

Adaptation dynamique des optimisations : versions interchangeables

Lénaïc BAGNÈRES, Cédric BASTOUL

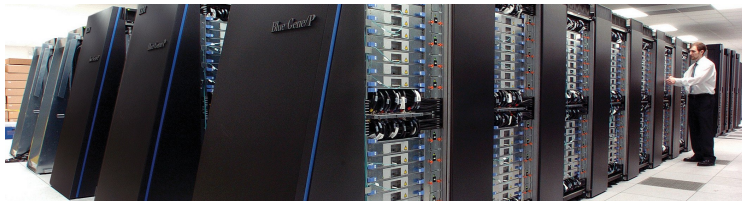
Laboratoire de Recherche en Informatique (LRI)
Inria Grand Large
Université Paris-Sud 11

Laboratoire ICube - ICPS
Inria CAMUS
Université de Strasbourg

7^e rencontres de la communauté française de compilation
6 décembre 2013

Introduction

- ▶ Once upon a time, life was easy for parallel programs...
 - One program running, one architecture, one context, one user



- ▶ In the multicore era, welcome to the jungle!
 - A large set of target architectures is possible
 - High dataset variations (user and/or architecture dependent)
 - Several parallel programs are executing at the same time



Adaptation dynamique
des optimisations :
versions
interchangeables

Lénaïc BAGNÈRES,
Cédric BASTOUL

Introduction

Execution Context &
Related Work
Our Approach

Building Switchable
Scheduling

Conclusion

References

Execution Context & Related Work

PROGRAM MUST ADAPT (TRANSPARENTLY) TO THE EXECUTION CONTEXT

- ▶ Target architecture
(e.g. number of cores, size of caches, accelerator type)
- ▶ Dataset properties
(e.g. dataset size)
- ▶ System load
(e.g. number of running process, processor load, memory access)

RELATED WORK

- ▶ Choice based on compile time profiling [Pradelle et al., 2011]
Possible prediction error, switching not possible during kernel execution
- ▶ LeTSeE [Pouchet et al., 2008]
Chooses the best version for a given architecture and context
- ▶ One application with VMAD [Jimborean et al., 2012]
Select the best version at runtime, high overhead, user develops and chooses versions
- ▶ EvolveTile [Tavarageri et al., 2011]
Perform a dynamic tile size selection
- ▶ Number of threads depends on the external workload [Emani et al., 2013]

Our Approach: Switchable Scheduling

Adaptation dynamique
des optimisations :
versions
interchangeables

Lénaïc BAGNÈRES,
Cédric BASTOUL

IN A NUTSHELL:

- ▶ Prepare a pertinent set of versions
- ▶ Versions have meeting points where it is possible to switch from one to another
- ▶ A runtime analysis drives to the "best" version at each meeting point

PROPERTIES:

- ▶ Static / dynamic approach
- ▶ Low overhead
- ▶ Benchmarking computations contribute to the final solution:
no backtrack

Introduction

Execution Context &
Related Work

Our Approach

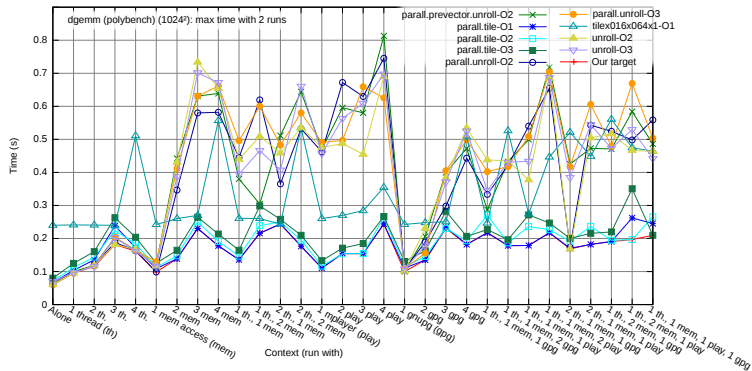
Building Switchable
Scheduling

Conclusion

References

Building Switchable Scheduling: Static Part

- ▶ Pre-selection of versions:
 - Switchable versions building engine based on PoCC and Pluto
 - Generate versions: sequential, parallel, tiled, etc.
- ▶ Pre-selected pertinents contexts
 - Architecture (number of cores, size of caches)
 - Aggressive processor use or/and memory access processes
- ▶ Selecting pertinents versions:
 - Execution on artificial contexts
 - Selection of a restricted set of complementary "well-behaving" versions



