

Mapping pipeline skeletons onto heterogeneous platforms

Anne Benoit and Yves Robert

GRAAL team, LIP
École Normale Supérieure de Lyon

January 2007

Introduction and motivation

- Mapping applications onto parallel platforms
Difficult challenge
- Heterogeneous clusters, fully heterogeneous platforms
Even more difficult!
- Structured programming approach
 - Easier to program (deadlocks, process starvation)
 - Range of well-known paradigms (pipeline, farm)
 - Algorithmic skeleton: help for mapping

Mapping pipeline skeletons onto heterogeneous platforms

Introduction and motivation

- Mapping applications onto parallel platforms
Difficult challenge
- Heterogeneous clusters, fully heterogeneous platforms
Even more difficult!
- Structured programming approach
 - Easier to program (deadlocks, process starvation)
 - Range of well-known paradigms (pipeline, farm)
 - Algorithmic skeleton: help for mapping

Mapping pipeline skeletons onto heterogeneous platforms

Introduction and motivation

- Mapping applications onto parallel platforms
Difficult challenge
- Heterogeneous clusters, fully heterogeneous platforms
Even more difficult!
- Structured programming approach
 - Easier to program (deadlocks, process starvation)
 - Range of well-known paradigms (pipeline, farm)
 - Algorithmic skeleton: help for mapping

Mapping pipeline skeletons onto heterogeneous platforms

Introduction and motivation

- Mapping applications onto parallel platforms
Difficult challenge
- Heterogeneous clusters, fully heterogeneous platforms
Even more difficult!
- Structured programming approach
 - Easier to program (deadlocks, process starvation)
 - Range of well-known paradigms (pipeline, farm)
 - Algorithmic skeleton: help for mapping

Mapping pipeline skeletons onto heterogeneous platforms

Introduction and motivation

- Mapping applications onto parallel platforms
Difficult challenge
- Heterogeneous clusters, fully heterogeneous platforms
Even more difficult!
- Structured programming approach
 - Easier to program (deadlocks, process starvation)
 - Range of well-known paradigms (pipeline, farm)
 - Algorithmic skeleton: help for mapping

Mapping pipeline skeletons onto heterogeneous platforms

Introduction and motivation

- Mapping applications onto parallel platforms
Difficult challenge
- Heterogeneous clusters, fully heterogeneous platforms
Even more difficult!
- Structured programming approach
 - Easier to program (deadlocks, process starvation)
 - Range of well-known paradigms (pipeline, farm)
 - Algorithmic skeleton: help for mapping

Mapping pipeline skeletons onto heterogeneous platforms

Introduction and motivation

- Mapping applications onto parallel platforms
Difficult challenge
- Heterogeneous clusters, fully heterogeneous platforms
Even more difficult!
- Structured programming approach
 - Easier to program (deadlocks, process starvation)
 - Range of well-known paradigms (pipeline, farm)
 - Algorithmic skeleton: help for mapping

Mapping pipeline skeletons onto heterogeneous platforms

Why restrict to pipelines?

- Chains-on-chains partitioning problem
 - no communications
 - identical processors
- Extensions (**done**)
 - with communications
 - with heterogeneous processors/links
 - goal: assess complexity, design heuristics
- Extensions (**current work**)
 - deal with DEALs
 - deal with DAGs

Why restrict to pipelines?

- Chains-on-chains partitioning problem
 - no communications
 - identical processors
- Extensions (**done**)
 - with communications
 - with heterogeneous processors/links
 - goal: assess complexity, design heuristics
- Extensions (**current work**)
 - deal with DEALs
 - deal with DAGs

Why restrict to pipelines?

- Chains-on-chains partitioning problem
 - no communications
 - identical processors
- Extensions (**done**)
 - with communications
 - with heterogeneous processors/links
 - goal: assess complexity, design heuristics
- Extensions (**current work**)
 - deal with DEALs
 - deal with DAGs

Chains-on-chains

Load-balance contiguous tasks

5 7 3 4 8 1 3 8 2 9 7 3 5 2 3 6

- Back to Bokhari and Iqbal partitioning papers
- See survey by Pinar and Aykanat, JPDC 64, 8 (2004)

Chains-on-chains

Load-balance contiguous tasks

5 7 3 4 8 1 3 8 2 9 7 3 5 2 3 6

With $p = 4$ processors?

- Back to Bokhari and Iqbal partitioning papers
- See survey by Pinar and Aykanat, JPDC 64, 8 (2004)

Chains-on-chains

Load-balance contiguous tasks

5 7 3 4 8 1 3 8 2 9 7 3 5 2 3 6

With $p = 4$ processors?

5 7 3 4 | 8 1 3 8 | 2 9 7 | 3 5 2 3 6

- Back to Bokhari and Iqbal partitioning papers
- See survey by Pinar and Aykanat, JPDC 64, 8 (2004)

Chains-on-chains

Load-balance contiguous tasks

5 7 3 4 8 1 3 8 2 9 7 3 5 2 3 6

With $p = 4$ processors?

