MAP-I doctoral program

Multi-core and Many-core Computing

1. Theme, motivation and context

At the end of last century CPUs have doubled in performance roughly every 18 months, due to both an increase in design complexity and to faster CPU clock speeds, supported by advances in chip fabrication technology. However, there are barriers to further significant improvements in operating frequency due to voltage leakage across internal chip components and heat dissipation limits [1]. Intel abandoned the development of Pentium processor at 4GHz in 2004.

To further increase performance on pipelined superscalar processors that already included vector extensions at the ISA level, designers moved into multithreaded CPUs. These employ hardware-level context switching between threads, to reduce the idle time of resources in complex superscalar processors. Designers then integrated more than one processing core into a single chip, and we are currently moving into 2-figure-core CPU devices.

In multi-core architectures the memory hierarchy is assuming an increasing level of complexity. The known latency to access off-chip memory is increased by the concurrent access from the different cores. So far, designers have increased the on-chip cache and are testing private/shared arrangements among the on-chip 3 cache levels, but researchers are still arguing on the best approaches, and new moves are towards a system view with no shared memory, with severe implications on data locality.

Assuming that Moore's law will hold in the multi-core era, it is expected a doubling of the number of cores per chip every two years, leading shortly to many-core CPU based devices. And these devices are populating standard laptops and desktops.

To further increase the complexity to efficiently design software for these devices, computing platforms are becoming highly heterogeneous, mixing devices with few but complex cores, with others to support massively parallelism based on simpler cores, aiming a wide range of applications, including graphics, regular and irregular numerical applications, network processing, and cryptography. These range from dedicated or general-purpose graphic processing units (GPU or gpGPU, namely those from NVidia and AMD), to FPGA-based solutions (from Xilinx and Altera) to ASIC-based solutions. Novel programming paradigms, tools and methods are being developed to also explore these massively parallel architectures, some as extensions to popular environments (CUDA, OpenCL), others as completely new approaches (Chapel).

To efficiently orchestrate the computational power available in these highly parallel computing devices, the new generations of software developers require deeper knowledge of the underlined system architecture, and access to adequate development tools to evaluate the overall system performance in these heterogeneous environments. Software tuning is highly recommended in most cases, while robust and user-friendly solutions are not yet developed.

It is our responsibility to adequately educate these new generations of software developers, integrating into the desktop programming concepts that are being applied in High Performance Computing, and adding the new requirements for data locality.

Related courses in other institutions:

1. University of Illinois at Urbana-Champaign
   ECE 498 AL: Applied Parallel Programming
   http://courses.ece.illinois.edu/ece498/al/index.html
   Instructor: Sanjay J Patel

2. RICE University
   COMP 522 Multi-core Computing
   http://www.cs.rice.edu/~johnmc/comp522/
   Instructor: John Mellor-Crummey

3. Georgia Tech
   CS 8803-MC, CS 4803-MC: Multicore Computing
   http://www.cc.gatech.edu/~bader/COURSES/GATECH/CS8803-Spring2007/
   Instructor: Dr. David A. Bader

4. Georgia Tech
   CS 8803 MCA – Manycore Computer Architecture, Spring 2009
   http://www.cc.gatech.edu/grads/p/pkuber/SylS09.pdf
   Instructor: Tom Conte

2. Objectives and Learning Outcomes

Programming for multi-core and many-core architectures require two different programming models mainly due to the memory organization. Current multi-core CPU-based devices have on-chip shared memory, which supports course-grain shared memory parallelism. On the other hand, many-core devices usually have hundreds of simple computing cores organized in segments where memory can only be shared inside each segment. This programming model is closer to the data parallel model.

We identify two sets of interesting component parts that can help to structure course material in **multi-core and many-core computing**. Each set is covered in several lectures that close with a case study.

The work will be challenging and rewarding, enhancing each student's ability to work in this rapidly advancing field and giving participants experience to multi-core and many-core programming. Students will find themselves to:

- be familiar with the architecture of multi-core devices, including the multi-layer memory hierarchy and the inter-core communication facilities;
- be familiar with the architecture of many-core devices, mainly the GPU-type, including their internal organization and structure;
- be able to develop hybrid programs for multi-core and many-core architectures;
- be familiar with the emerging tools and environments for multi-core and many-core devices;
- be familiar with parallel program analysis and debugging tools.