System load context: 2mm without and with strong system load on Xeon¹

Adaptation dynamique
des optimisations :
versions
interchangeables

Lénaïc BAGNÈRES,
Cédric BASTOUL

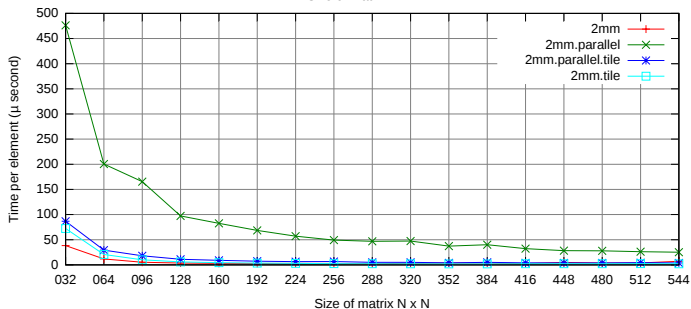
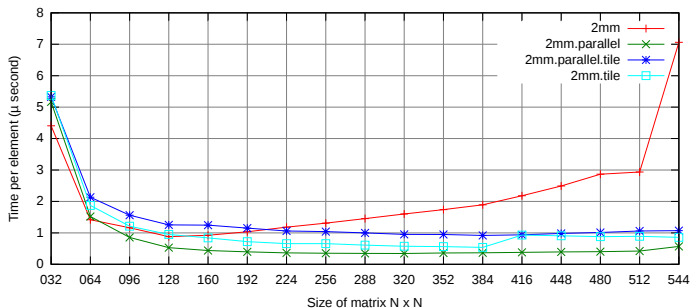
Introduction

Building Switchable
Scheduling

Static Part
Dynamic Part

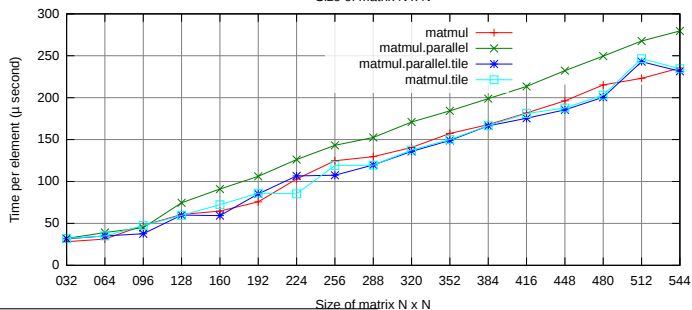
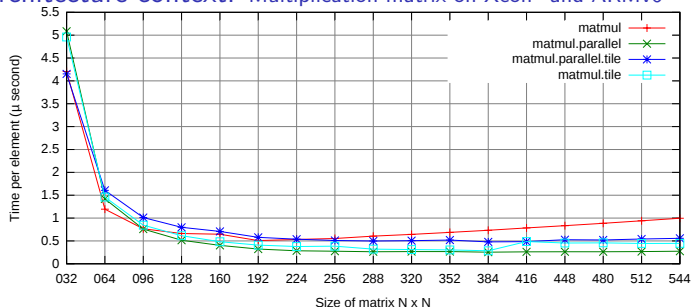
Conclusion

References



¹Xeon: Two Intel Xeon CPU E5645 6 cores 2.40GHz, Linux 3.2, GCC 4.6.3

Architecture context: Multiplication matrix on Xeon² and ARMv6³



²Xeon: Two Intel Xeon CPU E5645 6 cores 2.40GHz, Linux 3.2, GCC 4.6.3

³ARMv6: LG GW620, ARMv6 1 core 600Mhz, Linux 2.6.32, GCC ARM-Linux GNU ABI 4.7.2

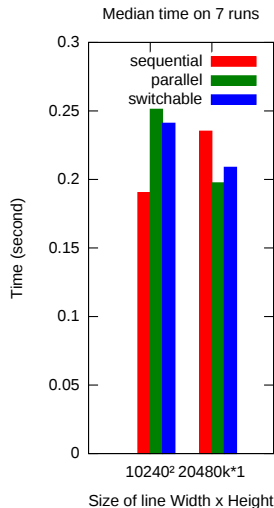


Building Switchable Scheduling: Dynamic Part

(a) selecting a version and switching at runtime

```
while (i < N)
// Sampling
for (v = 0 → versions.size())
    t0 = now()
    // Execute one version
    versions[v](i, i + sample_size)
    t1 = now()
    times[v] = t1 - t0
// Computing
// Execute previous best version
version[min(times)](i, i + compute_size)
```

(b) actual performance on line bench



- ▶ Lessons learned so far:
 - Already useful to converge to the best version with respect to the context
 - Low overhead techniques demonstrated essential (but we are not good enough...)
 - Too often, we still get one version to rule them all (no switching after convergence)

- ▶ Ongoing work and todo list:
 - More aggressive selection of versions
 - Refinement of switchable version set study at compile time using collective tuning ([Fursin, 2009], <http://ctuning.org>)
 - Refinement of switchable version set at runtime using dynamic tests à la [Pradelle et al., 2011]
 - Background pre-sampling



Emani, M., Wang, Z. and O'Boyle, M. (2013).

