

Resource Allocation for Multiple Concurrent In-network Stream-processing Applications

Anne Benoit, Henri Casanova, Veronika Rehn, Yves Robert

▶ **To cite this version:**

Anne Benoit, Henri Casanova, Veronika Rehn, Yves Robert. Resource Allocation for Multiple Concurrent In-network Stream-processing Applications. Seventh International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms, Aug 2009, Delft, Netherlands. inria-00424049

HAL Id: inria-00424049

https://inria.hal.science/inria-00424049

Submitted on 13 Oct 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Resource Allocation for Multiple Concurrent In-network Stream-processing Applications

Anne Benoit[1], Henri Casanova[2], Veronika Rehn-Sonigo[1], and Yves Robert[1]

[1] LIP, ENS Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France
{Anne.Benoit,Veronika.Sonigo,Yves.Robert}@ens-lyon.fr
[2] University of Hawai'i at Manoa, 1680 East-West Road, Honolulu, HI 96822, USA
henric@hawaii.edu

Abstract. This paper investigates the operator mapping problem for in-network stream-processing applications. In-network stream-processing is the application of one or several trees of operators, in steady-state, to data that are continuously updated at different locations in the network. The goal is to generate final results at a desired rate. Different operator trees may share common subtrees, so that intermediate results could be reused in different applications. This work provides complexity results for different instances of the basic problem and proposes several polynomial-time heuristics. Quantitative comparison of the heuristics in simulation demonstrates the importance of mapping operators to appropriate processors, and allows us to identify a heuristic that achieves good results in practice.

1 Introduction

We consider applications structured as trees of operators, where leaves correspond to basic data objects distributed in a network. Each internal node in the tree denotes the aggregation and combination of the data from its children, which in turn generates new data that is used by the node's parent. The computation is complete when all operators have been applied up to the root node, thereby producing a final result. We consider the scenario in which the basic data objects are constantly being updated, meaning that the tree of operators must be applied continuously. The goal is to produce final results at some desired rate.

The above problem is called *stream processing* [1] and arises in several domains. One such domain is the acquisition and refinement of data from a set of sensors [2]. For instance, [2] outlines a video surveillance application in which the sensors are cameras located at different locations over a geographical area. Another example arises in the area of network monitoring [3,4]. In this case routers produce streams of data pertaining to forwarded packets. More generally, stream processing can be seen as the execution of one of more "continuous queries" in the relational database sense of the term (e.g., a tree of join and select operators). Many authors have studied the execution of continuous queries on data streams [5,6].

In practice, the execution of the operators must be distributed over the network. In some cases the servers that produce the basic objects may not have the computational capability to apply all operators. Besides, objects must be combined across devices, thus requiring network communication. Sending all basic objects to a central compute server often proves unscalable due to network bottlenecks, or due to the central server not providing sufficient computational power. The alternative is to distribute the execution by mapping each node in the operator tree to one or more servers in the network, including servers that produce and update basic objects and/or servers that are only used for applying operators. One then talks of *in-network stream-processing*. Several in-network stream-processing systems have been developed [7,4]. These systems all face the same question: where should operators be mapped in the network?

In this paper we study the operator-mapping problem for *multiple concurrent in-network stream-processing applications*. The problem for a single application was studied in [8] for an ad-hoc objective function that trades off application delay and network bandwidth consumption. In a recent paper [9] we have studied a more general objective function, enforcing the constraint that the rate at which final results are produced, or *throughput*, is above a given threshold. This corresponds to a Quality of Service (QoS) requirement of the application that should be met while using as few resources as possible. In this paper we extend the work in [9] in two ways. First, we study a "non-constructive" scenario, i.e., we are given a set of compute and network elements, and we attempt to use as few resources as possible. Instead, in [9], we studied a "constructive" scenario in which resources could be purchased and the objective was to spend as little money as possible. Second, while in [9] we studied the case of a single application, in this paper we focus on multiple concurrent applications that contend for the servers, each with its own QoS requirement. Indeed, with several applications from several users running concurrently, it is more likely to share an existing set of resources for a common deployment, hence the call for the non-constructive scenario. Higher performance and reduced resource consumption is possible by reusing common sub-expression between operator trees when applications share basic objects [10]. We consider target platforms that are either fully homogeneous, or with a homogeneous network but heterogeneous servers, or fully heterogeneous. We formalize operator mapping problems for multiple in-network stream-processing applications and give their complexity; and we propose heuristics to solve the problems and evaluate them in simulation.

