



Pipelined fission for stream programs with dynamic selectivity and partitioned state



B. Gedik*, H.G. Özsema, Ö. Öztürk

Department of Computer Engineering, Bilkent University, 06800, Ankara, Turkey

HIGHLIGHTS

- Formalizes the pipelined fission problem for streaming applications.
- Models the throughput of pipelined fission configurations.
- Develops a three-stage heuristic algorithm to quickly locate a close to optimal pipelined fission configuration.
- Experimentally evaluates the solution and demonstrate its efficacy.

ARTICLE INFO

Article history:

Received 19 February 2015

Received in revised form

2 December 2015

Accepted 3 May 2016

Available online 14 May 2016

Keywords:

Data stream processing

Auto-parallelization

Pipelining

Fission

ABSTRACT

There is an ever increasing rate of digital information available in the form of online data streams. In many application domains, high throughput processing of such data is a critical requirement for keeping up with the soaring input rates. Data stream processing is a computational paradigm that aims at addressing this challenge by processing data streams in an on-the-fly manner, in contrast to the more traditional and less efficient store-and-then process approach. In this paper, we study the problem of automatically parallelizing data stream processing applications in order to improve throughput. The parallelization is automatic in the sense that stream programs are written sequentially by the application developers and are parallelized by the system. We adopt the asynchronous data flow model for our work, which is typical in Data Stream Processing Systems (DSPS), where operators often have dynamic selectivity and are stateful. We solve the problem of *pipelined fission*, in which the original sequential program is parallelized by taking advantage of both *pipeline parallelism* and *data parallelism* at the same time. Our pipelined fission solution supports *partitioned stateful* data parallelism with dynamic selectivity and is designed for shared-memory multi-core machines. We first develop a cost-based formulation that enables us to express pipelined fission as an optimization problem. The bruteforce solution of this problem takes a long time for moderately sized stream programs. Accordingly, we develop a heuristic algorithm that can quickly, but approximately, solve the pipelined fission problem. We provide an extensive evaluation studying the performance of our pipelined fission solution, including simulations as well as experiments with an industrial-strength DSPS. Our results show good scalability for applications that contain sufficient parallelism, as well as close to optimal performance for the heuristic pipelined fission algorithm.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

We are experiencing a data deluge due to the ever increasing rate of digital data produced by various software and hardware sensors present in our highly instrumented and interconnected world. This data often arrives in the form of continuous streams.

Examples abound, such as ticker data [41] in financial markets, call detail records [7] in telecommunications, production line diagnostics [3] in manufacturing, and vital signals [35] in healthcare. Accordingly, there is an increasing need to gather and analyze data streams in near real-time, detect emerging patterns and outliers, and take automated action. Data stream processing systems (DSPSs) [37,20,36,1,5] enable carrying out these tasks in a natural way, by taking data streams through a series of analytic operators. In contrast to the traditional store-and-process model of data management systems, DSPSs rely on the process-and-forward model and are designed to provide high throughput and timely response.

* Corresponding author.

E-mail addresses: bgedik@bilkent.edu.tr (B. Gedik), habibe.ozsema@bilkent.edu.tr (H.G. Özsema), ozturk@bilkent.edu.tr (Ö. Öztürk).

<http://dx.doi.org/10.1016/j.jpdc.2016.05.003>

0743-7315/© 2016 Elsevier Inc. All rights reserved.

Since performance is one of the fundamental motivations for adopting the stream processing model, optimizing the throughput of stream processing applications is an important goal of many DSPSs. In this paper, we study the problem of *pipelined fission*, that is automatically finding the best configuration of combined pipeline and data parallelism in order to optimize application throughput. Pipeline parallelism naturally occurs in stream processing applications [39]. As one of the stages is processing a data item, the previous stage can concurrently process the next data item in line. Data parallelization, aka. *fission*, involves replicating a stage and concurrently processing different data items using these replicas. Typically, data parallelism opportunities in streaming applications need to be discovered (to ensure safe parallelization) and require runtime mechanisms, such as splitting and ordering, to enforce sequential semantics [33,15].

Our goal in this paper is to determine how to distribute processing resources among the data and pipeline parallel aspects within the stream program, in order to best optimize the throughput. While pipeline parallelism is very easy to take advantage of, the amount of speed-up that can be obtained is limited by the pipeline depth. On the other hand, data parallelism, when applicable, can be used to achieve higher levels of scalability. Yet, data parallelism has limitations as well. First, the mechanisms used to establish sequential semantics (e.g., ordering) have overheads that increase with the number of replicas used. Second, and more importantly, since data parallelism is applied to a subset of operators within the chain topology, the performance is still limited by other operators for which data parallelism cannot be applied (e.g., because they are arbitrarily stateful). The last point further motivates the importance of pipelined fission, that is the need for performing combined pipeline and data parallelism.

The setting we consider in this paper is multi-core shared-memory machines. We focus on streaming applications that possess a *chain topology*, where multiple stages are organized into a series, each stage consuming data from the stage before and feeding data into the stage after. Each stage can be a *primitive* operator, which is an atomic unit, or a *composite* [22] operator, which can contain a more complex sub-topology within. In the rest of the paper, we will simply use the term operator to refer to a stage. The pipeline and data parallelism we apply are all at the level of these operators.

Our work is applicable to and is designed for DSPSs that have the following properties:

- **Dynamic selectivity:** If the number of input data items consumed and/or the number of output data items produced by an operator are **not** fixed and may change depending on the contents of the input data, the operator is said to have dynamic selectivity. Operators with dynamic selectivity are prevalent in data-intensive streaming applications. Examples of such operators include data dependent filters, joins, and aggregations.
- **Backpressure:** When a streaming operator is unable to consume the input data items as fast as they are being produced, a bottleneck is formed. In a system with backpressure, this eventually results in an internal buffer to fill up, and thus an upstream operator blocks while trying to submit a data item to the full buffer. This is called backpressure, and it recursively propagates up to the source operators.
- **Partitioned processing:** A stream that multiplexes several sub-streams, where each sub-stream is identified by its unique value for the partitioning key, is called a *partitioned stream*. An operator that independently processes individual sub-streams within a partitioned stream is called a *partitioned operator*. Partitioned operators could be stateful, in which case they maintain independent state for each sub-stream. DSPSs that support partitioned processing can apply fission for partitioned stateful operators—an important class of streaming operators [34,4].

There are several challenges in solving the pipelined fission problem we have outlined. First, we need to formally define what a valid parallelization configuration is with respect to the execution model used by the DSPS. This involves defining the restrictions on the mapping between threads and parallel segments of the application. Second, we need to model the throughput as a function of the pipelined fission configuration, so as to compare different pipelined fission alternatives among each other. Finally, even for a small number of operators, processor cores, and threads, there are combinatorially many valid pipelined fission configurations. It is important to be able to quickly locate a configuration that provides close to optimal throughput. There are two strong motivations for this. The first is to have a fast edit–debug cycle for streaming applications. The second is to have low overhead for dynamic pipelined fission, that is being able to update the parallelization configuration at run-time. Note that, the optimal pipelined fission configuration depends on the operator costs and selectivities, which are often data dependent, motivating dynamic pipelined fission. In this paper, our focus is on solving the pipelined fission problem in a reasonable time, with high accuracy with respect to throughput.

Our solution involves three components. First, we define valid pipelined fission configurations based on application of *fusion* and *fission* on operators. Fusion is a technique used for minimizing scheduling overheads and executing stream programs in a streamlined manner [25,13]. In particular, series of operators that form a *pipeline* are fused and executed by a dedicated thread, where buffers are placed between successive pipelines. On the other hand, using fission, series of pipelines that form a *parallel region* are replicated to achieve data parallelism.

Second, we model concepts such as operator compatibility (used to define parallel regions), backpressure (key factor in defining throughput), and system overheads like the thread switching and replication costs (factors impacting the effectiveness of parallelization), and use these to derive a formula for the throughput.

