AutoTune: Automatic Online Tuning

Renato Miceli
Irish Centre for High-End Computing (ICHEC)
renato.miceli@ichec.ie

Who am I?

- Mr. Renato Miceli, Brazilian
  - Computational Scientist and GPU Developer at ICHEC – Irish Centre for High-End Computing
  - PhD candidate at Université de Rennes 1, France
  - Interests in Software Engineering, Distributed and Parallel Computing
  - Main technical person and contact at ICHEC for the AutoTune Project

About ICHEC

- Irish Centre for High-End Computing
- The Irish HPC resources centre
- One keyword: Enablement
- GPU computing
  - Since 2009
  - 7th NVIDIA CRC
  - 2nd HMPP CoC

http://www.ichec.ie/

Motivation

- Pervasiveness of multi-/many-core processors
- Complexity in programming and optimizing codes for such hybrid architectures:
  - Energy consumption
  - Inter-process communication
  - Load balancing
  - Data locality
  - Memory access
  - Single-core performance

Motivation

- Programming and optimizing is guided…
  - By experience
  - By the output of profiling
  - By the output of tracing
- Number of possible tunings are immense
  - Exponential explosion

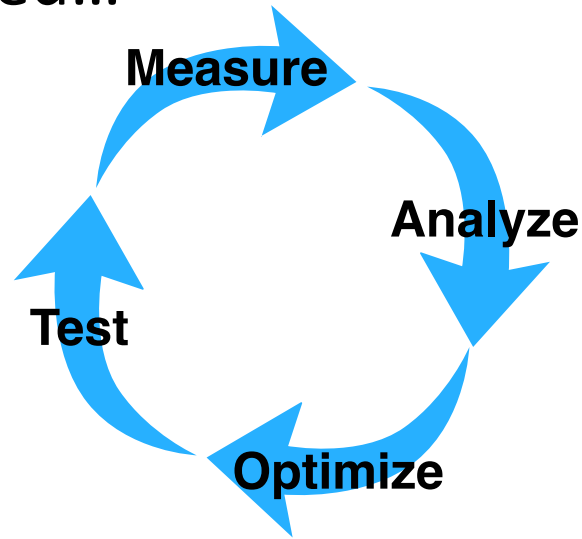

These tasks are only partially automated by existing tools

Research Objectives

AUTOTUNE
www.autotune-project.eu

- Develop heuristics for auto-tuning of C/Fortran codes…
  - For performance (execution runtime)
  - For energy consumption (watt-hours)
- … on various computer architectures…
  - HPC and parallel servers
  - Homogeneous and heterogeneous
  - Multicore and GPU accelerated systems
- … for a variety of parallel paradigms
  - MPI, HMPP, parallel patterns

Approach: Take Periscope…

AUTOTUNE
www.autotune-project.eu

- Performance analysis tool
  - **Automatic** (no human intervention)
  - **Online** (no application restart, no trace files)
  - **Distributed** (processes analyses in parallel, low network utilisation)
  - **Highly scalable** (run on 65k Blue Gene cores, able to run on 100k cores)
  - **Portable** (Fortran/C with MPI & OMP on x86, Intel Itanium2, IBM Power6, BlueGene P, Cray, and virtually every other parallel architecture)

Instruments, runs and analyses an application live

… extend its current features…

AUTOTUNE

www.autotune-project.eu

Architecture of Periscope

Developed by Technische Universität München, Periscope is an automatic performance analysis tool for highly parallel applications written in MPI and/or OpenMP. Periscope performs program analysis while the application is executing (online) using a set of analysis agents in a distributed fashion. It is available for all major parallel architectures and has production-level quality. Its scalability has been demonstrated on large BlueGene systems by analyzing applications running on up to 65,000 cores.

1. User
Defines a C/Fortran project on Eclipse with all source files, starts a performance analysis using the GUI and investigates the performance properties found

2. Front-End
Invokes the application, creates agent hierarchy and broadcasts a command for all analysis agents to start searching for performance properties

3. Agent hierarchy
The number of application processes and processors requested influence the total number of agents created. Processes and agents are mapped to the available processors such that communication is local regarding the physical interconnect network topology

4. Master Agent
Propagates commands from the front-end to analysis agents and provides a list of performance properties found to the front-end. Each found property points to code regions that may benefit from further tuning

5. Communication Agent
Transmits information between the Master Agent and the Analysis Agents

6a. Analysis Agent
Selects performance properties and required measurements to look, based on a search strategy, for performance issues in a subset of the application's processes. Search strategies define the order in which an analysis agent investigates the multidimensional search space of properties, program regions and processes

6b. Performance analysis
Conducted by analysis agents using search strategies. Many implemented strategies consist of multiple analysis steps and the performance properties found in one step are refined into more precise properties at the next step

7. MRI Monitor
Measures performance data requested by its analysis agent and controls the execution of the process, taking commands from the agent to halt/release the execution