5 7 3 4 | 8 1 3 8 | 2 9 7 | 3 5 2 3 6

- Back to Bokhari and Iqbal partitioning papers
- See survey by Pinar and Aykanat, JPDC 64, 8 (2004)


Rule of the game

- Map each pipeline stage on a single processor (no deals)
- Goal: minimize execution time
- Several mapping strategies



Rule of the game

- Map each pipeline stage on a single processor (no deals)
- Goal: minimize execution time
- Several mapping strategies



The pipeline application

Rule of the game

- Map each pipeline stage on a single processor (no deals)
- Goal: minimize execution time
- Several mapping strategies



ONE-TO-ONE MAPPING

Rule of the game

- Map each pipeline stage on a single processor (no deals)
- Goal: minimize execution time
- Several mapping strategies



Rule of the game

- Map each pipeline stage on a single processor (no deals)
- Goal: minimize execution time
- Several mapping strategies



GENERAL MAPPING

Major contributions

Theory Formal approach to the problem
Problem complexity
Integer linear program for exact resolution

Practice Heuristics for INTERVAL MAPPING on clusters
Experiments to compare heuristics and evaluate their absolute performance

Major contributions


Theory Formal approach to the problem
Problem complexity
Integer linear program for exact resolution

Practice Heuristics for INTERVAL MAPPING on clusters
Experiments to compare heuristics and evaluate their absolute performance

Outline


- 1 Framework
- 2 Complexity results
- 3 Heuristics
- 4 Experiments
- 5 Linear programming formulation
- 6 Conclusion

The application



- n stages S_k , $1 \leq k \leq n$
- S_k :
 - receives input of size δ_{k-1} from S_{k-1}
 - performs w_k computations
 - outputs data of size δ_k to S_{k+1}
- S_0 and S_{n+1} : virtual stages representing the outside world

The platform



- p processors P_u , $1 \leq u \leq p$, fully interconnected
- s_u : speed of processor P_u
- bidirectional link $link_{u,v} : P_u \rightarrow P_v$, bandwidth $b_{u,v}$
- **one-port** model: each processor can either send, receive or compute at any time-step
- P_{in} : input data – P_{out} : output data

Different platforms


Fully Homogeneous – Identical processors ($s_u = s$) and links ($b_{u,v} = b$): typical parallel machines

Communication Homogeneous – Different-speed processors ($s_u \neq s_v$), identical links ($b_{u,v} = b$): networks of workstations, clusters

Fully Heterogeneous – Fully heterogeneous architectures, $s_u \neq s_v$ and $b_{u,v} \neq b_{u',v'}$: hierarchical platforms, grids

Mapping problem: ONE-TO-ONE MAPPING

- $n \leq p$: map each stage S_k onto a distinct processor $P_{\text{alloc}(k)}$
- Period of $P_{\text{alloc}(k)}$: minimum delay between processing of two consecutive tasks




- Cycle-time of P_u : $\text{cycle}_u = \frac{\delta_{k-1}}{b_{t,u}} + \frac{w_k}{s_u} + \frac{\delta_k}{b_{u,v}}$
- Optimization problem: find the allocation function $\text{alloc} : [1, n] \rightarrow [1, p]$ which minimizes

$$T_{\text{period}} = \max_{1 \leq k \leq n} \text{cycle}_{\text{alloc}(k)}$$

(with $\text{alloc}(0) = \text{in}$ and $\text{alloc}(n+1) = \text{out}$)

Mapping problem: ONE-TO-ONE MAPPING

- $n \leq p$: map each stage S_k onto a distinct processor $P_{\text{alloc}(k)}$
- Period of $P_{\text{alloc}(k)}$: minimum delay between processing of two consecutive tasks



- Cycle-time of P_u : $\text{cycle}_u = \frac{\delta_{k-1}}{b_{t,u}} + \frac{w_k}{s_u} + \frac{\delta_k}{b_{u,v}}$
- Optimization problem: find the allocation function $\text{alloc} : [1, n] \rightarrow [1, p]$ which minimizes

$$T_{\text{period}} = \max_{1 \leq k \leq n} \text{cycle}_{\text{alloc}(k)}$$

(with $\text{alloc}(0) = \text{in}$ and $\text{alloc}(n+1) = \text{out}$)

Mapping problem: INTERVAL MAPPING

- Several consecutive stages onto the same processor
- Increase computational load, reduce communications
- Mandatory when $p < n$
- Partition of $[1..n]$ into m intervals $I_j = [d_j, e_j]$
(with $d_j \leq e_j$ for $1 \leq j \leq m$, $d_1 = 1$, $d_{j+1} = e_j + 1$ for $1 \leq j \leq m - 1$ and $e_m = n$)
- Interval I_j mapped onto processor $P_{\text{alloc}(j)}$

$$T_{\text{period}} = \max_{1 \leq j \leq m} \left\{ \frac{\delta_{d_{j-1}}}{b_{\text{alloc}(j-1), \text{alloc}(j)}} + \frac{\sum_{i=d_j}^{e_j} w_i}{s_{\text{alloc}(j)}} + \frac{\delta_{e_j}}{b_{\text{alloc}(j), \text{alloc}(j+1)}} \right\}$$