3. Detailed Syllabus

1. Review on architectures and parallel programming fundamentals:
   a. Multi-core and many-core processors architecture.
   b. Problem division, data dependencies (communication patterns), synchronization, granularity of parallelization, data locality analysis and scheduling.

2. Tools and environments for multi-core and many-core programming
   a. Programming APIs and languages for multi-core: OpenMP, Chapel, TBB (Intel Thread Building Blocks)
   b. Programming APIs and languages for many-core: CUDA, OpenCL.
   c. Environments for multi-core: Cetus [2], Intel Parallel Studio.

3. Characterization of parallel computing
   a. Execution models and computing models
   b. Efficiency and performance measures
   c. Scalability analysis (isoefficiency function)

4. Parallel program analysis
   a. Software instrumentation.
   b. Tools for parallel program debugging and performance tuning.

4. Teaching Methods and Evaluation

The overall aim is to give students a broad and well-balanced understanding of multi-core and many-core computing that will serve well as a foundation for more specific work or research.
The Course will include lectures and discussions on the principles, technologies, experience and exploitation of these architectures.
Students will develop mini projects (Course Work) during the semester to apply and develop expertise in the course content.
Invited lectures and researchers from Univ. Texas in Austin, Univ. Illinois Urbana-Champaign and NVidia will provide advanced seminars similar to the ones that have been organized at UMinho in 2009 and 2010, which will also support discussions on the case studies. These activities have the support of the UTA cooperation program with Portugal.
The grading will be composed by the Course Work (50%) and by Exam (50%).

5. Main Bibliography

[1] D. Patterson, J. Hennessy, ""Computer Organization and Design. The Hardware/Software Interface", 4th Ed., Morgan Kaufmann, 2009
[2] M. Quinn; “Parallel programming in C and OpenMP”, McGraw Hill, 2003
[3] Chapman B., Jost G., Kuck D., Pas R.; “Using OpenMP”, The MIT Press, 2008
[4] D. Kirk, W. Hwu, "Programming Massively Parallel Processors: A Hands-on

Approach", Morgan Kaufmann, 2010
[5] Sodan, A.C. Machina, J. Deshmeh, A. Macnaughton, K. Esbaugh, B.; "Parallelism via Multithreaded and Multicore CPUs", IEEE Computer, March 2010
[6] John A. Stratton , Sam S. Stone , Wen-Mei W. Hwu, "MCUDA: An Efficient Implementation of CUDA Kernels for Multi-core CPUs, Languages and Compilers for Parallel Computing": 21th Int. Workshop, LCPC 2008, Edmonton, Canada, July 2008, Revised Selected Papers, Springer-Verlag, Berlin, Heidelberg, 2008
[7] S. Che, J. Li, J.W. Sheaffer, K. Skadron, and J. Lach. Accelerating compute-intensive applications with GPUs and FPGAs. Proc. Symp. Application Specific Processors, 2008
[8] Dave, C. Hansang Bae Seung-Jai Min Seyong Lee Eigenmann, R. Midkiff, S.; "Cetus: A Source-to-Source Compiler Infrastructure for Multicores", IEEE Computer, December 2009, pp. 36-42.

B. **Instructors' team**

This team is based on 2 faculty members from U. Minho (Alebrto Proença) and U. Porto (Jorge Barbosa) that persistently run the course, together with other faculties from US that cover some particular modules.
At least two professors have already agreed to give some specific courses:

- **Keshav Pingali**, Professor at the Dep. Computer Science and at the Institute for Computational Engineering and Science at Univ. Texas in Austin, W.A."Tex" Moncrief Chair of Grid and Distributed Computing; he is currently the Director of the Focus Theme on Advanced Computing at CoLab@Austin, the cooperation program between POrtugal and the Univ. Texas in Austin;
- **Wen-mei Hwu**, the Walter J. ("Jerry") Sanders III-Advanced Micro Devices Endowed Chair in Electrical and Computer Engineering in the Coordinated Science Laboratory of the University of Illinois at Urbana-Champaign; Wen-mei Hwu has co-authored an important new textbook that breaks down the complexities of parallel programming and the GPU architecture to enable programmers to address the critical challenges of massive parallelism: *Programming Massively Parallel Processors: A Hands-on Approach*, co-authored with David Kirk, an NVIDIA Fellow and former Chief Scientist.