8. Application Process
The application should be highly parallel and use MPI and/or OpenMP

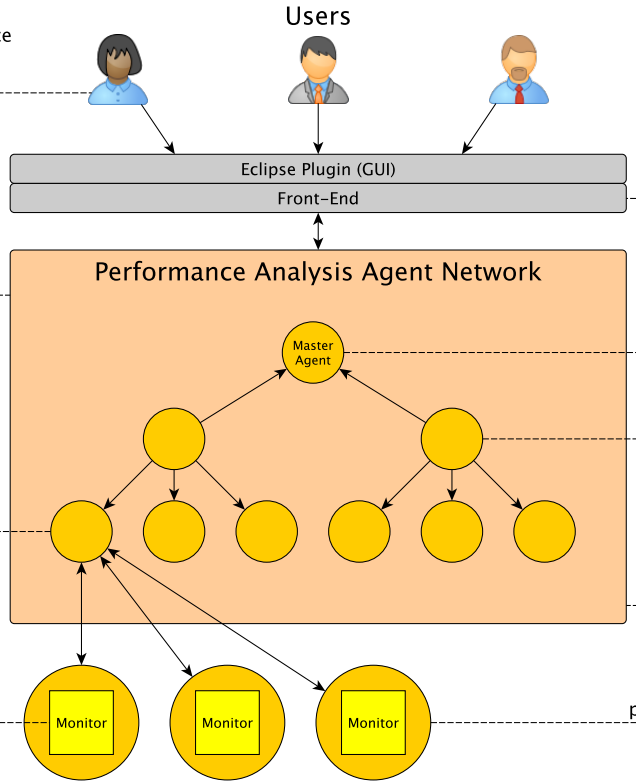

… to develop the PTF tool

AUTOTUNE
www.autotune-project.eu

- The Periscope Tuning Framework
  - Why framework? It can be further extended with a number of plugins
  - PTF will…

1. Invoke Periscope to analyze and capture issues
2. Select plugins whose tuning strategies will be executed
   - Compiler-based optimisation
   - HMPP tuning for GPUs
   - Parallel pattern tuning
   - MPI tuning
   - Energy efficiency tuning
3. Allow for the tuning recommendations to be automatically applied

Project Validation

1. Applications selection and manual tuning
2. Execution of PTF over the raw version
3. Comparison of improvements between raw and tuned versions
   - PTF will obtain at least 50% of manual improvements, or even surpass them (>100%);
   - PTF will require only a single or a few application runs, compared to effort timed in months for manual tuning.

Resources

- Cluster Stoney in Ireland
  - Provides Ireland's National GPU service
  - 24 compute nodes reserved for GPGPU computing
  - Each node has 2 NVIDIA Tesla M2090 (Fermi) cards attached
    - 512 CUDA cores
    - 6 GB local GDDR5 memory
    - Theoretical double-precision peak perf. of 665 GFlops

System	Stoney
Service	National
Manufacture / Model	Bull Novascale R422-E2
CPU	Intel Xeon X5560
CPU Clock	2.8 GHz
CPU Cores	512
Memory	3072 GB
Peak Performance	5.73 TFlop
Linpack Performance	5.14 TFlop
Interconnect	ConnectX Infiniband (DDR)
Storage (Formatted Capacity)	21 TB
Launched	Jul 2009

Resources

- Cluster SuperMUC in Germany
  - Under deployment by IBM
  - Expected to be fully operational in 2012
  - Its rough specs are…
    - Intel Xeon architecture
    - More than 110,000 cores
    - Peak performance of about 3 PFlops/s

And who is behind all this?

AUTOTUNE
www.autotune-project.eu

Technische Universität München, Germany

Universität Wien, Austria

CAPS Entreprise, France

Universitat Autònoma de Barcelona, Spain

Leibniz Computing Centre, Germany

Irish Centre for High-End Computing, Ireland

IBM, Germany

Timeline

AutoTune
www.autotune-project.eu

Month 12 | Month 24 | Month 32 | Month 36

Cycle 1 — Cycle 2 — Cycle 3 — Cycle 4

AutoTune Project

- Specification of the tuning model
- Results of the manual tuning of selected applications demonstrating the potential of the tuning techniques
- Detailed technical specification for all work packages
- Extended version of Periscope's monitor for the tuning plugins
- PTF Demonstrator demonstrating the integration of PA and tuning plugins

- Prototype versions of the tuning plugins
- Single plugin tuning strategies
- PA strategies for HMPP/OpenCL and energy efficiency
- PTF Integrated Prototype demonstrating single plugin tuning for the selected applications

- Final tuning plugins and combined plugin tuning strategies
- PTF Release

- Documentation of PTF
- Detailed evaluation
- Demonstration of the promised automatic improvements for the selected applications

We are here!

Latest News

- Project is very active
  - New things are constantly happening!
- Current status:
  - The tuning model and plugin strategies are being formalized
  - The repository of applications for validation is being finished
  - A PTF demo is planned for October/2012

Want to know more?

AUTOTUNE
www.autotune-project.eu

- Project poster at NVIDIA GTC 2012
  - May 14th in San Jose, USA
  - Available at GTC On-Demand website
- Tutorial at PACT-2012
  - September 19th ~ 23rd in Minneapolis, USA
- Visit the project website:
  - http://www.autotune-project.eu/
- Drop us an email
- Or else, come chat with me!

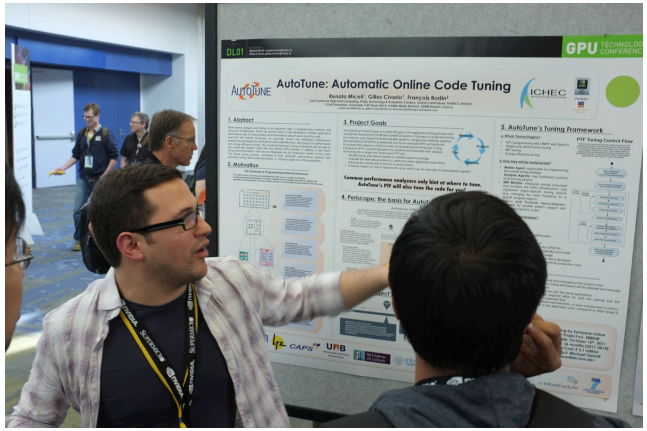

Any questions?

Thank you