Last, and most importantly, we develop a heuristic algorithm to quickly locate a pipelined fission configuration that provides close to optimal performance. The algorithm relies on three main ideas: The first is to form regions based on the longest compatible sequence principle, where compatible means that a formed region carries properties that make it amenable to data parallelism as a whole. The second is to divide regions into pipelines using a greedy bottleneck resolving procedure. This procedure performs iterative pipelining, using a variable utilization-based upper bound as the stopping condition. The third is another greedy step, which resolves the remaining bottlenecks by increasing the number of replicas of a region.

We evaluate the effectiveness of our solution based on extensive analytic experimentation. We also use IBM's SPL language and its runtime system to perform an empirical evaluation. Our SPL-based evaluation shows that we can quickly locate a pipelined fission configuration that is within 5%–10% of the optimal using our heuristic algorithm.

In summary, we make the following contributions:

- We formalize the pipelined fission problem for streaming applications that are organized as a series of stages and can potentially exhibit dynamic selectivity, backpressure, and partitioned processing.
- We model the throughput of pipelined fission configurations and cast the problem of locating the best configuration as a combinatorial optimization one.
- We develop a three-stage heuristic algorithm to quickly locate a close to optimal pipelined fission configuration and evaluate its effectiveness using analytical and empirical experiments.

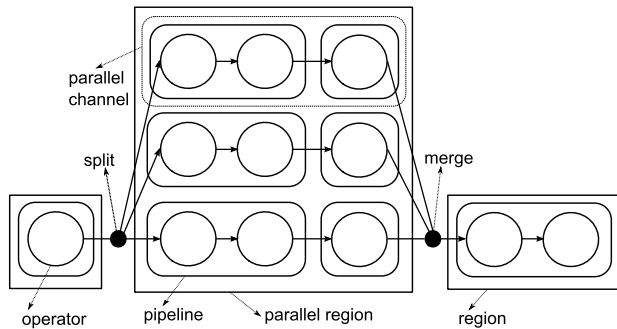


Fig. 1. Pipelined fission terminology.

2. Background

In this section, we summarize the terminology used for the pipelined fission problem, and outline a system execution model that will guide the problem formulation and solution used in the rest of the paper.

2.1. Terminology and definitions

A *stream graph* is a set of operators connected to each other via *streams*. As mentioned earlier, we consider graphs with *chain topology* in this work. Fig. 1 summarizes the terminology used to define our pipelined fission problem.

There are two operator properties that play an important role in pipelined fission, namely *selectivity* and *state*.

- Selectivity of an operator is the number of items it produces per number of items it consumes. It could be less than one, in which case the operator is *selective*; it could be equal to one, in which case the operator is *one-to-one*; or it could be greater than one, in which case the operator is *prolific*.
- State specifies whether and what kind of information is maintained by the operator across firings. An operator could be *stateless*, in which case it does not maintain any state across firings. It could be *partitioned stateful*, in which case it maintains independent state for each sub-stream determined by a *partitioning key*. Finally, an operator could be *stateful* without a special structure.

We name a series of operators fused together as a *pipeline*. A series of pipelines replicated as a whole is called a *parallel region*. Series of pipelines that fall between parallel regions form simple *regions*. Each replica within a parallel region is called a *parallel channel*. A parallel channel contains replicas of the pipelines and operators of a parallel region. In order to maintain sequential program semantics under selective operators, *split* and *merge* operations are needed before and after a parallel region, respectively. The split operation assigns sequence numbers to tuples and distributes them over the parallel channels, such as a hash-based splitter for a partitioned stateful parallel region. The merge operation unions tuples from different parallel channels and orders them based on their sequence numbers. A parallel region cannot contain a stateful (non-partitioned) operator [33] and thus such regions are formed by stateless and partitioned stateful operators.

Listing 1 shows a toy SPL (Stream Processing Language) [21] application for illustrating some of the concepts introduced. This application counts the number of appearances of each word in a file. It consists of 4 operators organized into a chain. The first operator, named `Lines`, is a source operator producing lines of text. The second operator, named `Words`, divides each line into words, and outputs one tuple for each word. Note that this operator has a selectivity value over 1. The exact selectivity value is not known at development time, as it is dependent on the data. As

```

stream<rstring line> Lines = FileSource() {
  param file: "in.txt";
}
stream<rstring word> Words = Custom(Lines) {
  logic
  onTuple Lines:
    for (rstring word in tokenize(line, "\t", false))
      submit({word=word}, Words);
  onPunct Lines:
    submit(currentPunct(), Words);
}
stream<rstring word, uint32 count> Counts = Aggregate(Words) {
  window
  Words: tumbling, punct(), partitioned;
  param
  partitionBy: word;
  output
  Counts: count = Sum(1u);
}
() as Results = FileSink(Counts) {
  param file: "counts.txt";
}

```

Listing 1: Sample SPL application.

such, this stream program does not follow the synchronous data flow model and cannot be scheduled statically at compile-time. The `Words` operator is stateless, as it does not maintain state across tuple firings. The next operator in line is the `Counts` operator, which performs a simple *Sum* aggregation over the window of tuples. Importantly, this operator is partitioned stateful and it is highly selective. Finally, the last operator is named `Results`, which is a file sink. In this application the *pipeline* formed by the `Words` and `Counts` operators can be made into a *parallel region*. Since the `Counts` operator is partitioned on the `word` attribute, we should have a hash-based *split* before the parallel region, and a re-ordering *merge* after it.

2.2. Execution model

A distributed stream processing middleware typically executes data flow graphs by partitioning them into basic units called *processing elements*. Each processing element contains a sub-graph and can run on a different host. For small and medium-scale applications, the entire graph can be mapped to a single processing element. Without loss of generality, in this paper we focus on a single multi-core host executing the entire graph. Our pipelined fission technique can be applied independently on each host when the whole application consists of multiple distributed processing elements.

In this paper, we follow an execution model based on the SPL runtime [20], which has been used in a number of earlier studies as well [39,33,15,32,13]. In this model, there are two main sources of threading, which contribute to the execution of the stream graph. The first one is *operator threads*. Source operators, which do not have any input ports, are driven by their own operator threads. When a source operator makes a submit call to send a tuple to its output port, this same thread executes the rest of the downstream operators in the stream graph. As a result, the same thread can traverse a number of operators, before eventually coming back to the source operator to execute the next iteration in its event loop. This behavior is because the stream connections in a processing element are implemented via function calls. Using function calls yields fast execution, avoiding scheduler context switches and explicit buffers between operators. This optimization is known as *operator fusion* [25,13].

The second source of threading is *threaded ports*. Threaded ports can be inserted at any operator input port. When a thread reaches a threaded port, it inserts the tuple at hand into the threaded port buffer, and goes back to executing upstream logic. A separate




Fig. 2. A chain topology with 5 operators.

thread, dedicated to the threaded port, picks up the queued tuples and executes the downstream operators. In pipelined fission, we use threaded ports to ensure that *each pipeline is run by a separate thread*. For instance, in Fig. 1 there are 8 threads. The scheduling of threads to the processor cores is left to the operating system.

The goal of our pipelined fission solution is to automatically determine a parallelization configuration, that is the pipelines, regions, and number of replicas, so as to maximize the throughput. It is important to note that our solution is designed for asynchronous data flow systems [36,31,38,20,37,5,1] that support dynamic selectivity and partitioned stateful operators. In contrast, synchronous data flow systems (SDF) [16,28] assume that the relative flow rates (selectivities) are specified at development time. They produce a static schedule at compile-time, which is executed at runtime. We assume a programming model that does not require specification of selectivities and accordingly, a runtime system that does not rely on a static schedule. We emphasize the speed of finding a parallelization configuration, as this has to be performed at runtime.¹ We model backpressure in our solution in order to handle rate differences at runtime. Furthermore, data parallelism in SDF systems is limited to stateless operators. Our work supports partitioned stateful operators [34,4], which are typical in asynchronous data flow systems.