Mapping problem: INTERVAL MAPPING

- Several consecutive stages onto the same processor
- Increase computational load, reduce communications
- Mandatory when $p < n$
- Partition of $[1..n]$ into m intervals $I_j = [d_j, e_j]$
(with $d_j \leq e_j$ for $1 \leq j \leq m$, $d_1 = 1$, $d_{j+1} = e_j + 1$ for $1 \leq j \leq m - 1$ and $e_m = n$)
- Interval I_j mapped onto processor $P_{\text{alloc}(j)}$

$$T_{\text{period}} = \max_{1 \leq j \leq m} \left\{ \frac{\delta_{d_{j-1}}}{b_{\text{alloc}(j-1), \text{alloc}(j)}} + \frac{\sum_{i=d_j}^{e_j} w_i}{s_{\text{alloc}(j)}} + \frac{\delta_{e_j}}{b_{\text{alloc}(j), \text{alloc}(j+1)}} \right\}$$

Mapping problem: INTERVAL MAPPING

- Several consecutive stages onto the same processor
- Increase computational load, reduce communications
- Mandatory when $p < n$
- Partition of $[1..n]$ into m intervals $I_j = [d_j, e_j]$
(with $d_j \leq e_j$ for $1 \leq j \leq m$, $d_1 = 1$, $d_{j+1} = e_j + 1$ for $1 \leq j \leq m - 1$ and $e_m = n$)
- Interval I_j mapped onto processor $P_{\text{alloc}(j)}$

$$T_{\text{period}} = \max_{1 \leq j \leq m} \left\{ \frac{\delta_{d_{j-1}}}{b_{\text{alloc}(j-1), \text{alloc}(j)}} + \frac{\sum_{i=d_j}^{e_j} w_i}{s_{\text{alloc}(j)}} + \frac{\delta_{e_j}}{b_{\text{alloc}(j), \text{alloc}(j+1)}} \right\}$$

Mapping problem: GENERAL MAPPING

- Not suiting the one-port model very well: can always be replaced by an INTERVAL MAPPING as good as the general one for *Communication Homogeneous* platforms
- Can be the optimal mapping for *Fully Heterogeneous* platforms in some particular cases
- More general, but requires threads and may lead to idle times and races with the one-port model

Outline

- 1 Framework
- 2 Complexity results**
- 3 Heuristics
- 4 Experiments
- 5 Linear programming formulation
- 6 Conclusion

Complexity results

	Fully Hom.	Comm. Hom.
One-to-one Mapping		
Interval Mapping		
General Mapping		

-
-
-
-

Complexity results

	Fully Hom.	Comm. Hom.
One-to-one Mapping	polynomial	polynomial
Interval Mapping		
General Mapping		

- Binary search polynomial algorithm for ONE-TO-ONE MAPPING
-
-
-

Complexity results

	Fully Hom.	Comm. Hom.
One-to-one Mapping	polynomial	polynomial
Interval Mapping	polynomial	NP-complete
General Mapping		

- Binary search polynomial algorithm for ONE-TO-ONE MAPPING
- Dynamic programming algorithm for INTERVAL MAPPING on Hom. platforms
-
-

Complexity results

	Fully Hom.	Comm. Hom.
One-to-one Mapping	polynomial	polynomial
Interval Mapping	polynomial	NP-complete
General Mapping	same complexity as Interval	

- Binary search polynomial algorithm for ONE-TO-ONE MAPPING
- Dynamic programming algorithm for INTERVAL MAPPING on Hom. platforms
- General mapping: same complexity as INTERVAL MAPPING
-

Complexity results

	Fully Hom.	Comm. Hom.
One-to-one Mapping	polynomial	polynomial
Interval Mapping	polynomial	NP-complete
General Mapping	same complexity as Interval	

- Binary search polynomial algorithm for **ONE-TO-ONE MAPPING**
- Dynamic programming algorithm for **INTERVAL MAPPING** on Hom. platforms
- General mapping: same complexity as **INTERVAL MAPPING**
- All problem instances NP-complete on *Fully Heterogeneous* platforms

Back to chains-on-chains

- Chains-on-chains + **homogeneous communications**:
polynomial 😊
- Chains-on-chains + **different-speed processors**:
NP-complete 😞

Back to chains-on-chains

- Chains-on-chains + **homogeneous communications**:
polynomial 😊
- Chains-on-chains + **different-speed processors**:
NP-complete 😞

One-to-one/Comm. Hom.: binary search algorithm

- Work with fastest n processors, numbered P_1 to P_n , where $s_1 \leq s_2 \leq \dots \leq s_n$
- Mark all stages \mathcal{S}_1 to \mathcal{S}_n as free
- **For** $u = 1$ **to** n
 - Pick up any free stage \mathcal{S}_k s.t. $\delta_{k-1}/b + w_k/s_u + \delta_k/b \leq T_{\text{period}}$
 - Assign \mathcal{S}_k to P_u , and mark \mathcal{S}_k as already assigned
 - If no stage found return "failure"
- **Proof:** exchange argument

