FIU FLORIDA INTERNATIONAL UNIVERSITY

Partnership for International Research and Education
A Global Living Laboratory for Cyberinfrastructure Application Enablement

LAGrid Latin American Grid

Project Title: Performance Prediction of Scientific Applications

Student: Javier Delgado, PhD student, Florida International University (FIU)
FIU Advisors: S. Masoud Sadjadi, PhD, FIU and Dr. Malek Adjouadi, PhD, FIU.
PIRE Industry Partner Advisor: Liana Fong, IBM
PIRE International Partner Advisor: Esteban Clua, PhD, Universidade Federal Fluminense (UFF)

National Science Foundation

I. Research Overview and Outcome

Performance Prediction

Motivation. When many resources are available for use, performance prediction provides a way to select the best resource(s) for the job. Criteria for "best resource" can be cost, execution time, reliability, etc. For long-running scientific applications, this is particularly important. This is especially true as we move towards a model of lease-based system usage, wherein there are several competing providers, both free (e.g. Teragrid) and paid (e.g. Amazon EC2) on which the application can run.

Methodology. Applications are executed on various systems with different configurations. Resource usage characteristics are profiled throughout the entire execution and then recorded. This data is then used for prediction. The prediction module is purely mathematical. User input determines the parameters that best model the application.

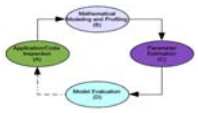

Methodology

Understanding the program being modeled helps us to obtain good prediction accuracy. At UFF, a thorough inspection of two WRF modules was conducted. We also examined the behavior of WRF, with different inputs, using different profiling tools. This gave us a better understanding of how WRF works, which helps us in selecting parameters that model its execution behavior.

Several "benchmarks" were run to see how WRF behaves with various inputs. The inputs we changed were: CPU architecture, compute node interconnection, WRF input domain size and resolution. At IBM, three different POWER systems were available. Some of these were virtual nodes, which allow setting of different clock speeds and memory availability. At UFF, we began to benchmark with certain WRF modules executing on the GPU.

In the *Prediction Results* section below, a chronological overview of the results obtained with this prediction paradigm is shown. As can be seen, modeling input domains is not trivial with WRF. In addition to fine-tuning the estimation of these parameters, we intend to use this paradigm on an up and coming execution platform - GPUs. My work at UFF was the first step in this direction.

Virtual Data Center with diverse resources

Marenostrum

- 2k+ nodes
- Power CPUs
- 4 cores each
- Myrinet
- Barcelona

Mind

- 16 nodes
- Intel Xeon
- Gigabit Ethernet
- Miami

Bluegene

- Thousands of ASIC nodes
- New York

Lincoln

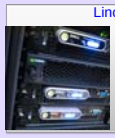

- 192 nodes
- 96 Teslas
- 4 cores each
- Infiniband
- Illinois

GPU Programming

Motivation. Traditional CPUs are no longer the only computational devices viable for HPC. Graphical Processing Units (GPUs), along with several other kinds of stream processors are being successfully used to speed up several scientific applications, including WRF. **Goals.** Our project in Brazil was to port an existing WRF module to work on GPUs. This is the first step in a larger project to be able to do performance prediction on GPU architectures.

Observations. Turned out to be an interesting software engineering problem. **Challenges.** Porting to CUDA requires rewriting some code so that the GPU is used. Since C is is the only language compatible with CUDA, the entire Fortran module had to be translated. **Results**. Developed methodology for porting from Fortran to CUDA.

Prediction Results

- Mind and GCB
  - Up to 16 nodes
  - 1 task per node
  - 20-100%CPU
  - Mean Error:

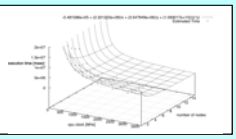

- Marenostrum
  - Up to 128 nodes
  - 1 task per node
  - 1 domain
  - Mean Error:

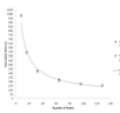

- 970MP
  - 1,2,&3 nodes
  - 1,2,&4 tasks each
  - 1 domain
  - Mean Error: 2.63%

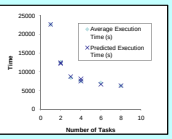

- Power 6
  - 1,2,4,8 tasks
  - 2 input domains
  - 1 output domain
  - Mean Error: 21.13%

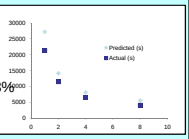

II. Industry and International Experience

Landmarks




Landmarks




Industry

Differences
- More business-oriented
- More meetings
- More secretive

Benefits
- Provided a new perspective
- Exposed industry interest in topics such as cloud computing
- Networking

Similarities
- Similar research topics
- Similar resources available
- Similar hours

Lifestyle
- Surrounded by "famous people"
- Creativity is encouraged
- No-nonsense

Academia

Benefits
- Cultural Experience
  - Love of music and sports
  - hospitality
- Learning about problems being solved at UFF
- Sharing my research with an international audience
- Strengthened collaboration and friendships

Differences
- Less HPC
- Greater sense of family

Similarities
- Large population
- Independent
- Openness to others

Research Evaluation

In industry, there is more of a concern for practical uses of research. By presenting my research approach and goals to various IBM researchers, some with many years of experience, I obtained valuable feedback that I had not gotten by merely presenting to fellow students or even at conferences.

Interoperability

At IBM, I was involved in some discussions regarding the use of IBM's scheduler, *Loadleveler*, with a metascheduler being developed at FIU. This introduced me to some of the problems that need to be worked out when interoperating software with more stringent usage policies.

Collaboration

At UFF, I learned that, like in the US, sometimes companies seek the cutting edge knowledge of academics to solve pressing problems. One example is *Petrobras*, a major Brazilian oil company. Graduate students at UFF are looking into ways to use CUDA to improve the software used by *Petrobras* to determine where to dig for oil.

Personal Impact

Last year, I got my first introduction to international research and lifestyle. This opportunity allowed me to see, first hand, what others are doing. This year, I got more international experience, but this time at a very different country. This time around, I also had the advantage of working very closely with a student from Brazil on the same project. In addition, I was able to experience six weeks at another kind of partner institution – a major company (IBM). This experience gave me a feel for how everything comes together. I saw how cutting-edge applications being looked at in academia are analyzed by companies like IBM (or *Petrobras*) for possible use in the business world.

III. Acknowledgement

The material presented in this poster is based upon the work supported by the National Science Foundation under Grant No. OISE-0730065. The author appreciates NSF grant OCI-0636031, for connecting me to my international advisor, and grant number HRD-0833093 for supporting me throughout my research.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.