Smart, adaptive mapping of parallelism in the presence of external workload.

In *Code Generation and Optimization (CGO)*, 2013 IEEE/ACM International Symposium on pp. 1–10,.



Fursin, G. (2009).

Collective Tuning Initiative: automating and accelerating development and optimization of computing systems.

In *Proceedings of the GCC Developers' Summit*.



Jimborean, A., Mastrangelo, L., Loechner, V. and Clauss, P. (2012).

VMAD: an Advanced Dynamic Program Analysis & Instrumentation Framework.

In *CC - 21st International Conference on Compiler Construction*, (O'Boyle, M., ed.), vol. 7210, of *Lecture Notes in Computer Science* pp. 220–237, Springer, Tallinn, Estonia.



Pouchet, L.-N., Bastoul, C., Cohen, A. and Cavazos, J. (2008).

Iterative optimization in the polyhedral model: Part II, multidimensional time.

In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'08)* pp. 90–100, ACM Press, Tucson, Arizona.



Pouchet, L.-N., Bastoul, C., Cohen, A. and Vasilache, N. (2007).

Iterative optimization in the polyhedral model: Part I, one-dimensional time.

In *IEEE/ACM Fifth International Symposium on Code Generation and Optimization (CGO'07)* pp. 144–156, IEEE Computer Society press, San Jose, California.



Pradelle, B., Clauss, P. and Loechner, V. (2011).

Adaptive Runtime Selection of Parallel Schedules in the Polytope Model.

In *19th High Performance Computing Symposium - HPC 2011 ACM/SIGSIM*, Boston, United States.



Tavarageri, S., Pouchet, L.-N., Ramanujam, J., Rountev, A. and Sadayappan, P. (2011).

Dynamic Selection of Tile Sizes.

In *18th annual IEEE International Conference on High Performance Computing (HiPC'11)* IEEE Computer Society Press, Bangalore, India.

Lénaïc BAGNÈRES,
Cédric BASTOUL

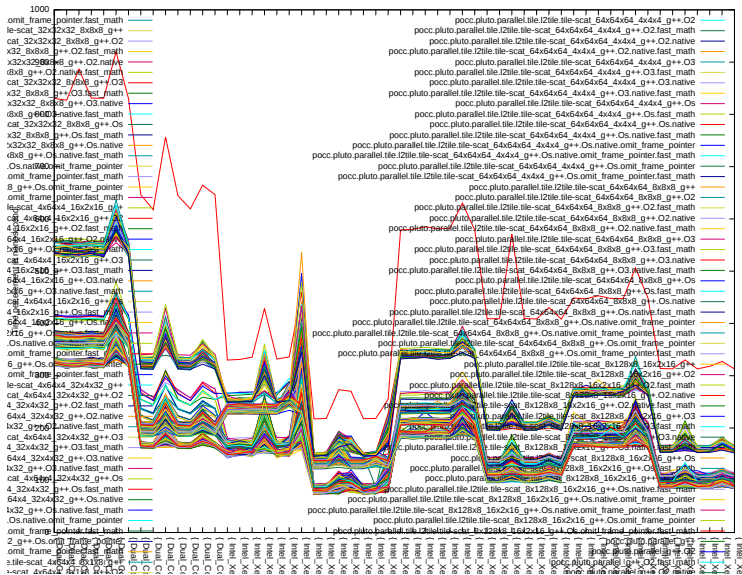
Introduction

Building Switchable
Scheduling

Conclusion

References

All version of benchmark jacobi_2d_imper (867 versions, 56 contexts)