2 Framework

Application Model – We consider $\mathcal{K}$ applications, each needing to perform several operations organized as a binary tree (see Fig. 1). Operators are taken from the set $\mathcal{OP} = \{op_1, op_2, \dots\}$, and operations are initially performed on basic objects from the set $\mathcal{OB} = \{ob_1, ob_2, \dots\}$. These basic objects are made available and continuously updated at given locations in a distributed network. Operators higher in the tree rely on previously computed intermediate results, and they may also require to download basic objects periodically.

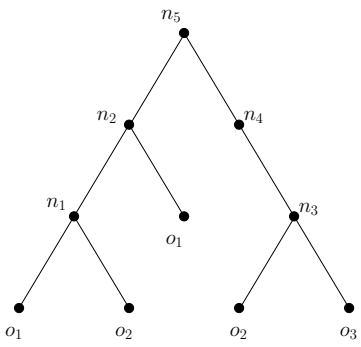


Fig. 1. Sample application structured as a binary tree of operators.

For an operator op_p we define $objects(p)$ as the index set of the basic objects in $\mathcal{OB}$ that are needed for the computation of op_p, if any; and $operators(p)$ as the index set of operators in $\mathcal{OP}$ whose intermediate results are needed for the computation of op_p, if any. We have $|objects(p)| + |operators(p)| \leq 2$.

The tree structure of application k is defined with a set of labeled nodes. The i^{th} internal node in the tree of application k is denoted as $n_i^{(k)}$, its associated operator is denoted as $op(n_i^{(k)})$, and the set of basic objects required by this operator is denoted as $ob(n_i^{(k)})$. Node $n_1^{(k)}$ is the root node. Let $op_p = op(n_i^{(k)})$ be the operator associated to node $n_i^{(k)}$. Then node $n_i^{(k)}$ has $|operators(p)|$ child nodes.

The applications must be executed so that they produce final results, where each result is generated by executing the whole operator tree once, at a target rate. We call this rate the application *throughput*, $\rho^{(k)}$, specified as a QoS requirement for each application. Each operator in the tree of the k^{th} application must compute (intermediate) results at a rate at least as high as $\rho^{(k)}$. Conceptually, operator op_p executes two concurrent threads in steady-state. **(1)** It periodically downloads (or continuously stream) the most recent copies of the basic objects in $objects(p)$, if any. Basic object ob_j has size d_j (in bytes) and needs to be downloaded by the processors that use it for application k with frequency $f_j^{(k)}$. This consumes an amount of bandwidth of $rate_j^{(k)} = d_j \times f_j^{(k)}$ on each involved network link and network card. If a processor requires object ob_j for several applications with different update frequencies, it downloads the object only once at the maximum required frequency $rate_j = \max_k\{rate_j^{(k)}\}$. **(2)** It receives intermediate results computed by $operators(p)$, if any, and performs computation using basic objects it is continuously downloading and/or data received from other operators. The computation of operator op_p requires w_p operations, and produces an output of size δ_p.

Platform Model – The distributed network is a fully connected graph (i.e., a clique) interconnecting a set of processors $\mathcal{P}$. Operators are mapped onto these processors. Some processors also hold and update basic objects. Processor $P_u \in \mathcal{P}$ is interconnected to the network via a network card with maximum bandwidth B_u. The network link between two distinct processors P_u and P_v is bidirectional and has bandwidth $b_{u,v} (= b_{v,u})$, shared by communications in both directions. Processor $P_u \in \mathcal{P}$ has compute speed s_u. Processors that only provide basic objects and cannot compute are simply given compute speed 0. Resources operate under the full-overlap, bounded multi-port model [11]: Processor P_u can simultaneously compute, send, and receive data. With the "multi-port" assump-

tion, each processor can send/receive data simultaneously on multiple network links. The "bounded" assumption enforces that the total transfer rate of data sent/received by processor P_u is bounded by its network card bandwidth, B_u.

