・ ロ ト ・ 雪 ト ・ 目 ト ・ 日 ト

Accelerating a local search algorithm for large instances of the independent task scheduling problem with the GPU

Frederic Pinel¹ Johnatan Pecero¹ Pascal Bouvry¹ Samee U. Khan²

> ¹Computer Science and Communications University of Luxembourg

²Electrical and Computer Engineering North Dakota State University

International Research Workshop on Advanced High Performance Computing Systems



э

Adaptation

Large inst

・ロン ・ 理 と ・ 理 と ・

Perspectives

Outline

Millicomputing

Independent task mapping

Initial algorithm

Adaptation

Large instances

Perspectives



э

Millicomputing

Millicomputing

- ARM cortex-A9 based, Panda boards (OMAP4430)
- Compare with data center server



Figure: Canonical ARM build farm





- Low energy (5-20W ?)
- Limited cooling
- Reduced contention
- Better programs
- SLA



・ロン ・ 理 と ・ 理 と ・

・ロット (雪) ・ (日) ・ (日)

Problems

- Programming model
 - Message passing (Erlang)
 - Pipelining (message queues)
 - Independent tasks
- Mapping (spawn multiple task instances)
- Task characterization
- Networking



Perspectives

Problem definition

- Mapping independent tasks on a distributed system
- Hypothesis: estimated time to complete (ETC)
- Makespan, combines several perspectives:
 - User: flowtime
 - Provider: load balance, energy (low machine heterogeneity)
- Fast algorithm: runtime, wallclock



э

・ ロ ト ・ 雪 ト ・ 目 ト ・ 日 ト

・ロン ・ 理 と ・ 理 と ・

Contribution

- Mapping heuristic
- · Parallel version for large instances on the GPU
- Lessons learned



Perspec

Background



Figure: Cellular genetic algorithm (CGA)



・ロト ・ 聞 ト ・ ヨ ト ・ ヨ ト

Perspective

Parallel CGA

- SIMD
- Clusters of sequential machines
- Multicore machines
- Communication models



Figure: Population decomposition



Starting point

- Parallel asynchronous cellular genetic algorithm
- Initialized with heuristic (Min-Min)
- With local search (5 iter)
- Problems: 512 tasks × 16 cores

・ ロ マ ・ 雪 マ ・ 雪 マ ・ 日 マ





Adaptation

- Simplified algorithm
- Min-Min, incremental formulation
- Local search (30 iter), complete-state formulation



э

・ロット (雪) (日) (日)

ces Per

Results



Figure: Consistent, high-h. tasks, low-h. machines



・ロト ・ 聞 ト ・ ヨ ト ・ ヨ ト

n Large

(日)

es Pers

Perspectives

Results



Figure: Semi-consistent, high-h. tasks, low-h. machines



Large in:

(日)

Perspectives

Results



Figure: Consistent, low-h. tasks, low-h. machines



(日)

Results



Figure: Semi-consistent, low-h. tasks, low-h. machines



Min-Min runtime



Figure: Runtime Min-Min (minimum in array)



Perspective

Min-Min on GPU



Figure: Parallel reduction in Min-Min



◆□ ▶ ◆□ ▶ ◆□ ▶ ◆□ ▶ ●

Min-Min runtime



Figure: Runtime Min-Min



・ロト ・ 聞 ト ・ ヨ ト ・ ヨ ト

Quality of solution



Figure: Makespan



・ロト ・聞ト ・ヨト ・ヨト

Performance







э

Perspectives

- Failure
- Solution \rightarrow Feedback \rightarrow Loop
- Learning process



æ

・ロン ・ 理 と ・ 理 と ・

Large insta

・ ロ ト ・ 雪 ト ・ 目 ト ・ 日 ト

Perspectives

Machine learning opportunities

- Learn from problem instance
 - Classify problem instance
 - Task characterization
- Learn mapping rules
 - Adapt (parameters, heuristics)
 - Algorithm (oracle: solved instances)
 - Data placement
 - Networking



э

Adaptation

・ロト ・聞 と ・ ヨ と ・ ヨ と

Perspectives



Thank you.



æ