One-to-one/Comm. Hom.: binary search algorithm

- Work with fastest n processors, numbered P_1 to P_n , where $s_1 \leq s_2 \leq \dots \leq s_n$
- Mark all stages \mathcal{S}_1 to \mathcal{S}_n as free
- **For** $u = 1$ **to** n
 - Pick up any free stage \mathcal{S}_k s.t. $\delta_{k-1}/b + w_k/s_u + \delta_k/b \leq T_{\text{period}}$
 - Assign \mathcal{S}_k to P_u , and mark \mathcal{S}_k as already assigned
 - If no stage found return "failure"
- **Proof:** exchange argument

Interval, Fully Hom.: dynamic programming algorithm

- $c(i, j, k)$: optimal period to map stages S_i to S_j using exactly k processors
- Goal: $\min_{1 \leq k \leq p} c(1, n, k)$

$$c(i, j, k) = \min_{\substack{q+r=k \\ 1 \leq q \leq k-1 \\ 1 \leq r \leq k-1}} \left\{ \min_{i \leq \ell \leq j-1} \{ \max(c(i, \ell, q), c(\ell+1, j, r)) \} \right\}$$

$$c(i, j, 1) = \frac{\delta_{i-1}}{b} + \frac{\sum_{k=i}^j w_k}{s} + \frac{\delta_j}{b}$$

$$c(i, j, k) = +\infty \text{ if } k > j - i + 1$$

- **Proof:** search over all possible partitionings into two subintervals, using every possible number of processors for each interval
- Complexity: $O(n^3 p^2)$

Interval, Fully Hom.: dynamic programming algorithm

- $c(i, j, k)$: optimal period to map stages S_i to S_j using exactly k processors
- Goal: $\min_{1 \leq k \leq p} c(1, n, k)$

$$c(i, j, k) = \min_{\substack{q+r=k \\ 1 \leq q \leq k-1 \\ 1 \leq r \leq k-1}} \left\{ \min_{i \leq \ell \leq j-1} \{ \max(c(i, \ell, q), c(\ell+1, j, r)) \} \right\}$$

$$c(i, j, 1) = \frac{\delta_{i-1}}{b} + \frac{\sum_{k=i}^j w_k}{s} + \frac{\delta_j}{b}$$

$$c(i, j, k) = +\infty \text{ if } k > j - i + 1$$

- **Proof:** search over all possible partitionings into two subintervals, using every possible number of processors for each interval
- Complexity: $O(n^3 p^2)$

Interval, Fully Hom.: dynamic programming algorithm

- $c(i, j, k)$: optimal period to map stages S_i to S_j using exactly k processors
- Goal: $\min_{1 \leq k \leq p} c(1, n, k)$

$$c(i, j, k) = \min_{\substack{q+r=k \\ 1 \leq q \leq k-1 \\ 1 \leq r \leq k-1}} \left\{ \min_{i \leq \ell \leq j-1} \{ \max(c(i, \ell, q), c(\ell+1, j, r)) \} \right\}$$

$$c(i, j, 1) = \frac{\delta_{i-1}}{b} + \frac{\sum_{k=i}^j w_k}{s} + \frac{\delta_j}{b}$$

$$c(i, j, k) = +\infty \text{ if } k > j - i + 1$$

- **Proof:** search over all possible partitionings into two subintervals, using every possible number of processors for each interval
- Complexity: $O(n^3 p^2)$

Interval, Fully Hom.: dynamic programming algorithm

- $c(i, j, k)$: optimal period to map stages S_i to S_j using exactly k processors
- Goal: $\min_{1 \leq k \leq p} c(1, n, k)$

$$c(i, j, k) = \min_{\substack{q+r=k \\ 1 \leq q \leq k-1 \\ 1 \leq r \leq k-1}} \left\{ \min_{i \leq \ell \leq j-1} \{ \max(c(i, \ell, q), c(\ell+1, j, r)) \} \right\}$$

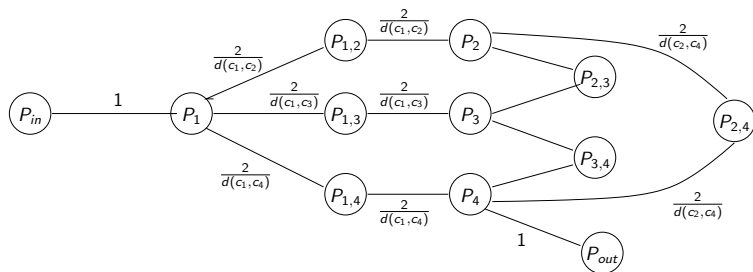
$$c(i, j, 1) = \frac{\delta_{i-1}}{b} + \frac{\sum_{k=i}^j w_k}{s} + \frac{\delta_j}{b}$$

$$c(i, j, k) = +\infty \text{ if } k > j - i + 1$$

- **Proof:** search over all possible partitionings into two subintervals, using every possible number of processors for each interval
- Complexity: $O(n^3 p^2)$