3. Problem formulation

In this section, we model the pipelined fission problem and present a brute-force approach to find a parallelization configuration that maximizes the throughput.

3.1. Application model

We start with modeling the topology, the operators, and the parallelization configuration.

Topology. We consider applications that have a chain topology. The operators that participate in the chain can be composite and have more complex topologies within, as long as they fit into one of the operator categories described below.

Let $O = \{o_i \mid i \in [1..N]\}$ be the set of operators in the application. Here, $o_i \in O$ denotes the i th operator in the chain. o_1 is the *source* operator and o_N is the *sink* operator. For $1 < i \leq N$, operator o_i has o_{i-1} as its *upstream* operator and for $1 \leq i < N$, o_i has o_{i+1} as its *downstream* operator. Fig. 2 shows an example chain topology with $N = 5$ operators.

Operators. For $o \in O$, $k(o) \in \{f, p, s\}$ denotes the operator kind: f is for *stateful*, p is for *partitioned stateful*, and s is for *stateless*. For a partitioned stateful operator o (that is $k(o) = p$), $a(o)$ specifies the partitioning key, which is a set of stream attributes. $s(o)$ denotes the *selectivity* of an operator, which can go over 1 for *prolific* operators—operators that can produce one or more tuples per input tuple consumed. As an example, the `Words` operator from Listing 1 is a prolific operator and the operator `Counts` is a partitioned stateful operator with a partitioning attribute of word and selectivity of less than 1. We use $c(o)$ to denote the per-tuple cost of an operator.

For $o \in O$, $f(o) : \mathbb{N}^+ \rightarrow \mathbb{R}$ is a *base scalability function* for operator o . Here, $f(o)(x) = y$ means that x copies of operator o will raise the throughput to y times the original, assuming no parallelization overhead. We have $k(o_i) = s \Rightarrow f(o_i) = f_i$, where f_i is the linear scalability function, that is $f_i(x) = x$. In other words, for stateless operators, the base scalability function is linear. For partitioned stateful operators, including parallel sources and sinks, bounded linear functions are more common, such as:

$$f_b(x; u) = \begin{cases} x & \text{if } x \leq u \\ u & \text{otherwise.} \end{cases}$$

Here, $f_b(\cdot; u)$ is a bounded linear scalability function, where u specifies the maximum scalability value. For partitioned stateful operators, the size of the partitioning key's domain could be a limiting factor on the scalability that could be achieved. For parallel sources and sinks, the number of distinct external sources and sinks could be a limiting factor (e.g., number of TCP/IP end points, number of data base partitions, etc.).

Parallelization configuration. Let us denote the set of threads used to execute the stream program as $T = \{t_i \mid i \in [1..|T|]\}$. The number of replicas for operator $o \in O$ is denoted by $r(o) \in \mathbb{N}^+$. Note that, we have $k(o) = f \Rightarrow r(o_i) = 1$, as stateful operators cannot be replicated.

Let us denote the j th replica of an operator o_i as $o_{i,j}$ and the set of all operator replicas as $V = \{o_{i,j} \mid o_i \in O \wedge j \in [1..r(o_i)]\}$. We define $m : V \rightarrow T$ as the *operator to thread mapping* that assigns operator replicas to threads. $m(o_{i,j}) = t$ means that operator o_i 's j th replica is assigned to thread t . An operator is assigned to a single thread, but multiple operators can be assigned to the same thread. There are a number of rules about this mapping that restrict the set of possible mappings to those that are consistent with the execution model we have outlined earlier. We first define additional notation to formalize these rules.

Given $O' \subseteq O$, we define a Boolean predicate $L(O')$ that captures the notion of a *sequence* of operators. Formally, $L(O') \equiv \{o_{i_1}, o_{i_2}\} \subset O' \Rightarrow \forall_{i_1 \leq i_2}, o_{i_1} \in O'$. There are two kinds of sequences we are interested in. The first one is called a *non-replicated sequence* and is defined as $L_s(O', r) \equiv L(O') \wedge \forall_{o \in O'}, r(o) = 1$. In a non-replicated sequence, all operators have a single replica. The second is called a *replicated sequence* and is defined as $L_p(O', r) \equiv L(O') \wedge \forall_{o \in O'}, (k(o) \neq f \wedge r(o) = 1) \wedge \bigcap_{o \in O', k(o)=p} a(o) \neq \emptyset$. That is, a group of operators are considered a replicated sequence if and only if they form a sequence, they do not include a stateful operator, they all have the same number of replicas, and if there are any partitioned stateful operators in the sequence, they have compatible partitioning keys.² We will drop r , that is the function that defines the replica counts for the operators, from the parameter list of the sequence defining predicates, L_s and L_p , when it is obvious from the context.

With these definitions, we list the following rules for the operator replica to thread mapping function, m :

- $m(o_{i_1, j_1}) = m(o_{i_2, j_2}) \Rightarrow j_1 = j_2$. I.e., operator replicas from different *channels* are not assigned to the same thread. Here, channel corresponds to the replica index.
- $t = m(o_{i_1, j}) = m(o_{i_2, j}) \Rightarrow \exists O' \text{ s.t. } \{o_{i_1}, o_{i_2}\} \subseteq O' \subset O \wedge (L_s(O') \vee L_p(O')) \wedge (\forall_{o_i \in O'}, m(o_{i,j}) = t)$. I.e., if two operator replicas are assigned to the same thread, they must be part of a replicated or non-replicated sequence and all other operator replicas in between these two on the same channel should be assigned to the same thread.

¹ It has been shown that lightweight profiling can be used to determine selectivities and costs at runtime [39].

² In practice, there is also the requirement that these keys are forwarded by the other operators in the sequence [33], but such details do not impact our modeling.




Fig. 3. Regions and pipelines.

- $m(o_{i_1,j}) = m(o_{i_2,j}) \Rightarrow \forall l \in [1..r(o_{i_1})]$, $m(o_{i_1,l}) = m(o_{i_2,l})$. I.e, if two operator replicas are assigned to the same thread, their sibling operator replicas should share their threads as well. For instance, if $o_{1,1}$ and $o_{2,1}$ both map to t_1 , then their siblings $o_{1,2}$ and $o_{2,2}$ should share the same thread, say t_2 .

Regions and Pipelines. The above rules divide the program into *regions* and these regions into sub-regions that we call *pipelines*, as shown in the example in Fig. 3.

In this example, we have 3 parallel regions: $P_1 = \{P_{1,1}, P_{1,2}\}$, $P_2 = \{P_{2,1}, P_{2,2}\}$, and $P_3 = \{P_{3,1}\}$. The first region P_1 has a single replica, that is $r(P_1) = 1$ and it consists of two pipelines, namely $P_{1,1}$ and $P_{1,2}$. The first pipeline has a single operator inside, whereas the second one has two operators. Concretely, we have $P_{1,1} = \{o_1\}$ and $P_{1,2} = \{o_2, o_3\}$. The second region has $r(P_2) = 3$, as there are 3 parallel channels, and it consists of 2 pipelines, namely $P_{2,1}$ and $P_{2,2}$. We have $P_{2,1} = \{o_4, o_5\}$ and $P_{2,2} = \{o_6\}$. Finally, the third region is $P_3 = \{P_{3,1}\}$, where $r(P_3) = 1$ and $P_{3,1} = \{o_7\}$.