Alberto José Proença: Professor at Dep. Computer Science, University of Minho, since January 2001. He lead the creation of the research center CCTC presented at FCT, where he is a member. Received the MSc and PhD degrees from University of Manchester, UK, in 1979 and 1982. He is currently the national Director of the Advanced Computing Program at CoLab, under the agreement between Portugal and the Univ. Texas in Austin. Current R&D interests lie on high performance and computer imaging, namely on the exploitation of system and devices capabilities to increase the performance of graphics/vision applications, including functional parallelism (vector co-processors), and CPU parallelism and accelerators in heterogeneous cluster environments.

Coordination of Projects:

SeARCH: Services and Advanced Research Computing with HTC/HPC clusters
REEQ/443/EEI/2005 (Infrastructures)
Jan'05– Dec'09
Budget: €286.6k

CROSS-Fire - Collaborative Resources Online to Support Simulations on Forest Fires: a Grid Platform to Integrate Geo-referenced Web Services for Real-Time Management
GRID/GRI/81795/2006 (FCT)
Aug'07 – Aug'10
Budget: €170k

CICH: Computer Imaging in the Cultural Heritage
Contract with a non-profit US Institution and a project proposal submitted to FCT
Jun'09 – Dec'10
Budget: €20k (proposal: €200k)

Selected Publications:

L.P. Santos, V. Coelho, P. Bernardes, A. Proença, "High Fidelity Walkthroughs in Archaeology Sites"; 6th Int. Symp. Virtual Reality, Archaeology and Cultural Heritage (VAST'2005); Italy, November 2005

J. Fernando, J. Sobral, A. Proenca. "JaSkel: "A Java Skeleton-Based Framework for Structured Cluster and Grid Computing", 6th IEEE Int. Symp. Cluster Computing and the Grid (CCGrid'2006), IEEE Computer Society, Singapore, May 2006

A. Oliveira, L.P. Santos, A. Proença, "Refinement Criteria for High Fidelity Interactive Walkthroughs"; 4th Int. Conf. Computer Graphics and Interactive Techniques (Graphite'2006), Australasia and South-east Asia, pp 453..460; ACM SIGGRAPH; Kuala Lumpur, Malaysia, December 2006

Sobral, J., Proença, A.; Enabling JaSkel Skeletons for Clusters and Computational Grids, Proc. IEEE Cluster (Cluster 2007), Austin, Texas, September 2007

Proença, A.J., Barbosa, J., Sobral, J.L.; Imaging Techniques to Simplify the PTM Generation of a Bas-Relief, Proc. 8th Int. Symposium on Virtual Reality, Archaeology and Cultural Heritage, VAST´2007, Brighton, UK, November 2007

Mudge, M., Malzbender, T., Chalmers, A., Scopigno, R., Davis, J., Wang, O., Gunawardane, P., Ashley, M., Doerr, M., Proenca, A. Barbosa, J., Image-Based Empirical Information Acquisition, Scientific Reliability, and Long-Term Digital Preservation for the Natural Sciences and Cultural Heritage. Full Day Tutorial, Eurographics 2008, April 14-18, 2008, Crete, Greece

A. Pina, B. Oliveira, J. Puga, R. Marques, A. Proença, "An OGC-WS Framework to Run FireStation on the Grid, Proc. 3rd Iberian Grid Infrastructure" Conf (IBERGRID´09), Valência, Spain, May 2009

A. Pina, B. Oliveira, J. Puga, A. Esteves, A. Proença.; FireStation on the grid - a step further on the adoption of OGC/SDI standards. Proc. Enabling Grids for E-sciencE Conf. (EGEE'09), Barcelona, Spain, September 2009

A. Esteves, M. Caldas, A. Pina, A. Proença.; An OGC/SOS Conformant Client to Manage Geospatial Data on the GRID, Proc. 4th Iberian Grid Infrastructure Conf (IBERGRID´10), Braga, Portugal, May 2010