Heterogeneous platforms: NP-complete

- Reduction from MINIMUM METRIC BOTTLENECK WANDERING SALESPERSON PROBLEM



Outline

- 1 Framework
- 2 Complexity results
- 3 Heuristics**
- 4 Experiments
- 5 Linear programming formulation
- 6 Conclusion

Greedy heuristics (1/2)

- Target clusters: *Communication Homogeneous* platforms and INTERVAL MAPPING
- $L = \lceil n/p \rceil$ consecutive stages per processor: set of intervals fixed
- $\lceil n/L \rceil$ processors used
- ONE-TO-ONE MAPPING when $n \leq p$ ($L = 1$)
- Rule applied in all greedy heuristics except random interval length

Greedy heuristics (2/2)

H1a-GR: **random** – Random choice of a free processor for each interval

H1b-GRIL: **random interval length** – Idem with random interval sizes: average length L , $1 \leq \text{length} \leq 2L - 1$

H2-GSW: **biggest $\sum w$** – Place interval with most computations on fastest processor

H3-GSD: **biggest $\delta_{in} + \delta_{out}$** – Intervals are sorted by communications ($\delta_{in} + \delta_{out}$)
in: first stage of interval; *out* – 1: last one

H4-GP: **biggest period on fastest processor** – Balancing computation and communication: processors sorted by decreasing speed s_u ; for current processor u , choose interval with biggest period
 $(\delta_{in} + \delta_{out})/b + \sum_{i \in \text{Interval}} w_i/s_u$

Sophisticated heuristics

- H5-BS121: binary search for ONE-TO-ONE MAPPING** – optimal algorithm for ONE-TO-ONE MAPPING. When $p < n$, application cut in fixed intervals of length L .
- H6-SPL: splitting intervals** – Processors sorted by decreasing speed, all stages to first processor. At each step, select used proc j with largest period, split its interval (give fraction of stages to j'): minimize $\max(\text{period}(j), \text{period}(j'))$ and split if maximum period improved.
- H7a-BSL and H7b-BSC: binary search (longest/closest)** – Binary search on period P : start with stage $s = 1$, build intervals (s, s') fitting on processors. For each u , and each $s' \geq s$, compute period $(s..s', u)$ and check whether it is smaller than P . **H7a**: maximizes s' ; **H7b**: chooses the closest period.

Outline

- 1 Framework
- 2 Complexity results
- 3 Heuristics
- 4 Experiments**
- 5 Linear programming formulation
- 6 Conclusion

Plan of experiments

- Assess performance of **polynomial heuristics**
- Random applications, $n = 1$ to 50 stages
- Random platforms, $p = 10$ and $p = 100$ processors
- $b = 10$ (comm. hom.), proc. speed between 1 and 20
- Relevant parameters: ratios $\frac{\delta}{b}$ and $\frac{w}{s}$
- Average over 100 similar random appli/platform pairs

Plan of experiments

- Assess performance of **polynomial heuristics**
- Random applications, $n = 1$ to 50 stages
- Random platforms, $p = 10$ and $p = 100$ processors
- $b = 10$ (comm. hom.), proc. speed between 1 and 20
- Relevant parameters: ratios $\frac{\delta}{b}$ and $\frac{w}{s}$
- Average over 100 similar random appli/platform pairs

Plan of experiments

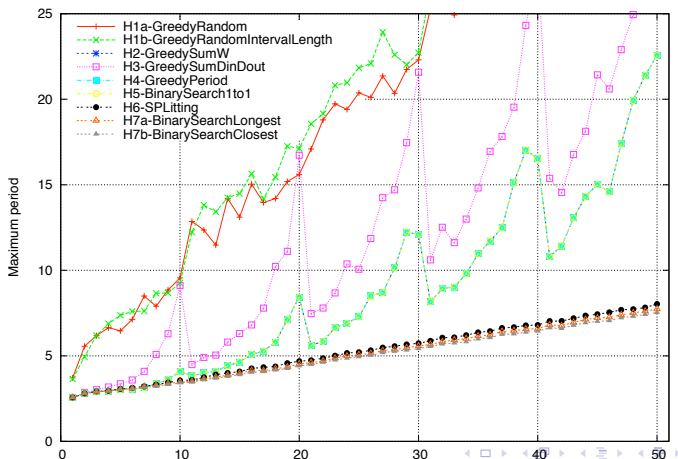
- Assess performance of **polynomial heuristics**
- Random applications, $n = 1$ to 50 stages
- Random platforms, $p = 10$ and $p = 100$ processors
- $b = 10$ (comm. hom.), proc. speed between 1 and 20
- Relevant parameters: ratios $\frac{\delta}{b}$ and $\frac{w}{s}$
- Average over 100 similar random appli/platform pairs

Experiment 1 - balanced comm/comp, hom comm

- $\delta_i = 10$, computation time between 1 and 20
- 10 processors
- 100 processors

