXX X XX XXXXXXX XXXX
Xxx XXXXXXX X XXXX
Xxx xx XXXXX X XXXX


Access to Nord cluster

ssh nct01YYY@nord1.bsc.es
Username: nct01YYY (specific value for YYY included in badge)
Password: ●●●●●●●●


PuTTY: SSH configuration



Before open session make sure...



MobaXterm: SSH configuration



System Overview: Nord III

- 9 iDataPlex compute racks. Each one composed of:
 - 84 IBM dx360 M4 compute nodes
 - 4 Mellanox 36-port Managed FDR10 IB Switches
 - 2 BNT RackSwitch G8052F (Management Network)
 - 2 BNT RackSwitch G8052F (GPFS Network)
 - 4 Power Distribution Units
- All IBM dx360 M4 node contain:
 - 2x E5-2670 SandyBridge-EP 2.6GHz cache 20MB 8-core
 - 500GB 7200 rpm SATA II local HDD
 - 8x 16G DDR3-1600 DIMMs (8GB/core) Total RAM: 128GB/node
- 1.9 PB of GPFS disk storage
- Interconnection Networks
 - Infiniband Mellanox FDR10: High bandwidth network used by parallel applications communications (MPI)
 - Gigabit Ethernet: 10GbitEthernet network used by the GPFS Filesystem.
- Operating System: Linux - SuSe Distribution 11 SP3

Assignment description

<https://pm.bsc.es/ftp/training/acm-ess-2019>

Table Of Contents

[Introduction](#)

[o1-Paraver](#)

[o2-Matrix multiplication](#)

[o3-TAMPI](#)

[o4-DLB](#)

[o5-IFSKER](#)

[Software Liability](#)

Quick search

Enter search terms or a module, class or function name.

Hybrid Task-Based Programming

- [Introduction](#)
- [o1-Paraver](#)
 - [Preliminary knowledge](#)
 - [Running Paraver for the first time](#)
 - [Navigating through Paraver](#)
 - [MPI+OpenMP example](#)
- [o2-Matrix multiplication](#)
- [o3-TAMPI](#)
 - [N-body benchmark](#)
 - [Heat diffusion](#)
- [o4-DLB](#)
- [o5-IFSKER](#)
- [Software Liability](#)

01 – Paraver Tutorial

- Getting familiar with Paraver
 - You can overlap this exercise with the next matrix multiplication example
 - Get traces, analyze them, get the corresponding histograms, etc.
 - Complete this sections with the Extrae/Paraver documentation [see more]:
https://tools.bsc.es/tools_manuals

01-Paraver

Preliminary knowledge

What is Paraver?

Paraver was developed to respond to the need to have a qualitative global perception of the application behavior by visual inspection and then to be able to focus on the detailed quantitative analysis of the problems.

Expressive power, flexibility and the capability of efficiently handling large traces are key features addressed in the design of Paraver. The clear and modular structure of Paraver plays a significant role towards achieving these targets.

02 – Matrix Multiplication example

- This code has been used through this tutorial
- Several version/techniques are included in the slides
 - Overlap computation in tasks and communication in master
 - Overlap between computation and communication in tasks (+ nested)
 - Can start computation of next block of C as soon as communication terminates
 - Can obtain parallelism without parallelizing fine grain the computation
- Now, you can try it, change it, test it

Play with it

My first hybrid program


```
...
for( it=0; it<nodes; it++ ) {

    mxm( m, n, a, B, (double (*)[n])&C[it][0]);

    callSendRecv(m, n, a, down, rbuf, up);

    //next C block circular
    i = (i+n)%m;
    //swap pointers
    ptmp=a; a=rbuf; rbuf=ptmp;

}
...
```



```
void mxm( int m, int n, double *a,
          double *b, double *c ) {
    ...
    for( i=0; i<n; i++) {
        dgemv(..., a, b, c);
        c+=n; a+=m; b+=n;
    }
}
```

```
void callSendRecv(int m, int n,
                  double (*a)[m], int down,
                  double (*rbuf)[m], int up)
{
    ...
    MPI_Sendrecv( a, ..., rbuf, ...);
}
```




20

My first hybrid program

```
...
for( it=0; it<nodes; it++ ) {
    #pragma omp task in([n][m]a, B) \
        inout(C[i:n][0:n])
    mxm( m, n, a, B, (double (*)[n])&C[it][0]);
    #pragma omp task in([n][m]a) out([n][m]rbuf)
    callSendRecv(m, n, a, down, rbuf, up);

    //next C block circular
    i = (i+n)%m;
    //swap pointers
    ptmp=a; a=rbuf; rbuf=ptmp;

}
#pragma omp taskwait
...
```



```
void mxm( int m, int n, double *a,
          double *b, double *c ) {
    ...
    for( i=0; i<n; i++) {
        dgemv(..., a, b, c);
        c+=n; a+=m; b+=n;
    }
}
```

```
void callSendRecv(int m, int n,
                  double (*a)[m], int down,
                  double (*rbuf)[m], int up)
{
    ...
    MPI_Sendrecv( a, ..., rbuf, ...);
}
```