Mapping Model and Constraints – The objective is to map internal nodes of application trees onto processors. If only one node is mapped to processor P_u, while P_u computes for the t-th final result it sends to its parent (if any) intermediate results for the $(t-1)$-th final result and it receives data from its children (if any) for computing the $(t+1)$-th final result. All three activities are concurrent. If several nodes are mapped to P_u the same overlap happens, but possibly on different result instances. A basic object can be duplicated, and thus available and updated at multiple processors. We assume that such duplication is achieved in some application-specific manner (e.g., via a distributed database that enforces sufficient data consistency). In this case, a processor can choose among multiple data sources for a basic object (or perform a local access if the basic object is available locally.)

We use an allocation function, a, to denote the mapping of the nodes onto the processors in $\mathcal{P}$: $a(k,i) = u$ if node $n_i^{(k)}$ is mapped to processor P_u. Conversely, $\bar{a}(u)$ is the index set of nodes mapped on P_u: $\bar{a}(u) = \{(k,i) \mid a(k,i) = u\}$. Also, we denote by $a_{op}(u)$ the index set of operators mapped on P_u: $a_{op}(u) = \{p \mid \exists (k,i) \in \bar{a}(u) \;\; op_p = op(n_i^{(k)})\}$. We introduce the following notations:

• $Ch(u) = \{(p,v,k)\}$ is the set of (operator, processor, application) tuples such that processor P_u needs to receive an intermediate result computed by operator op_p, which is mapped to processor P_v, at rate $\rho^{(k)}$; operators op_p are children of $a_{op}(u)$ in the operator tree.

• $Par(u) = \{(p,v,k)\}$ is the set of (operator, processor, application) tuples such that P_u needs to send to P_v an intermediate result computed by operator op_p at rate $\rho^{(k)}$; $p \in a_{op}(u)$ and the sending is done to the parents of op_p in the operator tree.

• $Do(u) = \{(j,v,k)\}$ is the set of (object, processor, application) tuples where P_u downloads object ob_j from processor P_v at rate $\rho^{(k)}$.

Given these notations, we can express constraints for the application throughput: each processor must compute and communicate fast enough to respect the prescribed throughput of each application with nodes allocated to it (Eq. 1). Note that each operator is computed only once at the maximum required throughput.

$$\forall P_u \in \mathcal{P} \quad \sum_{p \in a_{op}(u)} \left(\max_{(k,i) \in \bar{a}(u) \mid op(n_i^{(k)}) = op_p} \left(\rho^{(k)} \right) \frac{w_p}{s_u} \right) \leq 1 \; . \tag{1}$$

Communication occurs only when child and parent nodes are mapped on different processors. An operator computing for several applications may send/receive results to/from different processors. If the parent/child nodes corresponding to the different applications are mapped onto the same processor, the communication is done only once, at the most constrained throughput. In expressions below $v \neq u$ since we neglect intra-processor communications.

P_u must have enough bandwidth capacity to perform all its basic object downloads, to support downloads of the basic objects it may hold, and also to perform all communication with other processors, all at the required rates (Eq. 2). The first term corresponds to basic object downloads; the second term corresponds to download of basic objects from other processors; the third term corresponds to inter-node communications when a node is assigned to P_u and its parent node is assigned to another processor; and the last term corresponds to inter-node communications when a node is assigned to P_u and some of its children nodes are assigned to another processor.

$$\forall P_u \in \mathcal{P} \sum_{\substack{(j,v,k) \\ \in Do(u)}} rate_j^{(k)} + \sum_{P_v \in \mathcal{P}} \sum_{\substack{(j,u,k) \\ \in Do(v)}} rate_j^{(k)} + \sum_{\substack{(p,v,k) \\ \in Ch(u)}} \delta_p \rho^{(k)} + \sum_{\substack{(p,v,k) \\ \in Par(u)}} \delta_p \rho^{(k)} \leq B_u \quad (2)$$