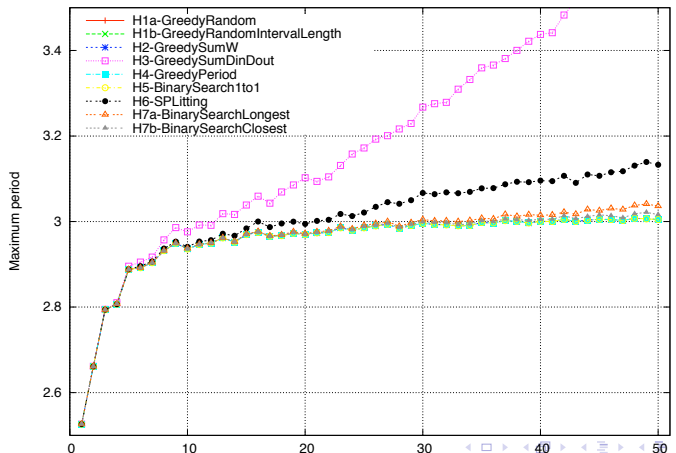
Experiment 1 - balanced comm/comp, hom comm

- $\delta_i = 10$, computation time between 1 and 20
- 10 processors
- 100 processors



Experiment 1 - balanced comm/comp, hom comm

- $\delta_i = 10$, computation time between 1 and 20
- 10 processors
- 100 processors




Experiment 2 - balanced comm/comp, het comm

- communication time between 1 and 100
- computation time between 1 and 20


Experiment 2 - balanced comm/comp, het comm

- communication time between 1 and 100
- computation time between 1 and 20



Experiment 2 - balanced comm/comp, het comm

- communication time between 1 and 100
- computation time between 1 and 20




Experiment 3 - large computations

- communication time between 1 and 20
- computation time between 10 and 1000

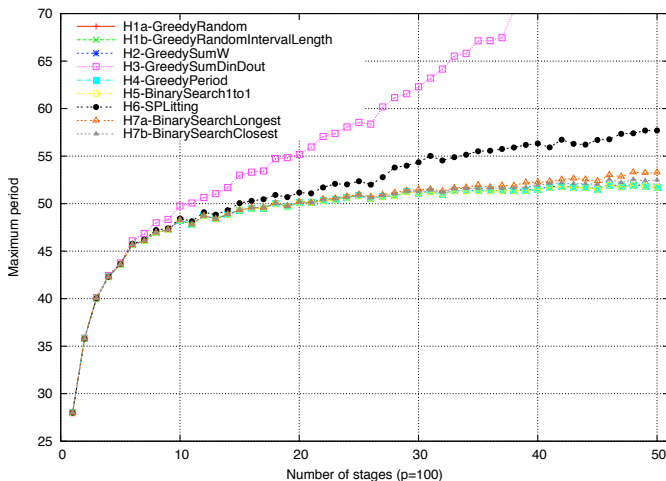
Experiment 3 - large computations

- communication time between 1 and 20
- computation time between 10 and 1000



Experiment 3 - large computations

- communication time between 1 and 20
- computation time between 10 and 1000




Experiment 4 - computations

- communication time between 1 and 20
- computation time between 0.01 and 10

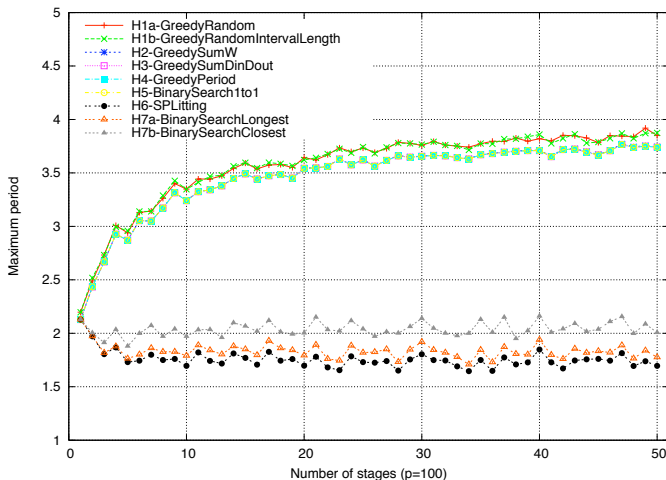
Experiment 4 - computations

- communication time between 1 and 20
- computation time between 0.01 and 10



Experiment 4 - computations

- communication time between 1 and 20
- computation time between 0.01 and 10



Summary of experiments

- Much more efficient than random mappings
- Three dominant heuristics for different cases
- Insignificant communications (*hom. or small*) and *many* processors: H5-BS121 (ONE-TO-ONE MAPPING)
- Insignificant communications (*hom. or small*) and *few* processors: H7b-BSC (clever choice where to split)
- Important communications (*het. or big*): H6-SPL (splitting choice relevant for any number of processors)

Summary of experiments

- Much more efficient than random mappings
- Three dominant heuristics for different cases
- Insignificant communications (**hom. or small**) and **many** processors: H5-BS121 (ONE-TO-ONE MAPPING)
- Insignificant communications (**hom. or small**) and **few** processors: H7b-BSC (clever choice where to split)
- Important communications (**het. or big**): H6-SPL (splitting choice relevant for any number of processors)

Outline

- 1 Framework
- 2 Complexity results
- 3 Heuristics
- 4 Experiments
- 5 Linear programming formulation**
- 6 Conclusion