26



- Can obtain parallelism without parallelizing fine grain the computation

03 – TAMPI Exercises

- **N-Body simulation** kernel: system of bodies interacting one each other
- **Heat diffusion** kernel: Gauss-Seidel method solving the heat equation
- Goals of this exercise
 - Analyze execution, is there comm. serialization?
 - Adapt the code to exploit TAMPI
 - Compare performance results with the initial version
 - Compare traces before/after changes

Task-Aware MPI (TAMPI)


- External **library** that targets hybrid OpenMP + MPI applications
- With this library, applications can benefit from...
 - **Taskifying** both computation and communication phases
 - Safe **parallel** execution of **communication** tasks
 - Exposing **more parallelism** (removing artificial deps.)
 - Natural **overlap** of communications and computations
- Compatible with **mainstream MPI** implementations
- Currently supported by
 - **OmpSs-2** (v. 2018.11 or greater)
 - **Derivative** implementation of LLVM **OpenMP** (yet to be released)



8

Non-Blocking Mode (OpenMP & OmpSs-2)

- Targets the efficient execution of **non-blocking MPI operations** from inside tasks, but **avoiding the pause/resume** of those tasks!
- So, we allow tasks to **bind their completion** to the **finalization** of one or more MPI requests
 - usually implies the release of task dependencies, the freeing of its data structures, etc
- **TAMPI_Iwait** and **TAMPI_Iwaitall**
 - New **asynchronous** and **non-blocking** functions that bind the completion of the calling task to the finalization of the passed MPI requests
 - The calling task will complete when (1) it finishes its execution AND (2) all requests bound during its execution complete
- Users should declare the proper data **dependencies** on the comm. **data buffers**!



14

04 – DLB Exercises


- **LULESH**: The Shock Hydrodynamics Challenge Problem
- Goals of this exercise
 - Analyze execution, is there imbalance?
 - Adapt program in order to use DLB
 - Compare performance results with the initial version
 - Compare traces before/after changes

Exploiting malleability

- Dynamic Load Balance & Resource management
 - Intra/inter process/application



- Library (DLB)
 - Runtime interception (MPIP, OMPT, ...)
 - API to hint resource demands
 - Core reallocation policy



- Opportunity to fight Amdahl's law
 - Productive / Easy !!!
 - Nx1
 - Hybridize imbalanced regions


"LeWL: A Runtime Balancing Algorithm for Nested Parallelism". M.Garcia et al. ICPP09
 "Hints to improve automatic load balancing with LeWL for hybrid applications" JPDC2014



67

Towards the throughput age

- Dynamic resource sharing/management




- Configuration independence

- Amount of resources is what really matters

- Side effects

- Nx1 can be better than pure MPI !!!
- hope for lazy programmers



70

05 – IFSKER Practice

- **IFSKernel** code: a kernel of the OpenIFS application, used in global forecast and climate reanalyzes (from ECMWF)
- Consider using other interesting codes:
 - NTCHEM, Your own application! Send us an email¹ to confirm “interest”
- Goals of this exercise
 - Analysis Phase: check program’s arguments and/or runtime’s options impacting on the behavior (e.g., number of threads, schedulers, etc.). Scalability/parallel efficiency. Traces. Histograms.
 - Argument: Is there imbalance? Is there comm. serialization? Others?
 - Determine (and justify) what is your approach:
 - Consider to adapt the code to exploit DLB or TAMPLI
 - Is there any other technique you consider it can improve the code?
 - Compare performance results and traces before/after your contribution
 - Highlights. Conclusions. Extensions/next-steps.

¹ Send an email to [htbp-submit \[at\] bsc.es](mailto:htbp-submit@bsc.es)

05 – IFSKER Practice (submit)

- Write a 5-10 pages document report. It should include:
 - Initial performance analysis (how the application behaves)
 - Scaling/efficiency factors according to...
 - Paraver traces/histograms
 - Others,...
 - Proof of concept description (your plan)
 - Implementation description (what you have done)
 - Parallelization, taskyfication, synchronization, etc.
 - Experiment results (how the application behaves now)
 - Version comparison, discussion, compare them with initial analysis
 - Experiences and conclusions (what you have learnt)
- Submit a single **.tar.gz** file (< 20 MB) including the document, source codes, traces, etc. to **htbp-submit [at] bsc.es** before **September 15th 2019**



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



**EXCELENCIA
SEVERO
OCHOA**

Thanks!