Given the thread mapping function m and the replica function r , the set of regions formed is denoted by $\mathcal{P}(m, r)$ or \mathcal{P} for short. To find the first region, $P_1 \in \mathcal{P}$, we start from the source operator o_1 and locate the *longest* sequence of operators $O' \subset O$ s.t. $o_1 \in O' \wedge (L_s(O') \vee L_p(O'))$. We can apply this process successively, starting from the next operator in line that is not part of the current set of regions, until the set of all regions, \mathcal{P} , is formed. The pipelines in a given region are formed by grouping operators whose replicas for a parallel channel are assigned to the same thread by the mapping m .

For each pipeline $P_{i,j} \in P_i \in \mathcal{P}$, there are $r(P_i)$ replicas and a different thread executes each pipeline replica. Then the total number of threads used is given by $\sum_{P_i \in \mathcal{P}} r(P_i) \cdot |P_i|$. In the example above, we have 9 threads and 13 operator replicas.

3.2. Modeling the throughput

Our goal is to define the throughput of a given configuration \mathcal{P} . Once the throughput is formulated, we can cast our problem as an optimization one, where we aim to find the thread mapping function (m) and the operator replica counts (r) that maximize the throughput.

To formalize the throughput, we start with a set of helper definitions. We denote the kind of a region as $k(P_i)$, and define:

$$k(P_i) = \begin{cases} f & \text{if } \exists o_k \in P_{i,j} \in P_i \text{ s.t. } k(o_k) = f \\ s & \text{if } \forall o_k \in P_{i,j} \in P_i \ k(o_k) = s \\ p & \text{otherwise.} \end{cases} \quad (1)$$

For instance, in Fig. 3, if o_4 and o_6 are stateless, and o_5 is partitioned stateful, then the parallel region P_2 becomes partitioned stateful (p). As another example, if o_2 is stateful, then the region P_1 becomes stateful.

We denote the selectivity of a pipeline $P_{i,j}$ as $s(P_{i,j}) = \prod_{o_k \in P_{i,j}} s(o_k)$, the selectivity of a region P_i as $s(P_i) = \prod_{P_{i,j} \in P_i} s(P_{i,j})$,

and the selectivity of the entire flow \mathcal{P} as $s(\mathcal{P}) = \prod_{P_i \in \mathcal{P}} s(P_i)$. We denote the cost of a pipeline as $c(P_{i,j})$ and define it as:

$$c(P_{i,j}) = \sum_{o_k \in P_{i,j}} s_k(P_{i,j}) \cdot c(o_k). \quad (2)$$

Here, $s_k(P_{i,j}) = \prod_{o_l \in P_{i,j}, l < k} s(o_l)$ is the selectivity of the sub-pipeline up to and excluding operator o_k .

Region throughput. We first model a region's throughput in isolation, assuming no other regions are present in the system. Let $R(P_i)$ denote the maximum input throughput supported by a region under this assumption. And let $R_j(P_i)$ denote the output throughput of the first j pipelines in the region assuming the remaining pipelines have zero cost. Furthermore, let $R(P_{i,j})$ denote the input throughput of the pipeline $P_{i,j}$ if all other pipelines had zero cost (making it the bottleneck of the system).

We have $R_0(P_i) = \infty$ and also for $j > 0$:

$$R_j(P_i) = \begin{cases} s(P_{i,j}) \cdot R_{j-1}(P_i) & \text{if } R_{j-1}(P_i) < R(P_{i,j}) \\ s(P_{i,j}) \cdot R(P_{i,j}) & \text{otherwise.} \end{cases} \quad (3)$$

In essence, Eq. (3) models backpressure. If the input throughput of a pipeline, when considered alone, is higher than the output throughput of the sub-region formed by the pipelines before it, then the latter throughput is used to compute the pipeline's output throughput when it is added to the sub-region. This represents the case when the pipeline in question is not the bottleneck. The other case is when the pipeline's input throughput, when considered alone, is lower than the output throughput of the sub-region formed by the pipelines before it. In this case, the former throughput is used to compute the pipeline's output throughput when it is added to the sub-region. This represents the case when the pipeline in question is the bottleneck within the sub-region. Modeling the backpressure is important, as most real-world data stream processing systems rely on it. As a concrete example, the lack of back-pressure in the popular open-source stream processing system Storm [36] has resulted in the development of Heron [29].

The throughput of a pipeline by itself, that is $R(P_{i,j})$, can be represented as:

$$R(P_{i,j}) = (c(P_{i,j}) + h(P_{i,j}))^{-1}, \quad (4)$$

where $h(P_{i,j})$ is the cost of switching threads between sub-regions, defined as:

$$h(P_{i,j}) = \delta \cdot (\mathbf{1}(j > 1) + \mathbf{1}(j < |P_i|) \cdot s(P_{i,j})). \quad (5)$$

Here, δ is the *thread switching overhead* due to the queues involved in-between. The input overhead is incurred for the pipelines except the first one, and the output overhead is incurred for the pipelines except the last one.

With these definitions at hand, we can define the input throughput $R(P_i)$ as the output throughput of the region divided by the region's selectivity. That is:

$$R(P_i) = R_{|P_i|}(P_i) / s(P_i). \quad (6)$$

Parallel region throughput. The next step is to compute the throughput of a parallel region. For that purpose, we first define an *aggregate scalability function* $f\langle P_i \rangle$ for the region P_i as:

$$f\langle P_i \rangle(x) = \min_{o_k \in P_i, j \in P_i} f\langle o_i \rangle(x). \quad (7)$$

The aggregate scalability function for a region simply takes the smallest scalability value from the scalability functions of the constituent operators within the region.

We denote the *parallel throughput* of a region P_i as $R^*(P_i)$ and define it as follows:

$$R^*(P_i) = \left(c_p \cdot \log_2(r(P_i)) + \frac{1}{R(P_i) \cdot f\langle P_i \rangle(r(P_i))} \right)^{-1}. \quad (8)$$

Here, c_p is the *replication cost factor* for a parallel region. Recall that a parallel region needs to reorder tuples. In the presence of selectivity, this often requires attaching sequence numbers to tuples and re-establishing order at the end of the parallel region. The re-establishment of order takes time that is logarithmic in the number of channels, per tuple. However, such processing typically has a low constant compared to the cost of the operators.

Let $R^+(P_i)$ be the parallel throughput of the region when it is considered within the larger topology that contains the other regions, albeit assuming that all other regions have zero cost. We have:

$$R^+(P_i) = \left(h(P_i) + \frac{1}{R^*(P_i)} \right)^{-1}. \quad (9)$$

Here, $h(P_i)$ is the cost of switching threads between regions, which can be expressed as:

$$h(P_i) = \delta \cdot (\mathbf{1}(i > 1) + \mathbf{1}(i < |\mathcal{P}|) \cdot s(P_i)). \quad (10)$$

Throughput of a program. Given these definitions, we are ready to define the input throughput of a program, denoted as $R(\mathcal{P})$. We follow the same approach as we did for regions formed out of pipelines.

Let us define the output throughput of the first k regions as $R_k(\mathcal{P})$, assuming the downstream regions have zero cost. We have $R_0(\mathcal{P}) = \infty$, and for $i > 0$:

$$R_i(\mathcal{P}) = \begin{cases} s(P_i) \cdot R_{i-1}(\mathcal{P}) & \text{if } R_{i-1}(\mathcal{P}) < R^+(P_i) \\ s(P_i) \cdot R^+(P_i) & \text{otherwise.} \end{cases} \quad (11)$$

By dividing the output throughput of the program to its selectivity, we get:

$$R(\mathcal{P}) = R_{|\mathcal{P}|}(\mathcal{P})/s(\mathcal{P}). \quad (12)$$

Bounded throughput. So far we have computed the *unbounded* throughput. In other words, we have assumed that each thread has a core available to itself. However, in practice, there could be more threads than the number of cores available. For instance, replicating a region with 3 pipelines 3 times will result in 9 threads, but the system may only have 8 cores. However, replicating the region 2 times will result in an underutilized system that has only 6 threads and thus not all cores can be used.

