



Minimizing rental cost for multiple recipe applications in the Cloud

F. Hana, L. Marchal, J.-M. Nicod, L. Philippe, V. Rehn-Sonigo and H. Sabbah



LIP-ENS Lyon – FEMTO-ST institute - UFC/ENSMM Besançon

Nashville - May 18th, 2016


Introduction and motivation




Introduction and motivation



Introduction and motivation




Introduction and motivation



Overall objective: To provision just enough resources to reach the target throughput for a given DAG-based streaming application

Application framework

Each workflow application φ^j produces the same result Φ .



- Each task φ_i^j has a task type q
- Target throughput ρ
- Each application can be run at a different throughput ρ_j
- $\rho = \sum_j \rho_j$


Target platform

One processor type per task type

- c_q : Rental cost for type q
- r_q : Throughput of type q

Application framework

Each workflow application φ^j produces the same result Φ .



- Each task φ_i^j has a task type q
- Target throughput ρ
- Each application can be run at a different throughput ρ_j
- $\rho = \sum_j \rho_j$

Target platform

One processor type per task type

- c_q : Rental cost for type q
- r_q : Throughput of type q



Find the cheapest configuration to reach the target throughput



MinCOST: Minimize the global rental cost C

Given

- an application described by J graphs
 - a platform described by processor cost c_q and throughput r_q
 - a target throughput ρ
- ⇒ select which graphs φ^j are used
- ⇒ choose their output throughput ρ_j ($\rho_j = 0$ if unused)
- ⇒ deduce the number of processors x_q of each type to reach the prescribed throughput within minimal cost.



Simple case

- The application is described by a single graph

General case

$$\rho = \sum_j \rho_j$$

- Black box application
- Application graphs without shared task types
- Application graphs with shared task types

Simple case: Single application graph

- One application described by one single graph φ^1
- $\forall q$, the number of machines x_q can be easily computed:

$$x_q = \left\lceil \frac{n_q}{r_q} \cdot \rho \right\rceil$$


- The associated cost C_q

$$C_q(\rho) = \left\lceil \frac{n_q}{r_q} \cdot \rho \right\rceil \times c_q$$

- The final cost C :

$$C(\rho) = \sum_{q=1}^Q C_q(\rho) = \sum_{q=1}^Q \left\lceil \frac{n_q}{r_q} \cdot \rho \right\rceil \times c_q$$

n_q : number of tasks, r_q : throughput, c_q : cost of type q



General case: Black box applications




- Each graph $\varphi^j = \varphi_1^j$ is one complex task

$$\forall j \text{ and } \forall j' (1 \leq j, j' \leq J) : t(1, j) = t(1, j') \Rightarrow j = j'$$

- Let ρ_q the output of $\varphi^j = \varphi^q$
- x_q can be found by solving the following linear program:

$$\left\{ \begin{array}{l} \text{Minimize } C(\rho) = \sum_{q=1}^Q x_q c_q \\ \text{Under the constraint } \sum_{q=1}^Q x_q \rho_q \geq \rho \end{array} \right.$$

⇒ This resembles a knapsack problem with repetition using negative weights and values



General case: Black box applications



Unbounded Knapsack Problem

Given n objects with value v_i and weight w_i , and a total capacity of W , how many copies of each object should we select to maximize the total value without exceeding weight W ?




ILP formulation with x_i the number of copies of item i included in the solution

$$\begin{cases} \text{Maximize } \sum x_i v_i \\ \text{Under the constraint } \sum x_i w_i \leq W \end{cases}$$

Our problem is thus equivalent to a knapsack problem where:

- Items have value $(-c_q)$ and weight $(-\rho_q)$
 - The total capacity is $(-\rho)$.
- ⇒ The knapsack problem is a (unary) NP-complete problem
- ⇒ There exists a pseudo-polynomial dynamic program (time complexity $O(J\rho)$)

Application graphs without shared task types



- Application Φ can be described by $\varphi^1, \dots, \varphi^j, \dots, \varphi^J$ with the same output result
- Each task φ_i^j from one graph φ^j has a different type from every other task of an other graph $\varphi^{j'}$

$$t(i, j) \neq t(i', j') \text{ with } 1 \leq j, j' \leq J \text{ and } j \neq j' \text{ and } 1 \leq i \leq l_j \text{ and } 1 \leq i' \leq l_{j'}$$