Jorge Manuel Gomes Barbosa: is a lecturer at the Department of Electrical and Computer Engineering of University of Porto, and a research member of LIACC (Laboratório de Inteligência Artificial e Ciência de Computadores), being the research activities related to parallel algorithms for biomedical applications, scheduling and performance modelling. He obtained his BSc degree in Computer Science from University of Porto Faculty of Engineering, in 1992, and his MSc degree in Digital System from the University of Manchester Institute of Science and Technology, in 1993, and the PhD in Computer Science from the University of Porto Faculty of Engineering, in 2001. From 2003 to present, he is responsible of the discipline "Parallel Computing" of the MSc course in Informatics of Faculty of Engineering of the University of Porto. He participates in the Cost Action IC0805 "Open Network for High-Performance Computing on Complex Environments".

Supervision

- Ph.D. students:
  - Tito Vieira, "Utility Computing on Scientific Cloud Infrastructures", ProDEI, FEUP, 2009.
  - Carlos Costa, "Dynamic parallel task scheduler for multi-user jobs on heterogeneous clusters", ProDEI, FEUP, 2009.
  - Altino Sampaio, "Virtualization Management on a Scientific Cloud Computing Infrastructure", ProDEI, FEUP, 2010.
  - Tiago Santos, "Communication APIs for Virtual Clusters", ProDEI, FEUP, 2010.

Financed Research Projects:

- Rede Grid de Imagiologia Cerebral, GRID/GRI/81833/2006, Team budget 31560 Euros.
- GERES-med: repositórios em GRID para aplicações médicas, GRID/GRI/81819/2006, Team budget 61560 Euros.
- Re-equipment project: in 2002 collaborated in the proposal presented by Biomedical Engineering Institute to acquire a cluster infrastructure with 72 processors, which was installed in December 2006.

Publications:

- [1] J. Barbosa, Belmiro Moreira, "Dynamic scheduling of a batch of jobs on heterogeneous clusters", Submitted to Parallel Computing, Elsevier, 2010.
- [2] J. Barbosa, Belmiro Moreira, "Dynamic job scheduling on heterogeneous clusters", in International Symposium on Parallel and Distributed Computing, IEEE CS, pp.3-10, 2009
- [3] J. Barbosa, M. Monteiro, "A List Scheduling Algorithm for Scheduling Multi-user Jobs on Clusters", in High Performance Computing for Computational Science - VECPAR 2008, 8th International Conference, LNCS 5336, pp.123-136, 2008
- [4] J. Barbosa, J. Tavares, A.J. Padilha, "Optimizing Dense Linear Algebra Algorithms on Heterogeneous Machines" Algorithms and Tools for Parallel Computing On Heterogeneous Clusters, Nova Science Publisher, N.Y., pp 17-31, ISBN: 1-60021-049-X, 2007
- [5] R. Nóbrega, J. Barbosa, A.P. Monteiro, "BioGrid Application Toolkit: a Grid-based Problem Solving Environment Tool for Biomedical Data Analysis",

VECPAR'06 - 7th International Meeting on High Performance Computing for Computational Science, Rio de Janeiro, pp. 1-13, 2006
- [6] J. Barbosa, C. Morais, R. Nóbrega, A.P. Monteiro, "Static scheduling of dependent parallel tasks on heterogeneous clusters", Heteropar05 - 4th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks, IEEE CS Press, Boston, EUA, pp. 1-8, 2005
- [7]J. Barbosa, C. Morais, A.J. Padilha, "Simulation of data distribution strategies for LU factorization on heterogeneous machines", 12th Heterogeneous Computing Workshop, IEEE CS Press, Nice, pp. 1-8, 2003

Scientific Activities

- Since 2003, member of the Scientific Committee of the "Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks – Heteropar", integrated on IEEE Cluster conference.
- Member of the Scientific Committee of the IASTED International Conference on Parallel and Distributed Computing and Networks, PDCN 2007/8/9/10.
- Reviewer of scientific papers submitted to the "Vecpar'08 - International Meeting on High Performance Computing for Computational Science".
- Reviewer of IEEE International Symposium on Parallel and Distributed Computing, 2009.
- Reviewer of the international journal Parallel Computing - Systems & Applications, Elsevier, 2010.