Let C denote the number of cores in the system. We denote the *bounded* throughput of a program with parallelization configuration of m (the thread mapping function) and r (the operator replica counts) as $R(\mathcal{P}(m, r), C)$. The bounded throughput is simply computed as the unbounded throughput divided by the *utilization* times the number of cores. Formally,

$$R(\mathcal{P}, C) = R(\mathcal{P}) \cdot \frac{C}{U(\mathcal{P})}. \quad (13)$$

Here $U(\mathcal{P})$ is the utilization for the unbounded throughput. Eq. (13) simply scales the unbounded throughput by multiplying




Fig. 4. # of parallel program configurations as a function of # of threads (M) and # of operators ($|O|$).

it with the ratio of the maximum utilization that can be achieved (which is C) to the unbounded utilization. We assume that the cost due to scheduling of threads by the operating system is negligible. For instance, if the unbounded throughput is 3 units, but results in a utilization value of 6 and the system has only 4 cores, then the bounded throughput is given by $3 \cdot (4/6) = 2$ units.

The computation of the utilization, $U(\mathcal{P})$, is straightforward. We already have a formula for the input throughput of the program, which can be used to compute the input throughputs of the parallel regions and the pipelines. Multiplying input throughputs of the pipelines with the pipeline costs would give us the utilization, after adding the overheads for the thread switching and scalability. Overall utilization can be expressed as:

$$U(\mathcal{P}) = R(\mathcal{P}) \cdot \sum_{1 \leq i < j} s_i(\mathcal{P}) \cdot (c_p \cdot \log_2(r(P_i)) + h(P_i) + \sum_{1 \leq j < |P_i|} s_j(P_i) \cdot (h(P_{i,j}) + c(P_{i,j}))). \quad (14)$$

Here, $s_j(P_i) = \prod_{1 \leq l < j} s(P_{i,l})$ is the selectivity of the region P_i up to and excluding the j th pipeline, and $s_j(\mathcal{P}) = \prod_{1 \leq l < j} s(P_l)$ is the selectivity of the program \mathcal{P} up to and excluding the j th region.

Optimization. Our ultimate goal is to find $\text{argmax}_{m,r} R(\mathcal{P}(m, r), C)$, where m is subject to the rules we have outlined earlier. One way to solve this problem is to combinatorially generate all possible parallel configurations. This can be achieved via a recursive procedure that takes the maximum number of threads M and the set of operators O as input, and generates all valid parallelization configurations of the operators that uses at most M threads. Let us denote the set of configurations generated by such a generator as $D(O, M)$. Then we can compute $\text{argmax}_{(m,r) \in D(O,M)} R(\mathcal{P}(m, r), C)$ as the optimal configuration. There are two problems with this approach. First, and the more fundamental one, is that, the computation of $D(O, M)$ takes a very long time even for a small number of threads and operators; and this time grows exponentially, since the number of variations increases exponentially (both with increasing number of operators and maximum number of threads).

Fig. 4 shows the number of parallel configurations as a function of the number of operators and the maximum number of threads used. Second, we need to pick a reasonable value for M , which is typically greater than C . It can be taken as a constant times the number of cores, that is $k \cdot C$. Unfortunately, using a large constant will result in an excessively long running time for the configuration generation algorithm. On the other hand, using a small constant will have the risk of finding a sub-optimal solution.

Algorithm 1: PIPELINEDFISSION(O, M, α)

Data: O : operators (with their costs, c ; selectivities, s ; and state kinds, k), M : number of cores, α : fusion cost threshold
Result: Pipelined fission configuration
 $R \leftarrow \text{CONFIGUREREGIONS}(O, \alpha)$ \triangleright Configure regions
 $t \leftarrow 0$; $P \leftarrow \emptyset$; $r \leftarrow \emptyset$ \triangleright Initialize best settings
for $s \in 0.1 \cdot [0..10]$ **do** \triangleright Range of utilization scalers
 $P' \leftarrow \text{CONFIGUREPIPELINES}(R, s \cdot M)$ \triangleright Configure pipelines
 $r' \leftarrow \text{CONFIGUREREPLICAS}(R, P', M)$ \triangleright Configure # of replicas
 $t' \leftarrow \text{COMPUTETPUT}(R, P', r')$ \triangleright Compute the throughput
 if $t' > t$ **then** $t \leftarrow t'$; $P \leftarrow P'$; $r \leftarrow r'$
return (R, P, r) \triangleright Return the final configuration

4. Heuristic solution

In this section we present an algorithm to quickly solve the pipelined fission problem that was formalized in Section 3. Our algorithm is heuristic in nature and trades off throughput optimality to achieve reasonable performance in terms of solution time. Despite this, our results, presented later in Section 5, show that not only does our algorithm achieve close to optimal throughput, but also it outperforms optimal versions of fission-only and pipelining-only alternatives.

4.1. Overview

Algorithm 1 presents our solution, which consists of three phases, namely (i) *region configuration*, (ii) *pipeline configuration*, and (iii) *replica configuration*. The region configuration phase divides the chain of operators into chains of regions. This is done based on the compatibility of the successive operators in terms of their state, while avoiding the creation of small regions that cannot achieve effective parallelization. The second and third phases are used to configure pipeline and data parallelism, respectively. That is, pipeline configuration creates pipelines within regions, and replica configuration determines the number of replicas for the regions. These two phases are run multiple times, each time with a different amount of CPU utilization reserved for them, but always summing up to the number of CPUs available in the system. In particular, we range the fraction of the CPU utilization reserved for pipelining from 0% to 100%, in increments of 10%. The reason for running the pipeline and region configuration phases with differing shares of CPU utilization is that, we do not know, a priori, how much parallelism is to be reserved for pipelining versus how much for fission, in order to achieve the best performance with respect to throughput. Among the multiple runs of the second and the third phases, we pick the one that gives the highest throughput as our final pipelined fission solution. Internally, pipeline configuration phase and replica configuration phase work similarly. In pipeline configuration, we repeatedly locate the bottleneck pipeline and divide it. In replica configuration, we repeatedly locate the bottleneck region and increase its replica count. In what follows, we further detail the three phases of the algorithm.

4.2. Region configuration

Algorithm 2 presents the region configuration phase, where we divide the chain of operators into a chain of regions. The algorithm consists of two parts. In the first part, we form effectively parallelizable regions. This may leave out some operators unassigned. In the second phase, we merge the consecutive unassigned operators into regions as well.

The first for loop in Algorithm 2 represents the first step. We form regions by iterating over the operators. We keep accumulating operators into the current region, as long as the operators are

Algorithm 2: CONFIGUREREGIONS(O, α)

