

# BDSC-Based Automatic Task Parallelization: Experiments

Dounia KHALDI

CRI, Mathématiques et systèmes  
MINES ParisTech

Septièmes rencontres de la communauté française de compilation,  
Dammarie-les-Lys, France  
December 04, 2013



- *Anyone can build a fast CPU. The trick is to build a fast system.*  
Attributed to Seymour Cray
- Parallelism handling:
  - Parallel software developed by converting sequential programs by hand
  - Automatic task parallelism extraction: **Scheduling** problem
  - Resource constraints: memory requirements, processor features...
  - Scientific, signal and image processing benchmarks

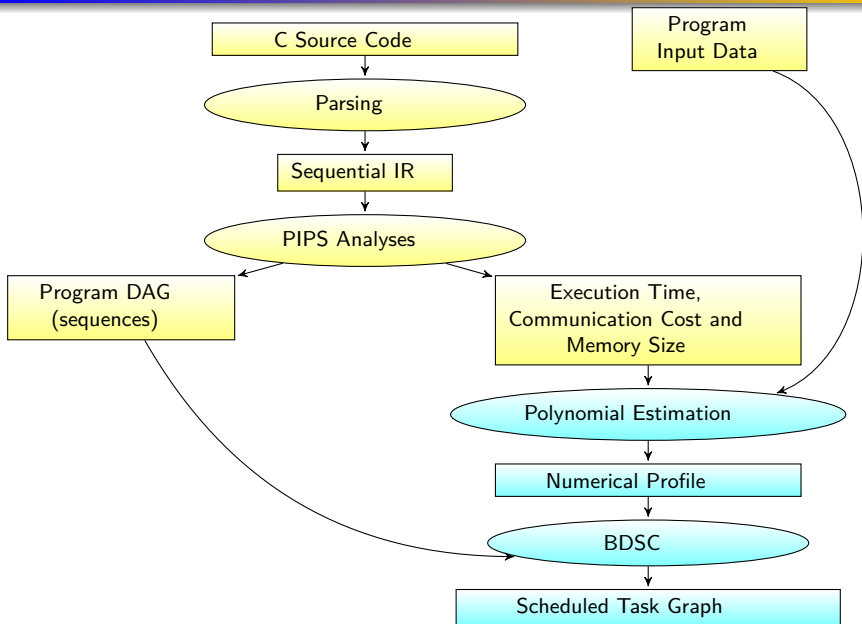
## *Automatic Resource-Constrained Static Task Parallelization*

- **BDSC**: *a memory-constrained, number of processor-bounded extension of DSC*
- *Experimentation on shared and distributed memory systems*

- 1 BDSC: A Memory-Constrained, Number of Processor-Bounded Extension of DSC
- 2 Experimental Evaluations with PIPS
- 3 Conclusion and Future Work

# Parallelization Process


blue indicates contributions; an ellipse, a process; and a rectangle, results



- 1 BDSC: A Memory-Constrained, Number of Processor-Bounded Extension of DSC
  - List-Scheduling Heuristics
  - Dominant Sequence Clustering (DSC)
  - BDSC: A Resource-Constrained Extension of DSC
- 2 Experimental Evaluations with PIPS
- 3 Conclusion and Future Work

# List-Scheduling Heuristics

- Priorities are computed for all unscheduled vertices using:
  - Top level ( $tlevel(\tau)$ ): length of the longest path from entry to  $\tau$
  - Bottom level ( $blevel(\tau)$ ): length of the longest path from  $\tau$  to exit




task	tlevel	blevel
<i>D</i>	0	7
<i>C</i>	3	2
<i>A</i>	0	5
<i>B</i>	4	3

- Vertex  $\tau$  with the highest priority is selected for scheduling
- $\tau$  is added to the cluster (logical process) with the earliest start-time

# A List-Scheduling Heuristic: Dominant Sequence Clustering (DSC)

- DSC (Dominant Sequence Clustering) [Yang and Gerasoulis 1994]
- Task list-scheduling heuristic for an unbounded number of clusters
- $\text{priority}(\tau) = \text{tlevel}(\tau) + \text{blevel}(\tau)$
- $\text{zeroing}(\tau_p, \tau)$  puts  $\tau$  in the cluster of a predecessor  $\tau_p \Rightarrow$  reduces  $\text{tlevel}(\tau)$  by setting to zero the cost of the edge  $(\tau_p, \tau)$



step	task	tlevel	blevel	prio	scheduled tlevel		
					$\kappa_0$	$\kappa_1$	$\kappa_2$
1	D	0	7	7	0*		
2	C	3	2	5	2	3*	
3	A	0	5	5			0*
4	B	4	3	7	2*		4

- Complexity:  $\mathcal{O}(n^2 \log(n))$

# A List-Scheduling Heuristic: Dominant Sequence Clustering (DSC)

- DSC algorithm weaknesses for our purpose:
  - Unbounded number of clusters
  - Number of clusters is not predefined → blind clustering
  - Memory size is not predefined → blind clustering
  - Creates long idle slots in already existing clusters