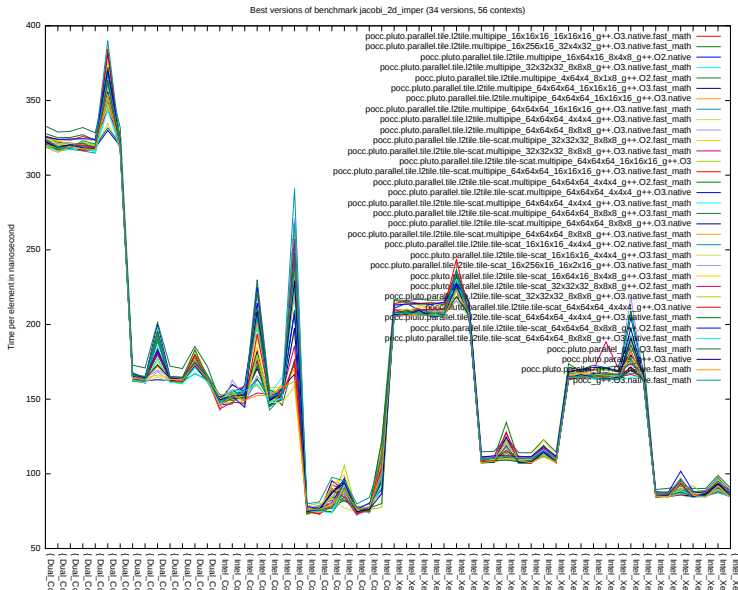
Lénaïc BAGNÈRES,
Cédric BASTOUL

Introduction

Building Switchable
Scheduling

Conclusion

References



Lénaïc BAGNÈRES,
Cédric BASTOUL

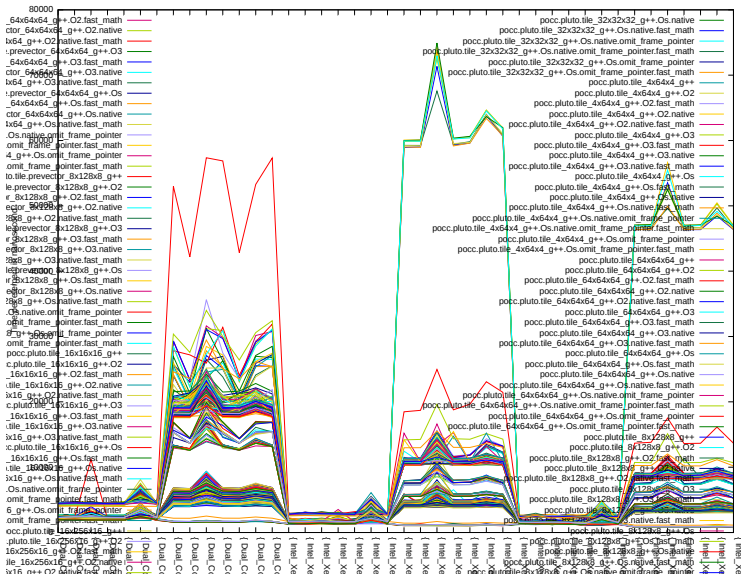
Introduction

Building Switchable
Scheduling

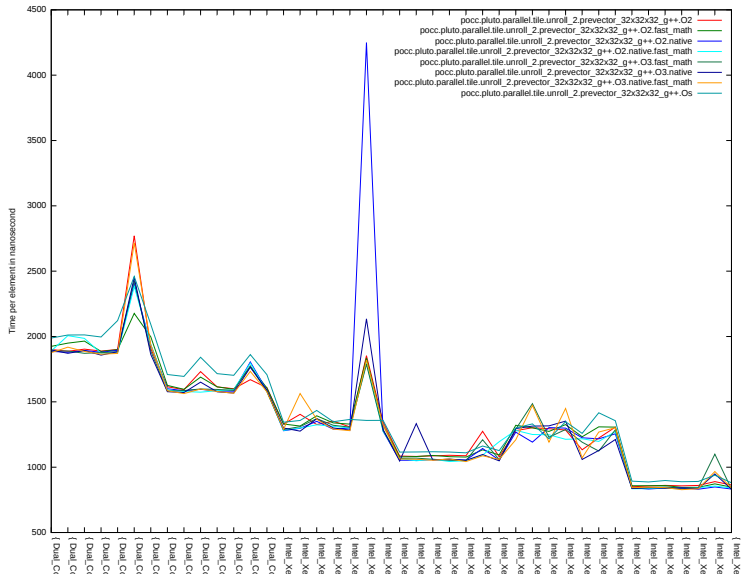
Conclusion

References

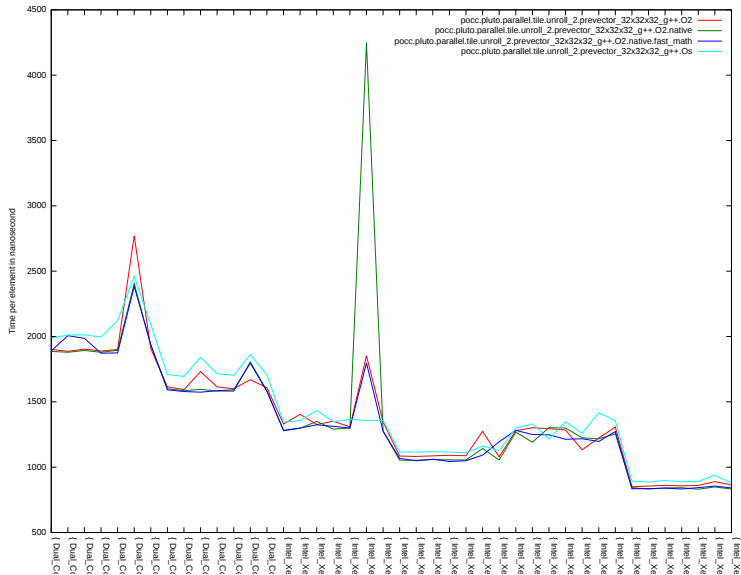
All version of benchmark lu (2142 versions, 42 contexts)



Best versions of benchmark lu (8 versions, 42 contexts)



Minimal versions of benchmark lu (4 versions, 42 contexts)