Data: O : operators, α : fusion cost threshold
Result: Regions
 $R \leftarrow \{\}$ \triangleright The list of regions that will hold the final result
 $C \leftarrow \{\}$ \triangleright The list of operators in the current potential region
for $i \leftarrow 1$; $i \leq |O|$; $i \leftarrow i + 1$ **do** \triangleright For each operator
 \triangleright The current region borrows the operator's properties
 if $k(o_i) \neq f \wedge (|C| = 0 \vee k(C) = s)$ **then**
 $k(C) \leftarrow k(o_i)$ \triangleright Update the region's kind
 if $k(o_i) = p$ **then** \triangleright If o_i is partitioned
 $a(C) \leftarrow a(o_i)$ \triangleright Update current region's key
 \triangleright The current region stays partitioned, possibly with a broadened key
 else if $k(o_i) = p \wedge a(o_i) \subseteq a(C)$ **then**
 $a(C) \leftarrow a(o_i)$ \triangleright Update current region's key
 \triangleright The current region and the operator are incompatible
 else if $k(o_i) \neq s$ **then**
 if $\text{COMPUTECOST}(C) > \alpha$ **then** \triangleright Region is costly enough
 $R \leftarrow R \cup \{C\}$ \triangleright Materialize the region in R
 $d(o) \leftarrow 1, \forall o \in C$ \triangleright Mark region's operators as assigned
 $C \leftarrow \{\}$ \triangleright Reset the current region
 if $k(o_i) \neq f$ **then** \triangleright Parallelizable operator
 $i \leftarrow i - 1$ \triangleright Redo iteration with empty current region
 continue
 $C \leftarrow C \cup \{o_i\}$ \triangleright Add the operator to the current region
 \triangleright Handle the pending region at loop exit
 if $\text{COMPUTECOST}(C) > \alpha$ **then** \triangleright Region is costly enough
 $R \leftarrow R \cup \{C\}$ \triangleright Materialize the region in R
 $d(o) \leftarrow 1, \forall o \in C$ \triangleright Mark region's operators as assigned
 \triangleright Merge all consecutive unassigned ops to a region
 $C \leftarrow \{\}$ \triangleright Reset the current region
 for $i \leftarrow 1$; $i \leq |O|$; $i \leftarrow i + 1$ **do** \triangleright For each operator
 if $d(o_i) = 1 \wedge |C| > 0$ **then** \triangleright We have a complete run
 $R \leftarrow R \cup \{C\}$ \triangleright Materialize the region in R
 $C \leftarrow \{\}$ \triangleright Reset the current region
 else \triangleright Run of unassigned operators continues
 $C \leftarrow C \cup \{o_i\}$ \triangleright Add the operator to the current region
return R \triangleright The final set of regions

not stateful or incompatible. Stateless operators are always compatible with the current region. Partitioned stateful operators are only compatible if their key is the same as the key of the active region so far, or broader (has less attributes). In the latter case, the region's key is updated accordingly. When an incompatible operator is encountered, the current region that is formed so far is completed. However, this region is discarded if its overall cost is below the *fusion cost threshold*, α . The motivation behind this is that, if a region is too small in terms of its cost, parallelization overhead will dominate and effective parallelization is not attainable.

Once a region is completed, the algorithm continues with a fresh region, starting from the next operator in line (the one that ended the formation of the former region). The first step of the algorithm ends, when all operators are processed. In the second step, the operators that are left without a region assignment are handled. Such operators are either stateful or cannot form a sufficiently costly region with the other operators around them. In the second step, consecutive operators that are not assigned a region are put into their own region. However, these regions cannot benefit from parallelization in the pipeline and replica configuration phases that are described next.

4.3. Pipeline configuration

Algorithm 3 describes the pipeline configuration phase. We start with each region being a pipeline and iteratively split

Algorithm 3: CONFIGUREPIPELINES(R, M)

Data: R : regions, M : number of cores
Result: Pipeline configuration
 $P \leftarrow R$ \triangleright Initialize the set of pipelines to regions
 \triangleright Find the bottleneck pipeline (C), and compute the total utilization (U)
 $\langle C, U \rangle \leftarrow \text{FINDBOTTLENECKPIPELINE}(R, P)$
while $U \leq M$ **do** \triangleright System is not fully utilized
 \triangleright Find the best split for the pipeline (maximizes throughput)
 $o_i \leftarrow \text{argmax}_{o_k \in C} \text{COMPUTETPUT}(\{o_j \in C \mid j < k\}, \{o_j \in C \mid j \geq k\})$
 $C_0 \leftarrow \{o_j \in C \mid j < i\}$ \triangleright First half of the best split
 $C_1 \leftarrow \{o_j \in C \mid j \geq i\}$ \triangleright Second half of the best split
if $\text{COMPUTETPUT}(\{C_0, C_1\}) \leq \text{COMPUTETPUT}(C)$ **then**
 break \triangleright No further improvement is possible
 $P \leftarrow P \setminus \{C\} \cup \{C_0, C_1\}$ \triangleright Split the pipeline
 $\langle C, U \rangle \leftarrow \text{FINDBOTTLENECKPIPELINE}(P)$ \triangleright Re-eval. for next iter.
return P \triangleright The final set of pipelines

Algorithm 4: CONFIGUREREPLICAS(R, P, M)

Data: R : regions, P : pipelines, M : number of cores
Result: Set of number of replicas of each region
 $r[C] \leftarrow 1, \forall C \in R$ \triangleright Initialize the replica counts to 1
 \triangleright Find the bottleneck region (C), and compute the total utilization (U)
 $\langle C, U \rangle \leftarrow \text{FINDBOTTLENECKREGION}(R, P, r)$
while $U \leq M$ **do** \triangleright System is not fully utilized
 $t \leftarrow \text{CALCULATETPUT}(R, P, r)$ \triangleright Baseline throughput
 $r[C] \leftarrow r[C] + 1$ \triangleright Increase the channel count
 if $t \geq \text{CALCULATETPUT}(R, P, r)$ **then** \triangleright Throughput decreased
 $r[C] \leftarrow r[C] - 1$ \triangleright Revert back
 break \triangleright No further improvement is possible
 $\langle C, U \rangle \leftarrow \text{FINDBOTTLENECKREGION}(R, P)$ \triangleright Re-eval. for next iter.
return r \triangleright Return the replica counts

the bottleneck pipeline. The FINDBOTTLENECKPIPELINE procedure is used to find the bottleneck pipeline. This procedure simply computes the *unbounded* throughput of the program as new pipelines are successively added, using the formalization from Section 3, and selects the last pipeline that resulted in a reduction in the unbounded throughput as the bottleneck one. It then reports this bottleneck pipeline, together with the utilization of the current configuration. If the utilization is above or equal to the total utilization reserved for the pipeline configuration phase (recall Algorithm 1), then the iteration is terminated and the pipeline configuration phase is over. Otherwise, i.e., if there is room available for an additional pipeline, we find the best split within the bottleneck pipeline. This is done by considering each operator as a split point and picking the split that provides the highest unbounded throughput. However, if the unbounded throughput of this split configuration of two consecutive pipelines is lower compared to the original single pipeline (which might happen for low cost pipelines due to the impact of thread switching overhead), we again terminate the pipeline configuration phase. This is because, if the bottleneck pipeline cannot be improved, then no overall improvement is possible.

4.4. Replica configuration

Algorithm 4 describes the replica configuration phase. It is similar in structure to the pipeline configuration phase. However, it works on regions, rather than pipelines. It iteratively finds the bottleneck region and increases its replica count. The FINDBOTTLENECKREGION procedure is used to find the bottleneck region. This

procedure simply computes the *unbounded* throughput of the program as new regions are successively added, using the formalization from Section 3, and selects the last region that resulted in a reduction in the unbounded throughput as the bottleneck one. It then reports this bottleneck region, together with the utilization of the current configuration. If the utilization is above the number of CPUs available, then the iteration is terminated and the replica configuration phase is over. Otherwise, i.e., if there is room available for an additional parallel channel, we increment the replica count of the bottleneck region. However, if the unbounded throughput of this parallel region with an incremented replica count has a lower unbounded throughput compared to the original parallel region (which might happen for low cost regions due to the impact of replication cost factor and thread switching overhead), we again terminate the region configuration phase. This is because if the bottleneck region cannot be improved, then no overall improvement is possible.

5. Evaluation

In this section, we evaluate our heuristic solution and showcase its performance in terms of the achieved throughput, as well as the time it takes to locate a parallelization configuration. We perform two kinds of experiments. First, we evaluate our pipelined fission solution using model-based experiments under varying workload and system settings. Second, we evaluate our algorithm using stream programs written in IBM's SPL language [20] and executed using the IBM InfoSphere Streams [12] runtime.

In our experiments, we compare our solution against four different approaches, namely: optimal, sequential, fission-only, and pipelining-only.