- This problem is unary NP-complete (includes Black Box applications)
- ⇒ There exists a pseudo-polynomial dynamic program to solve it



A dynamic program to solve this problem

- Let $C(\rho, j)$ be the optimal platform cost to reach ρ using the first j application graphs

$$C(\rho, j) = \begin{cases} \sum_{i=1}^{l_j} \left\lceil \frac{n_{t(i,1)}^1}{r_{t(i,1)}} \cdot \rho \right\rceil \times c_{t(1,k)} & \text{if } j = 1 \\ \min_{0 \leq \rho_j \leq \rho} \left(C(\rho - \rho_j, j - 1) + \right. \\ \quad \left. \sum_{i=1}^{l_j} \left\lceil \frac{n_{t(i,j)}^j}{r_{t(i,j)}} \cdot \rho_j \right\rceil \times c_{t(i,j)} \right) & \text{otherwise} \end{cases}$$




⇒ The solution is given by $C(\rho, J)$



Complexity analysis

- As $\forall q, r_q \in \mathbb{N}, \forall j \rho_j \in \mathbb{N}$
 - There exists a finite number of ρ_j to test in the previous formulation
 - To compute $C(\rho, j)$, all $C(\rho', j')$ with $\rho' \leq \rho$ and $j' \leq j$ has to be computed
 - ⇒ The complexity of the elementary computation is $O(\rho l)$
 - ⇒ The complexity of computing $C(\rho)$ is $O(\rho^2 l J)$

Application graphs with shared task types



- One application is described by several graphs which share task types

$$\exists \varphi^j, \varphi^{j'} (1 \leq j, j' \leq J, j \neq j'), \exists i (1 \leq i \leq l_j), \\ \exists i' (1 \leq i' \leq l_{j'}) \text{ s.t. } t(i, j) = t(i', j')$$

⇒ A processor may be shared between several application graphs



ILP formulation

Minimizing $C(\rho) = \sum_{q=1}^Q x_q \cdot c_q$

under the constraints

- ρ has to be at least the sum of ρ_j

$$\sum_{j=1}^J \rho_j \geq \rho$$

- For each type q we have to provision enough resources (x_q)

$$\forall q \ x_q \cdot r_q \geq \sum_{j=1}^J \left(\sum_{i=1}^{l_j} \rho_j \right),$$

with $q = t(i, j)$ and $x_q \in \mathbb{N}$

Application graphs with shared task types



ILP formulation

Minimizing $C(\rho) = \sum_{q=1}^Q x_q \cdot c_q$

under the constraints

- ρ has to be at least the sum of ρ_j

$$\sum_{j=1}^J \rho_j \geq \rho$$

- For each type q we have to provision enough resources (x_q)

$$\forall q \ x_q \cdot r_q \geq \sum_{j=1}^J \left(\sum_{i=1}^{l_j} \rho_j \right),$$

with $q = t(i, j)$ and $x_q \in \mathbb{N}$

The complexity of this case is still open
unary or binary NP-complete



Heuristic approaches

$$\rho = (\rho_1, \rho_2, \dots, \rho_J)$$

1. H0: random
2. H1: best graph
3. H2: random walk
4. H31: stochastic descent
5. H32/H32Jump: steepest gradient



Heuristic approaches

$$\rho = (\rho_1, \rho_2, \dots, \rho_J)$$

1. H0: random
2. H1: best graph
3. H2: random walk
4. H31: stochastic descent
5. H32/H32Jump: steepest gradient

$$\rho = (\rho_1, \rho_2, \dots, \rho_J)$$



Heuristic approaches

$$\rho = (\rho_1, \rho_2, \dots, \rho_J)$$

1. H0: random
2. H1: best graph
3. H2: random walk
4. H31: stochastic descent
5. H32/H32Jump: steepest gradient

$$\rho = (0, \dots, \rho, \dots, 0)$$



Heuristic approaches

$$\rho = (\rho_1, \rho_2, \dots, \rho_J)$$

1. H0: random
2. H1: best graph
3. H2: random walk
 - φ_{j1} and φ_{j2} are randomly chosen

$$(\dots, \rho_{j1}, \dots, \rho_{j2}, \dots) \rightarrow (\dots, \rho_{j1} - \delta, \dots, \rho_{j2} + \delta, \dots)$$

$$(\dots, \rho_{j1}, \dots, \rho_{j2}, \dots) \rightarrow (\dots, 0, \dots, \rho_{j2} + \rho_{j1}, \dots) \text{ if } \rho_{j1} < \delta$$