Finally, the link between processor P_u and processor P_v must have enough bandwidth capacity to support all possible communications between the nodes mapped on both processors, as well as the object downloads (Eq. 3).

$$\forall P_u, P_v \in \mathcal{P} \sum_{\substack{(j,v,k) \\ \in Do(u)}} rate_j^{(k)} + \sum_{\substack{(j,u,k) \\ \in Do(v)}} rate_j^{(k)} + \sum_{\substack{(p,v,k) \\ \in Ch(u)}} \delta_p \rho^{(k)} + \sum_{\substack{(p,v,k) \\ \in Par(u)}} \delta_p \rho^{(k)} \leq b_{u,v} \quad (3)$$

Optimization Problems – The goal is to achieve a prescribed throughput for each application while minimizing a cost function. Several relevant problems can be envisioned. Proc-Nb minimizes the number of used processors; Proc-Power minimizes the compute capacity and/or the network card capacity of used processors (e.g., a linear function of both criteria); BW-Sum minimizes the sum of the used bandwidth capacities; and BW-Max minimizes the maximum percentage of bandwidth used on all links. Different platform types may be considered depending on resource heterogeneity. We consider the fully homogeneous case ($s_u = s$, $B_u = B$ and $b_{u,v} = b$), which we term Hom. The case in which network links can have various bandwidths is termed Het.

3 Complexity

Problem Proc-Nb is NP-complete in the strong sense even for a simple case: a Hom platform and a single application ($|\mathcal{K}| = 1$), that is structured as a left-deep tree [12], in which all operators take the same amount of time to compute and produce results of size 0, and in which all basic objects have the same size. We refer the reader to [9] for the proof. It turns out that the same proof holds for Proc-Power on a Hom platform.

The BW-Max problem is NP-hard because downloading objects with different rates on two processors is the same as the NP-hard 2-Partition problem [13]. Here is a sketch of the straightforward proof for a single application. Consider an application in which all operators produce zero-size results, and in which each basic object is used only by one operator. Consider three processors, with one of them holding all basic objects but unable to compute any operator. The two remaining processors are able to compute all the operators, and they are

connected to the first one with identical network links. Such an instance can be easily constructed. The goal is to partition the set of operators in two subsets so that the bandwidth consumption on the two network links is as equal as possible. This is exactly the 2-Partition problem.

The BW-SUM problem can be reduced to the NP-hard Knapsack problem [13]. Here is a proof sketch for a single application. Consider the same application as for the proof of the NP-hardness of BW-MAX above. Consider two identical processors, A and B, with A holding all basic objects. Not all operators can be executed on A and a subset of them need to be executed on B. Such an instance can be easily constructed. The problem is then to determine the subset of operators that should be executed on A. This subset should satisfy the constraint that the computational capacity of A is not exceeded, while maximizing the bandwidth cost of the basic objects associated to the operators in the subset. This is exactly the Knapsack problem.

All above problems can be solved via linear programming (see [14] for Integer Linear Program formulations). However, they cannot be solved in polynomial time (unless P=NP).

4 Heuristics

In this section we propose polynomial heuristics[3] for solving the PROC-POWER problem when considering only the compute capacities of used processors. We propose 5 *heuristics* to map application nodes to processors. Each heuristic can use one of 4 generic *processor selection strategies* to select which processor a node should be mapped to. We consider two processor selection strategies, each with a *blocking* and a *non-blocking* version. *Blocking* means that once chosen for a given operator op_1, a processor cannot be used later for another operator op_2 unless op_2 is a relative (i.e., father or child) of op_1. *Non-blocking* heuristics impose no such restriction. We obtain four strategies *(S1) Select the fastest processor (blocking)*; *(S2) Select the processor with the fastest network card (blocking)*; *(S3) Select the fastest processor (non-blocking)*; and *(S4) Select the processor with the fastest network card*. Note that the processor and network card speeds used for the selection are computed while accounting for operators that may have already been mapped to servers.