- *Sequential* solution is the configuration with no parallelism.
- *Optimal* solution is the configuration that achieves the maximum throughput among all possible parallel configurations.
- *Fission-only optimal* solution is the configuration that achieves the highest throughput among all possible parallel configurations that do not involve pipeline parallelism (that is, each parallel channel is executed by a single thread).
- *Pipelining-only optimal* solution is the configuration with the highest throughput among all possible parallel configurations that do not involve data parallelism.

5.1. Experimental setup

For the model-based experiments, we used the analytical model presented in Section 3 to compare alternative solutions. The five alternative solutions we study were all implemented in Java. The SPL experiments rely on the parallelization configurations generated by these solutions to customize the runtime execution of the SPL programs. The SPL programs are compiled down to C++ and executed on the Streams runtime [12].

We describe the experimental setup for the model based experiments in Table 1. Each model based experiment was repeated 1000 times, whereas SPL based experiments were repeated 50 times. All experiments were executed on a Linux system with 2 Intel Xeon E5520 2.27 GHz CPUs with a total of 12 cores and 48 GB of RAM.

We discuss the thread switching overhead and replication cost factor for the SPL experiments later in Section 5.3.

5.2. Model-based experiments

Streaming applications contain operators with diverse properties. Accordingly, the throughput of the topology is highly dependent on the properties of the operators involved. Hence, we evaluate our solution by varying operator selectivity, operator

Table 1
Experimental parameters: default values and ranges for model based experiments.

Name	Range	Default value
Operator cost mean	[50, 250]	200
Operator cost stdev	–	100
Number of operators	[1, 8]	8
Number of cores	[1, 12]	4
Selectivity mean	[0.1, 1]	0.8
Selectivity stdev	–	0.4
Stateless operator fraction	[0, 0.8]	0.4
Stateful operator fraction	[0, 0.8]	0.4
Partitioned stateful operator fraction	[0, 0.8]	0.2
Thread switching overhead	[10, 210]	1
Replication cost factor	[10, 190]	50

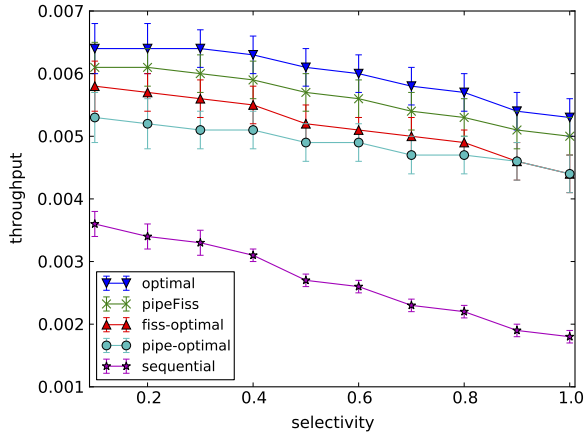


Fig. 5. The impact of selectivity.

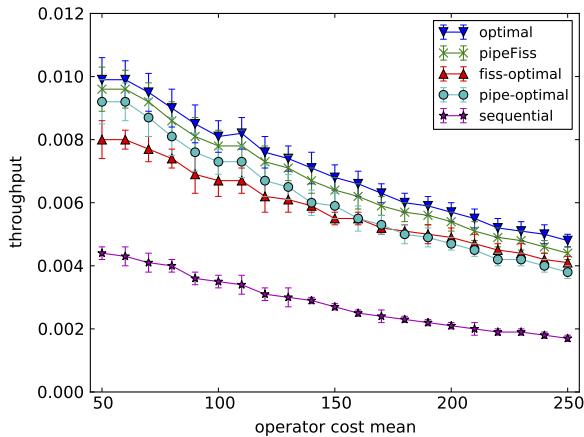


Fig. 6. The impact of operator cost.

cost, and operator kind with respect to state. Our default settings use an operator cost that is 200 times the thread switching overhead. When projected on SPL, this corresponds to a per-tuple operator cost of 19 μ s, which is quite reasonable based on our observation of real-world operator costs (see Fig. 17 for a sample real-world application and its operator costs). In addition, a variety of other factors impact the throughput of the topology, among which four most important ones are replication cost factor, thread switching overhead, number of cores, and the number of operators. Accordingly, we also perform experiments on these.

5.2.1. Operator selectivity

The impact of operator selectivity on the performance of our solution is shown in Fig. 5. The figure plots the throughput (y-axis) as a function of the mean operator selectivity (x-axis)




Fig. 7. The impact of the operator kind.

for different approaches. We observe that for the entire range of selectivity values, our solution outperforms the fission-only and pipelining-only optimal approaches, and provides up to 2.7 times speedup in throughput compared to the sequential approach. The throughput provided by our approach is also consistent within 5% of the optimal solution, for the entire range of selectivity values. Interestingly, we observe that the pipelining-only approach provides reduced performance compared to fission-only approach, for low selectivity values. This is because with reducing selectivity, the performance impact of the operators that are deeper in the pipeline reduces, which takes away the ability of pipelining to increase the throughput (as speedup due to pipelining is limited by the pipeline depth).

5.2.2. Operator cost mean

The impact of operator cost on the performance of our solution is shown in Fig. 6. The figure plots the throughput (y-axis) as a function of the mean operator cost (x-axis) for different approaches. Again we observe that the pipelined fission solution is quite robust, consistently outperforming fission-only and pipelining-only optimal solutions, and staying within 5% of the optimal solution. It achieves up to 2.5 times speedup in throughput compared to the sequential approach. One interesting observation is that, for smaller mean operator cost values, the performance of the fission-only approach is below the pipelining-only approach, but gradually increases and passes it as the mean operator cost increases. The reason is that, the fission optimization has a higher overhead due to the replication cost factor, and thus, for small operator costs, it is not beneficial to apply fission. As the operator cost increases, fission becomes more effective.

5.2.3. Operator kind

The impact of the operator kind on the performance of our solution is shown in Fig. 7. Recall that operators can be stateful, stateless, or partitioned stateful. Fig. 7 plots the throughput (y-axis) as a function of the fraction of stateless operators (x-axis) for different approaches. While doing this, we keep the fraction of partitioned stateful operators fixed at 0.2. We observe that the percentage of stateless operators do not impact the pipelining-only solution. The reason is that pipeline parallelism is applicable for both stateful and stateless operators. On the other hand, fission-only solution improves as the percent of the stateless operator increases. The reason is that data parallelism is not applicable for stateful operators. We also observe that our pipelined fission solution stays close to the optimal throughput the entire range of the stateless operator fraction. Again, pipelined fission clearly outperforms pipelining-only and fission-only approaches.

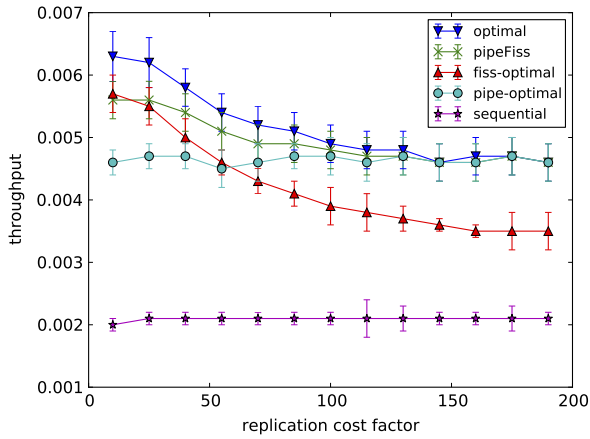


Fig. 8. The impact of replication cost factor.




Fig. 10. The impact of the number of cores.




Fig. 9. The impact of the thread switching overhead.