Integer linear programming

- Integer LP to solve INTERVAL MAPPING on *Fully Heterogeneous* platforms
- Many integer variables: no efficient algorithm to solve
- Approach limited to small problem instances
- Absolute performance of the heuristics for such instances

Linear program: variables

- $x_{k,u}$: 1 if S_k on P_u (0 otherwise)
- $y_{k,u}$: 1 if S_k and S_{k+1} both on P_u (0 otherwise)
- $z_{k,u,v}$: 1 if S_k on P_u and S_{k+1} on P_v (0 otherwise)
- $first_u$ and $last_u$: integer denoting first and last stage assigned to P_u (to enforce interval constraints)
- T_{period} : period of the pipeline
- Objective function: minimize T_{period}

Linear program: variables

- $x_{k,u}$: 1 if S_k on P_u (0 otherwise)
- $y_{k,u}$: 1 if S_k and S_{k+1} both on P_u (0 otherwise)
- $z_{k,u,v}$: 1 if S_k on P_u and S_{k+1} on P_v (0 otherwise)
- $first_u$ and $last_u$: integer denoting first and last stage assigned to P_u (to enforce interval constraints)
- T_{period} : period of the pipeline
- Objective function: minimize T_{period}

Linear program: variables

- $x_{k,u}$: 1 if S_k on P_u (0 otherwise)
- $y_{k,u}$: 1 if S_k and S_{k+1} both on P_u (0 otherwise)
- $z_{k,u,v}$: 1 if S_k on P_u and S_{k+1} on P_v (0 otherwise)
- $first_u$ and $last_u$: integer denoting first and last stage assigned to P_u (to enforce interval constraints)
- T_{period} : period of the pipeline
- **Objective function: minimize T_{period}**

Linear program: constraints

- $\forall k \in [0..n + 1], \quad \sum_u x_{k,u} = 1$
- $\forall k \in [0..n], \quad \sum_{u \neq v} z_{k,u,v} + \sum_u y_{k,u} = 1$
- $\forall k \in [0..n], \forall u, v \in [1..p] \cup \{in, out\}, u \neq v, x_{k,u} + x_{k+1,v} \leq 1 + z_{k,u,v}$
- $\forall k \in [0..n], \forall u \in [1..p] \cup \{in, out\}, \quad x_{k,u} + x_{k+1,u} \leq 1 + y_{k,u}$
- $\forall k \in [1..n], \forall u \in [1..p], \quad first_u \leq k.x_{k,u} + n.(1 - x_{k,u})$
- $\forall k \in [1..n], \forall u \in [1..p], \quad last_u \geq k.x_{k,u}$
- $\forall k \in [1..n - 1], \forall u, v \in [1..p], u \neq v,$
 $last_u \leq k.z_{k,u,v} + n.(1 - z_{k,u,v})$
- $\forall k \in [1..n - 1], \forall u, v \in [1..p], u \neq v, \quad first_v \geq (k + 1).z_{k,u,v}$

$$\forall u \in [1..p], \sum_{k=1}^n \left\{ \left(\sum_{t \neq u} \frac{\delta_{k-1}}{b_{t,u}} z_{k-1,t,u} \right) + \frac{w_k}{s_u} x_{k,u} + \left(\sum_{v \neq u} \frac{\delta_k}{b_{u,v}} z_{k,u,v} \right) \right\} \leq T_{\text{period}}$$

Linear program: constraints

- $\forall k \in [0..n + 1], \quad \sum_u x_{k,u} = 1$
- $\forall k \in [0..n], \quad \sum_{u \neq v} z_{k,u,v} + \sum_u y_{k,u} = 1$
- $\forall k \in [0..n], \forall u, v \in [1..p] \cup \{in, out\}, u \neq v, x_{k,u} + x_{k+1,v} \leq 1 + z_{k,u,v}$
- $\forall k \in [0..n], \forall u \in [1..p] \cup \{in, out\}, \quad x_{k,u} + x_{k+1,u} \leq 1 + y_{k,u}$
- $\forall k \in [1..n], \forall u \in [1..p], \quad first_u \leq k.x_{k,u} + n.(1 - x_{k,u})$
- $\forall k \in [1..n], \forall u \in [1..p], \quad last_u \geq k.x_{k,u}$
- $\forall k \in [1..n - 1], \forall u, v \in [1..p], u \neq v,$
 $last_u \leq k.z_{k,u,v} + n.(1 - z_{k,u,v})$
- $\forall k \in [1..n - 1], \forall u, v \in [1..p], u \neq v, first_v \geq (k + 1).z_{k,u,v}$

$$\forall u \in [1..p], \sum_{k=1}^n \left\{ \left(\sum_{t \neq u} \frac{\delta_{k-1}}{bt,u} z_{k-1,t,u} \right) + \frac{w_k}{s_u} x_{k,u} + \left(\sum_{v \neq u} \frac{\delta_k}{b_{u,v}} z_{k,u,v} \right) \right\} \leq T_{\text{period}}$$

