EasyPAP: a Framework for Learning Parallel Programming

Alice Lasserre, Raymond Namyst, Pierre-André Wacrenier firstname.lastname@u-bordeaux.fr

Dept. of Computer Science

A case for a comprehensive framework

- Parallel programming is not trivial
 - Debugging is entering a world of pain
 - Understanding (bad) performance is even more challenging
- Like many teachers, we progressively added visualization facilities to our lab applications
 - Increased student's motivation
 - Greatly helped to improve correctness
- EasyPAP goes further
 - Minimize time spent to become familiar with new problems
 - Enable quick OpenMP/MPI/OpenCL prototyping
 - Provide simple tools to analyze parallel behavior

EasyPAP: focus on parallelism!

- C library + utilities
 - Support for Pthreads, OpenMP, MPI, OpenCL
- Online rendering of 2D computations
 - Work distribution monitoring
- Trace visualization
 - Side-by-side comparison
- Plotting facilities
 - Thorough experiments & analysis



Online visualization and thread monitoring





Offline trace visualization and plotting facilities

Kernels and variants

- Students are provided with sequential implementations of various *kernels*
 - Mandelbrot Set, Game of Life, Abelian Sandpiles, Picture Blur
 - ✓ Just add a C file to create a new kernel, then compile & run
- They can design and experiment with as many *variants* as they can think of
 - Kernels and variants are selected on command line
 - ✓ Just add a function to create a new variant, then compile & run



Code instrumentation and monitoring



Off-line Trace Visualization

Task scheduling chart

Task-data mapping



Trace comparison

"diff" mode: iterations are re-aligned



Plotting facilities

- Experiments can easily be automated using scripts
 - No need to recompile
 - Each run records all experimental details in a CSV file
- Plotting (python) scripts
 - Ease graph selection
 - Make sure results are sound
 - Speedups automatically computed
 - Parameter consistency check



What are the main benefits?

• Focus on parallelism

- Implement many variants
- Experiments with multiple parameters

• Quicker and deeper understanding of

- Scheduling
 - Load balancing, data affinity
- Cache
 - Tiling, false sharing
- Synchronization
 - Race conditions, barriers, task dependencies
- Hardware specific optimizations
 - Code specialization, vectorization



CPU coverage map across multiple iterations revealing task-data affinity

EasyPAP documentation and download:

http://gforgeron.gitlab.io/easypap/