## Proposal

BDSC: A Memory-Constrained, Number of Processor-Bounded  
Extension of DSC



- ① Memory Constraint Warranty (MCW):
  - Do not exceed a memory threshold  $M$
  - Overapproximation of the amount of memory used in tasks
  - $\text{data\_size}(\text{cluster\_data}(\kappa) \cup \text{task\_data}(\tau)) \leq M$
- ② Bounded number of clusters  $P$ :
  - Number of cluster allocations do not exceed Threshold  $P$
  - Maintain the constraint MCW
  - $\text{argmin}_{\kappa \in \text{clusters}} \text{cluster\_time}(\kappa)$
- ③ Efficient cluster allocation by exploiting idle slots
- ④ Complexity:  $\mathcal{O}(n^3)$

# Related Work: Static Task Parallelization Tools

	blevel	tlevel	Resource constraints	Dependence		Execution time estimation	Communication time estimation	Memory model
				control	data			
BDSC Parallelization	✓	✓	✓		✓	✓	✓	Shared, distributed
Sarkar's work [Sarkar, 1989]	✓				✓	✓	✓	Shared, distributed
OSCAR [Kasahara et al., 1992]		✓		✓	✓	✓		Shared
Pedigree [Newburn and Shen, 1996]	✓	✓		✓	✓	✓		Shared
SPIR [Choi et al., 2009]		✓	✓	✓	✓	✓	✓	Shared

- 1 BDSC: A Memory-Constrained, Number of Processor-Bounded Extension of DSC
- 2 Experimental Evaluations with PIPS
  - Experimental Setting
  - BDSC vs. DSC
  - Comparative Study with Faust Parallelizing Compiler
- 3 Conclusion and Future Work

## ① Benchmarks





- *Thales ABF* (Adaptive Beam Forming), with 1,065 lines
- SPEC benchmark *equake*, with 1,432 lines
- *Harris* corner detector, with 105 lines
- NAS Parallel Benchmark *IS* (Integer Sort), with 1,076 lines



## ② Machines

- **Shared Memory:** host Linux (Ubuntu)  
2-socket AMD quadcore Opteron, 2.4 GHz  
 $M = 16\text{GB}$  of RAM  
gcc 4.6.3 -O3  
OpenMP 3.0  
Cluster  $\sim$  Thread
- **Distributed Memory:** host Linux (RedHat)  
6 dual-core processors Intel® Xeon®, 2.5 GHz  
 $M = 32\text{GB}$  of RAM per processor  
gcc 4.4.6 -O3  
Open MPI 1.6.2  
Cluster  $\sim$  Process

# ABF and equake



## Speedups with OpenMP and MPI







# NAS Parallel Benchmark IS

## Speedups with OpenMP and MPI: Different Class Sizes



# Faust Parallel Scheduling vs. BDSC

- Faust (Functional AUdio STream) [Orlarey et al., 2009]
- DSL for real-time audio signal processing and synthesis
- Generation of C or C++ with or without OpenMP directives
- `omp task` (BDSC) vs. `omp section` (Faust Parallelizing Compiler)
- Scheduling: BDSC vs. Faust topological ordering
- Speedups for two programs: Karplus32 and Freeverb





- 1 BDSC: A Memory-Constrained, Number of Processor-Bounded Extension of DSC
- 2 Experimental Evaluations with PIPS
- 3 Conclusion and Future Work

## ① BDSC-based hierarchical scheduling algorithm

- Memory constraint, bounded number of clusters, efficient cluster allocation
- BDSC-based task parallelization algorithm
- Communication, data and time cost models

## ② Experiments:

- BDSC-based automatic parallelization in PIPS
- Code generation in OpenMP and MPI
- Good speedups for coarse-grained parallelism

## ① BDSC Scheduling

- Handling of heterogeneous devices
- More precise cost models

## ② Parallel Code Generation

- More experimentation needed
- Solving communication generation problems (MPI)
- Hybrid task + data parallelism

# BDSC-Based Automatic Task Parallelization: Experiments

Dounia KHALDI

CRI, Mathématiques et systèmes  
MINES ParisTech

Septièmes rencontres de la communauté française de compilation,  
Dammarie-les-Lys, France  
December 04, 2013



# References I



Choi, Y., Lin, Y., Chong, N., Mahlke, S., and Mudge, T. (2009).

Stream Compilation for Real-Time Embedded Multicore Systems.

In *Proceedings of the 7th annual IEEE/ACM International Symposium on Code Generation and Optimization, CGO '09*, pages 210–220, Washington, DC, USA.



Kasahara, H., Honda, H., Mogi, A., Ogura, A., Fujiwara, K., and Narita, S. (1992).

A Multi-Grain Parallelizing Compilation Scheme for OSCAR (Optimally Scheduled Advanced Multiprocessor).

In *Proceedings of the Fourth International Workshop on Languages and Compilers for Parallel Computing*, pages 283–297, London, UK. Springer-Verlag.



Nandivada, V. K., Shirako, J., Zhao, J., and Sarkar, V. (2013).

A Transformation Framework for Optimizing Task-Parallel Programs.

*ACM Trans. Program. Lang. Syst.*, 35(1):3:1–3:48.



Newburn, C. J. and Shen, J. P. (1996).

Automatic Partitioning of Signal Processing Programs for Symmetric Multiprocessors.

In *IEEE PACT*, pages 269–280. IEEE Computer Society.



Sarkar, V. (1989).

*Partitioning and Scheduling Parallel Programs for Multiprocessors*.

MIT Press, Cambridge, MA, USA.