All our heuristics attempt to re-use results from common operator sub-trees across applications. For this purpose they try to add additional communications as show in Fig. 2 on an example. We consider the following 5 heuristics:

• **(H1) Random** – H1 randomly picks the next node to map and attempts to reuse sub-trees across applications. If the node's operator has not already been mapped, possibly for another application, but the node's parent, H1 tries to map the node to the same processor. If unsuccessful, it makes similar attempts with the node's children. Otherwise if the node's operator has already been mapped somewhere else in the forest, H1 tries to add a communication from the already

[3] To ensure the reproducibility of our results, the code for all heuristics is available on the web [15].

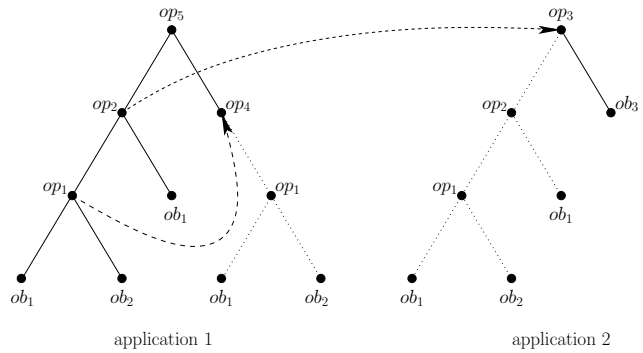


Fig. 2. Example for the reuse of nodes. op_1 is only computed once and its result is reused for the computation of op_2 and op_4. op_3 uses the result of op_2 in application 1 for its computation.

mapped operator to the father of the current node to reuse the common result. In this case, H1 marks the whole subtree (rooted at the operator) as mapped. Otherwise, H1 chooses a new processor according to the selected processor selection strategy. If unsuccessful, then H1 fails.

• **(H2) TopDownBFS** – H2 performs a breadth-first-search (BFS) traversal of all application trees, using an artificial node at which all application trees are rooted. For each node, H2 checks whether its operator has not been mapped yet and whether its father's has. In this case, H2 tries to map the operator on the same processor as its father, and in case of success continues the BFS traversal. If the node's operator has already been mapped, H2 tries to add a communication link between the mapped operator and the node's father: the mapped operator sends its result not only to its father but also to the node's father. If none of these two conditions holds, or if the mapping was not possible, H2 picks a processor according to the the processor selection strategy. If the mapping is successful, the BFS traversal continues, otherwise H2 fails.

• **(H3) TopDownDFS** – H4 uses the same mechanism as H2, but with a depth-first-search (DFS).

• **(H4) BottomUpBFS** – Like H2, H4 performs a BFS traversal of the application trees. For each node, H4 verifies whether it's operator has already been mapped. In this case a communication link is added (if possible), connecting the mapped operator and the node's father. If the operator is not yet mapped and if it has children, H4 tries to map the operator to one of its children's processors. If unsuccessful, or if the operator is at the bottom of a tree, H4 tries to map the operator onto a new processor chosen according to the processor selection strategy. If the mapping is successful, the traversal continues, otherwise H4 fails.

• **(H5) BottomUpDFS** – H5 is similar to H4, but uses a DFS traversal. This adds complexity as more cases need to be considered. For each node H5 checks if its operator has already been mapped and none of its children has. In this case H5 goes up in the tree until it reaches the last node n_1 such that there exists a node n_2 somewhere else in the forest whose operator is already mapped, and such that $op(n_1) = op(n_2)$. In this case H5 tries to add a communication between n_2