4. H31: stochastic descent
5. H32/H32Jump: steepest gradient



Heuristic approaches

$$\rho = (\rho_1, \rho_2, \dots, \rho_J)$$

1. H0: random
2. H1: best graph
3. H2: random walk
4. H31: stochastic descent
 - Same as H2 except that we keep the same solution as long as we do not obtain any improvement
5. H32/H32Jump: steepest gradient




Heuristic approaches

$$\rho = (\rho_1, \rho_2, \dots, \rho_J)$$

1. H0: random
2. H1: best graph
3. H2: random walk
4. H31: stochastic descent
5. **H32/H32Jump: steepest gradient**
 - H32/H32Jump same as H2 except we test all possible throughput fraction exchanges and keep the best until no more improvement is possible
 - H32Jump allows to explore solution that increases $C(\rho)$

Illustrating example




Processor	type	ρ	cost
P_1	t_1	10	10
P_2	t_2	20	18
P_3	t_3	30	25
P_4	t_4	40	33

Results

ρ	ILP				H1				H 2				H 32 JUMP			
	ρ_1	ρ_2	ρ_3	cost	ρ_1	ρ_2	ρ_3	cost	ρ_1	ρ_2	ρ_3	cost	ρ_1	ρ_2	ρ_3	cost
10	0	0	10	28	0	0	10	28	0	0	10	28	0	0	10	28
40	40	0	0	69	40	0	0	69	40	0	0	69	40	0	0	69
50	10	30	10	86	0	0	50	104	10	30	10	86	10	30	10	86
130	30	90	10	220	0	0	130	256	30	90	10	220	90	30	10	224
140	0	120	20	237	0	140	0	257	0	120	20	237	0	120	20	237
150	0	150	0	257	0	150	0	257	0	150	0	257	0	150	0	257
200	20	180	0	333	0	200	0	340	20	180	0	333	20	180	0	333


Experiments: small application graphs



ILP: Gurobi
Simulator: Python

Normalization of cost with the optimal solution
20 alternative graphs, between 5 and 8 tasks for each graph


Experiments: small application graphs



ILP H1 H2 H31 H32 H32Jump

Number of times where each algorithm finds the best 20 alternative graphs, between 5 and 8 tasks for each graph


Experiments: large application graphs



ILP H1 H2 H31 H32 H32Jump

Normalization of cost with the optimal solution
20 alternative graphs, between 50 and 100 tasks for each graph

Experiments: large application graphs



ILP H1 H2 H31 H32 H32Jump

Computation time for the heuristics
20 alternative graphs, between 100 and 200 tasks for each graph



- **Efficient ILP solver:**
 - Optimal solutions for small and medium sized problems
 - Fails for applications with more than 100 tasks
- The naive heuristic **H1** gives a good solution with minimal overhead
- More **sophisticated heuristics** only improve H1 up to 5%
- H1 approach gives solutions whose costs are asymptotically close to the optimal (ILP if possible)

Conclusion



Find the cheapest configuration to reach the target throughput for a given DAG based streaming application

- The issue was to find a **suitable distribution** between DAGs
- ⇒ We deduce the platform to rent on the Cloud (minimize the rental cost)

Conclusion



Find the cheapest configuration to reach the target throughput for a given DAG based streaming application

- The issue was to find a **suitable distribution** between DAGs
- ⇒ We deduce the platform to rent on the Cloud (minimize the rental cost)
- In some cases we exhibit algorithms to **optimally solve** the problem (even if NP-complete in the weak sens)
 - The complexity of the most general case remains open
- ⇒ **ILP** gives a characterization of an optimal solution

Conclusion



Find the cheapest configuration to reach the target throughput for a given DAG based streaming application

- The issue was to find a **suitable distribution** between DAGs
- ⇒ We deduce the platform to rent on the Cloud (minimize the rental cost)
- In some cases we exhibit algorithms to **optimally solve** the problem (even if NP-complete in the weak sens)
- The complexity of the most general case remains open
- ⇒ **ILP** gives a characterization of an optimal solution
- **Heuristics** with good performance (6% from the optimal and asymptotically optimal)



economical cost \leftrightarrow energy cost

Green computing

- How to take energy into account when we rent resources in the Cloud ?
- How to associate both economical and energetical criteria