Lénaïc BAGNÈRES,
Cédric BASTOUL

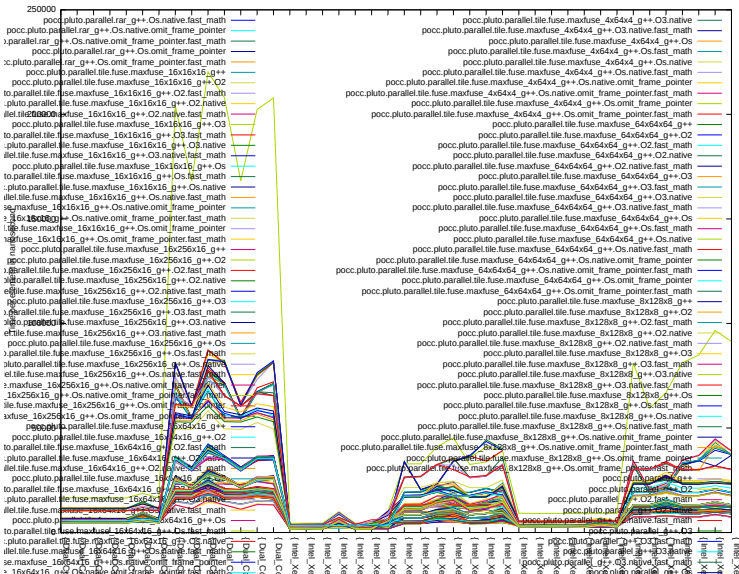
Introduction

Building Switchable
Scheduling

Conclusion

References

All version of benchmark matmul_kj (238 versions, 42 contexts)



Lénaïc BAGNÈRES,
Cédric BASTOUL

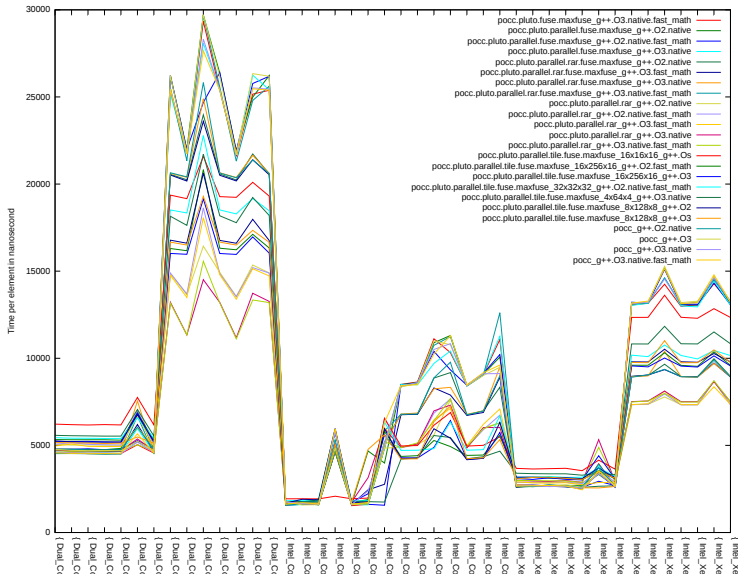
[Introduction](#)

[Building Switchable
Scheduling](#)

[Conclusion](#)

[References](#)

Best versions of benchmark matmul_ijk (24 versions, 42 contexts)



Minimal versions of benchmark matmul_ikj (16 versions, 42 contexts)