Linear program: constraints

- $\forall k \in [0..n + 1], \quad \sum_u x_{k,u} = 1$
- $\forall k \in [0..n], \quad \sum_{u \neq v} z_{k,u,v} + \sum_u y_{k,u} = 1$
- $\forall k \in [0..n], \forall u, v \in [1..p] \cup \{in, out\}, u \neq v, x_{k,u} + x_{k+1,v} \leq 1 + z_{k,u,v}$
- $\forall k \in [0..n], \forall u \in [1..p] \cup \{in, out\}, \quad x_{k,u} + x_{k+1,u} \leq 1 + y_{k,u}$
- $\forall k \in [1..n], \forall u \in [1..p], \quad first_u \leq k \cdot x_{k,u} + n \cdot (1 - x_{k,u})$
- $\forall k \in [1..n], \forall u \in [1..p], \quad last_u \geq k \cdot x_{k,u}$
- $\forall k \in [1..n - 1], \forall u, v \in [1..p], u \neq v,$
 $last_u \leq k \cdot z_{k,u,v} + n \cdot (1 - z_{k,u,v})$
- $\forall k \in [1..n - 1], \forall u, v \in [1..p], u \neq v, first_v \geq (k + 1) \cdot z_{k,u,v}$

$$\forall u \in [1..p], \sum_{k=1}^n \left\{ \left(\sum_{t \neq u} \frac{\delta_{k-1}}{b_{t,u}} z_{k-1,t,u} \right) + \frac{w_k}{s_u} x_{k,u} + \left(\sum_{v \neq u} \frac{\delta_k}{b_{u,v}} z_{k,u,v} \right) \right\} \leq T_{\text{period}}$$


Linear program: experiments

- $O(np^2)$ variables, as many constraints
- Experiments only on small problem instances
- Average over 10 instances of each application
- Use GLPK
- Largest experiment: $p = 8$, $n = 4$: 14-hour computation time
- Parameters similar to Experiment 1: homogeneous communications and balanced comm/comp

Linear program: experiments

- $O(np^2)$ variables, as many constraints
- Experiments only on small problem instances
- Average over 10 instances of each application
- Use GLPK
- Largest experiment: $p = 8$, $n = 4$: 14-hour computation time
- Parameters similar to Experiment 1: homogeneous communications and balanced comm/comp

Linear program: experiment $p = 8$




Linear program: experiment $p = 8$

n	LP	H5-BS121	H7b-BSC
1	2.576857	2.576882	2.576882
2	2.749913	2.749934	2.749934
3	2.879871	2.879900	2.883072
4	2.760960	2.760981	2.770690

Linear program: experiment $p = 4$

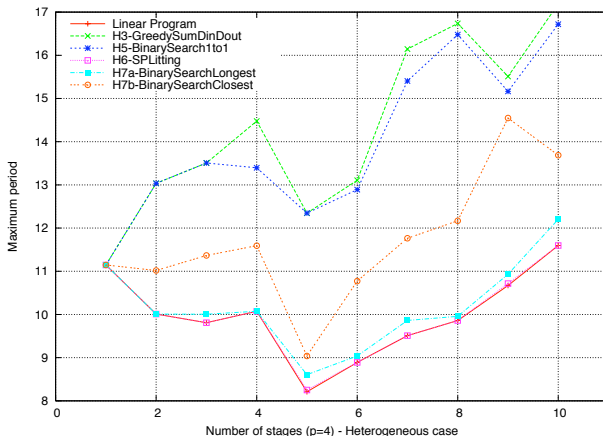
Homogeneous communications (Experiment 1)



H7b very close to the optimal (< 3% error)

Linear program: experiment $p = 4$

Heterogeneous communications (Experiment 2)



H6 very close to the optimal ($< 0.05\%$ error)

Outline

- 1 Framework
- 2 Complexity results
- 3 Heuristics
- 4 Experiments
- 5 Linear programming formulation
- 6 Conclusion**

Related work

- Scheduling task graphs on heterogeneous platforms– Acyclic task graphs scheduled on different speed processors [Topcuoglu et al.]. Communication contention: 1-port model [Beaumont et al.].
- Mapping pipelined computations onto special-purpose architectures– FPGA arrays [Fabiani et al.]. Fault-tolerance for embedded systems [Zhu et al.]
- Mapping pipelined computations onto clusters and grids– DAG [Taura et al.], DataCutter [Saltz et al.]
- Mapping skeletons onto clusters and grids– Use of stochastic process algebra [Benoit et al.]

Conclusion

Theoretical side – Complexity for different mapping strategies and different platform types

Practical side

- Optimal polynomial algorithm for ONE-TO-ONE MAPPING
- Design of several heuristics for INTERVAL MAPPING on *Communication Homogeneous*
- Comparison of their performance
- Linear program to assess the absolute performance of the heuristics, which turns out to be quite good

Future work

Short term

- Heuristics for *Fully Heterogeneous* platforms
- Extension to DAG-trees (a DAG which is a tree when un-oriented)
- Extension to stage replication
- LP with replication and DAG-trees

Longer term

- Real experiments on heterogeneous clusters, using an already-implemented skeleton library and MPI
- Comparison of effective performance against theoretical performance