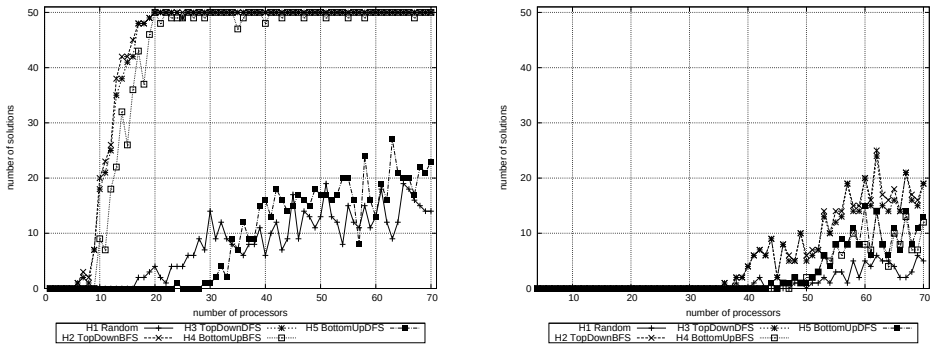
and the n_1's father to share a sub-tree. If the children have already been mapped H5 simply tries to map the operator to one of the children's processors. If this is not possible, or if the additional communication was not possible, or again if the operator has not been mapped anywhere in the forest, H5 tries to map the operator onto a new processor chosen according to the processor selection strategy. Otherwise H5 fails.

5 Experimental Results

We have conducted several experiments to assess the performance of the different heuristics described in Section 4. In particular, we are interested in the impact of node reuse on the number of solutions found by the heuristics. The application trees are fixed to a size of at most 50 operators, and in general we consider 5 concurrent applications. The following parameters are chosen randomly: The basic objects (leaves in the tree) are chosen among 10 different types. The size d of each object type varies between 3MB and 13MB. The download frequencies of objects for each application, f, as well as the application throughput, ρ, are such that $0 < f \leq 1$ and $1 \leq \rho \leq 2$. The operands of operators are also chosen randomly. The computation amount w_i for an operator lies between 0.5MFlop/sec and 1.5MFlop/sec, and the output size of each operator δ_i varies between 0.5MB and 1.5MB. We dispose of 30 processors, equipped with a network card of bandwidth between 50MB and 180MB each. We use the same range for processor compute speed : 50MIPS to 180MIPS. Processors are interconnected via heterogeneous communication links, whose bandwidths are between 60MB/s and 100MB/s. The 10 different types of objects are randomly distributed over the processors, where objects are maximal twice available on processors. When deciding about basic object downloads, we first try to download from processors which are already used in the mapping (with minimal available bandwidth), before downloading from an unused processor. Execution time and communication time are scaled units, thus execution time is the ratio between computation amount and processor speed, while communication time is the ratio between object size (or output size) and link bandwidth.

We study the relative performance of each heuristic compared to the best solution found by any heuristic. This allows to compare the cost, in amount of resources used, of the different heuristics. The relative performance for heuristic h is obtained by: $\frac{1}{|runs|}\sum_{r=1}^{|runs|} a_h(r)$, where $a_h(r) = 0$ if heuristic h fails in run r and $a_h(r) = \frac{cost_{best}(r)}{cost_h(r)}$. $cost_{best}(r)$ is the best solution cost returned over all heuristics for run r, and $cost_h(r)$ is the cost in the solution returned by heuristic h. The number of runs is fixed to 50 in all experiments. The complete set of results is available on the web [15].

Summary of Experiments – We have performed different test series, varying the number of processors, the number of applications and the application size. Also we tested the impact of the Communication-to-Computation Ratio (CCR), which is the ratio between the mean amount of communications and the mean



(a) Successful runs, strategy 3 (b) Successful runs without reuse, strategy 3

Fig. 3. Experiment: Increasing number of processors. Number of successful runs.

amount of computations. Finally we were interested in the influence of application similarities on the heuristics' performance. Due to lack of space, we resume our experimental results, but a detailed description is available in [14].

Our results show that a random approach for multiple applications performs considerably bad. Not reusing results from common subtrees dramatically limits the success rate (see Fig. 3) and also the quality of the solution in terms of cost (relative performance). The TopDown approach turns out to be the best, and in most cases BFS traversal achieves the best result. The BottomUp approach is only competitive using a BFS traversal. The DFS traversal seems unable to reuse results efficiently (it often finds itself with no bandwidth left to perform necessary communications.) Furthermore we see a strong dependency of the processor selection strategy on solution quality. The blocking strategies outperform the non-blocking strategies when the CCR is large. Overall, H2 in combination with strategy S3 proves to be a solid combination.

6 Conclusion

We have studied the operator mapping problem of multiple concurrent in-network stream-processing applications onto a collection of heterogeneous processors. These applications come as a set of operator trees, that have to continuously download basic objects at different sites of the network and at the same time have to process this data to produce some final result. We have identified four relevant optimization problems. All are NP-hard but can be formalized as integer linear programs. Focusing on one of these optimization problems, we have designed several polynomial-time heuristics, which we have evaluated in simulation. Our experiments show the importance of node reuse across applications. Reusing nodes leads to an important number of additional solutions, and also the quality of the solutions improves considerably. We conclude that top-down traversal of the application trees is more efficient than bottom-up traversal.

As future work, we could develop heuristics for the other optimization problems defined in Section 2. We could also envision a more general cost function $w_{i,u}$ (time required to compute operator i onto processor u), in order to express even more heterogeneity. This would lead to the design of more sophisticated heuristics. Also, we believe it would be interesting to add a storage cost for objects downloaded onto processors, which could lead to new objective functions.

References

1. Badcock, B., Babu, S., Datar, M., Motwani, R., Widom, J.: Models and issues in data stream systems. In: Proceedings of the Intl. Conf. on Very Large Data Bases. (2004) 456–467
2. Srivastava, U., Munagala, K., Widom, J.: Operator Placement for In-Network Stream Query Processing. In: Proceedings of the 24th ACM Intl. Conf. on Principles of Database Systems. (2005) 250–258
3. Cranor, C., Gao, Y., Johnson, T., Shkapenyuk, V., Spatscheck, O.: Gigascope: high-performance network monitoring with an SQL interface. In: Proc. of the ACM SIGMOD International Conference on Management of Data. (2002) 623–633
4. van Rennesse, R., Birman, K., Dumitriu, D., Vogels, W.: Scalable Management and Data Mining Using Astrolabe. In: Proceedings from the First International Workshop on Peer-to-Peer Systems. (2002) 280–294
5. Babu, S., Widom, J.: Continuous Queries over Data Streams. SIGMOD Record **30**(3) (2001)
6. Plale, B., Schwan, K.: Dynamic Querying of Streaming Data with the dQUOB System. IEEE Trans. on Parallel and Distributed Systems **14**(4) (2003) 422–432
7. Chen, L., Reddy, K., Agrawal, G.: GATES: a grid-based middleware for processing distributed data streams. High performance Distributed Computing, 2004. Proceedings. 13th IEEE International Symposium on (2004) 192–201
8. Pietzuch, P., Leflie, J., Shneidman, J., Roussopoulos, M., Welsh, M., Seltzer, M.: Network-Aware Operator Placement for Stream-Processing Systems. In: Proceedings of the 22nd International Conference on Data Engineering (ICDE'06). (2006) 49–60
9. Benoit, A., Casanova, H., Rehn-Sonigo, V., Robert, Y.: Resource Allocation Strategies for Constructive In-Network Stream Processing. In: Proceedings of APDCM'09, the 11th Workshop on Advances in Parallel and Distributed Computational Models, IEEE (2009)
10. Pandit, V., Ji, H.: Efficient in-network evaluation of multiple queries. In: HiPC. (2006) 205–216
11. Hong, B., Prasanna, V.: Distributed adaptive task allocation in heterogeneous computing environments to maximize throughput. In: Intl. Parallel and Distributed Processing Symposium IPDPS'2004, IEEE Computer Society Press (2004)
12. Ioannidis, Y.E.: Query optimization. ACM Computing Surveys **28**(1) (1996) 121–123
13. Garey, M.R., Johnson, D.S.: Computers and Intractability, a Guide to the Theory of NP-Completeness. W.H. Freeman and Company (1979)
14. Benoit, A., Casanova, H., Rehn-Sonigo, V., Robert, Y.: Resource Allocation for Multiple Concurrent In-Network Stream Processing Applications. Research Report 2009-07, LIP, ENS Lyon, France (February 2009)
15. Rehn-Sonigo, V.: Source Code for the Heuristics and diagrams of all experiments `http://graal.ens-lyon.fr/~vsonigo/code/query-multiapp/`.