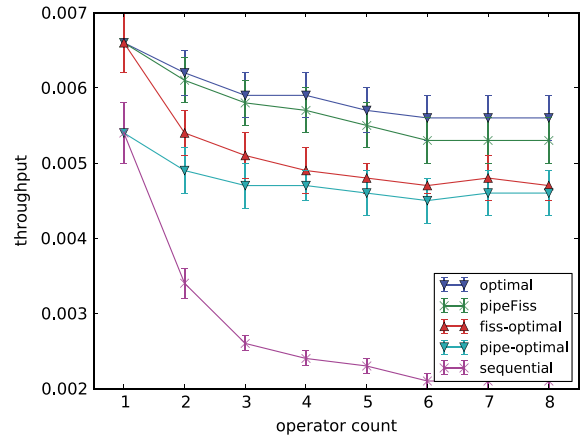


Fig. 11. The impact of the number of operators.

5.2.4. Replication cost factor

The impact of the replication cost factor on the performance of our solution is shown in Fig. 8. The figure plots the throughput (y-axis) as a function of the replication cost factor (x-axis) for different approaches. The results indicate that our solution considerably outperforms pipelining-only and fission-only approaches, providing up to 25% higher throughput compared to pipelining-only optimal approach and up to 30% higher throughput compared to fission-only optimal approach. Note that our pipelined fission approach provides performance as good as fission-only optimal approach when the replication cost factor is close to 0 and as good as pipelining-only optimal approach when the replication cost factor is very high. In effect, our solution switches from using fission to using pipelining as the replication cost factor increases. We also observe that the optimal solution’s throughput advantage is bigger for small replication cost factors, yet the gap with pipelined fission quickly closes as the replication cost factor increases. In the SPL based experiments presented later, we show that for realistic replication cost factors, our solution provides performance that is very close to the optimal.

5.2.5. Thread switching overhead

The impact of the thread switching overhead on the performance of our solution is shown in Fig. 9. The figure plots the throughput (y-axis) as a function of the thread switching overhead (x-axis) for different approaches. We observe that the throughput of all solutions, except the sequential one, decreases as the thread switching overhead increases. It is an expected result as all solutions benefit from parallelism via using threads, except the sequential solution. Again, our pipelined fission solution outperforms

pipelining-only and fission-only optimal solutions, and is able to stay close to the optimal performance throughout the entire range of thread switching overhead values. We also observe that as the thread switching overhead increases, all approaches start to get closer in terms of the throughput. This is due to the reducing parallelization opportunities, as a direct consequence of the high thread switching overhead values.

5.2.6. Number of cores

The impact of the number of cores on the performance of our solution is shown in Fig. 10. The figure plots the throughput (y-axis) as a function of the number of cores (x-axis) for different approaches. We observe that for all approaches, the throughput only increases to a certain degree, after which it stays flat. There are two reasons for not being able to achieve linear speedup: (i) not all operators are parallelizable, (ii) the thread switching and replication cost factor introduce overheads in parallelization. We again observe that our pipelined fission approach outperforms the pipelining-only and fission-only optimal solutions. With increasing number of cores, the gap between the optimal approach and alternatives increases, as the search space gets bigger. However, since the throughput flattens quickly, the increase in the gap eventually stops. At that point, our approach is still within 8% of the optimal.

5.2.7. Number of operators

The impact of the number of operators on the performance of our solution is shown in Fig. 11. The figure plots the throughput (y-axis) as a function of the number of operators (x-axis) for

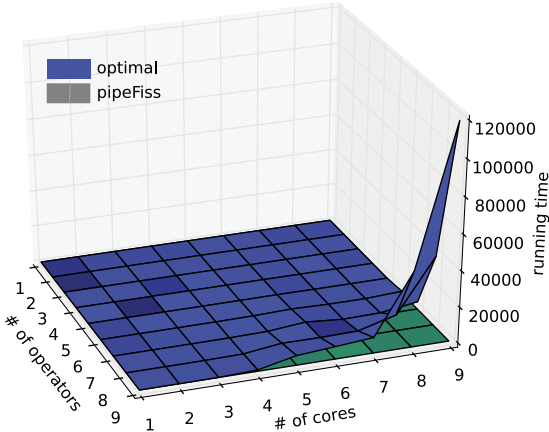


Fig. 12. Running time (in milliseconds).

different approaches. We observe that as the number of operators increases, the performance of the pipelining-only solution relative to the fission-only solution increases. The reason is that the pipeline parallelism cannot help a single operator, so it is not as effective for small number of operators. Our pipelined fission solution provides up to 18% higher throughput compared to the closest alternative. While the gap between the optimal solution and ours increases with increasing number of operators, eventually throughput flattens due to the fixed number of cores available. Importantly, our approach stays within 5% of the optimal solution.

5.2.8. Running time

We also evaluate the running time of our pipelined fission algorithm. Fig. 12 plots running time in terms of milliseconds, for our pipelined fission solution and the exhaustive optimal approach. Unfortunately, the running time of the optimal solution dramatically increases with increasing number of operators and cores. For 9 operators and 9 cores, the running time reaches 2 min, which makes it inapplicable for runtime adaptation. Furthermore, the time grows very quickly, reaching hours for 10 operators and 10 cores (not shown), and becomes practically unusable even for static optimization for larger setups. However, even if the number of operators and cores are high, our pipelined fission algorithm completes much faster (under 5 ms).

5.3. SPL experiments

In our second set of experiments, we use IBM's SPL language and its InfoSphere Streams runtime to evaluate the effectiveness of our solution. In order to perform this experiment, we need to determine the value of the replication cost factor and the thread switching overhead for the InfoSphere Streams runtime.

5.3.1. Thread switching overhead

For determining the thread switching overhead, we use a simple pipeline of two operators. We run this topology twice, once with a single thread and again with two threads. Let c be the cost of the operator. For the case of two threads, the throughput achieved, denoted as T_p , is given by:

$$T_p = 1/(c + \delta). \quad (15)$$

On the other hand, for the case of a single thread, the throughput achieved, denoted as T_s , is given by:

$$T_s = 1/(2 \cdot c). \quad (16)$$

By using T_s and T_p , we can compute the thread switching overhead, δ , without needing to know the operator cost c . More specifically,

$$\delta = \frac{1}{T_p} - \frac{1}{2 \cdot T_s}. \quad (17)$$

In order to calculate the thread switching overhead for our SPL experiments, we measure the throughput of the topology with and without pipeline parallelism for varying tuple sizes, and use Eq. (17) to compute the thread switching overhead. The use of different tuple sizes is due to the implementation of thread switching within the SPL runtime, which requires a tuple copy (the cost of which depends on the tuple size).

5.3.2. Replication cost factor

For determining the replication cost factor, we use a simple pipeline of three operators, where the first and the last operators are the source and the sink operators with no work performed and the middle operator has cost c . We then run this topology with different number of parallel channels used for the middle operator. Let n denote the number of channels used. We can formulate the throughput as:

$$T_p(n) = \left(\frac{2 \cdot \delta + c}{n} + \log_2 n \cdot c_p \right)^{-1}. \quad (18)$$

If we know the throughput for two different number of channels, say $T(n_1)$ and $T(n_2)$, then we can compute the replication cost factor, c_p , independent of other factors, such as the cost c , as follows:

$$c_p = \frac{\frac{1}{n_2 \cdot T_p(n_1)} - \frac{1}{n_1 \cdot T_p(n_2)}}{\frac{\log_2 n_1}{n_2} - \frac{\log_2 n_2}{n_1}}. \quad (19)$$

In order to calculate the replication cost factor for our SPL experiments, we measure the throughput of our sample topology with different number of replicas for varying tuple sizes, and use Eq. (19) to compute the replication cost factor.

By using the calculated thread switching overhead and replication cost factor values, we perform SPL experiments to evaluate our solution for varying operator count, selectivity, cost, and kind. Throughput is again our main metric for evaluation. The applications used for these experiments are similar to the ones from the model-based experiments, but are written using the SPL language. The operators used are *busy* operators that perform repeated multiplication operations to emulate work (the cost is the number of multiplications performed). To emulate selectivity, they draw a random number for each incoming tuple and compare it to the selectivity value to determine if the tuple should be forwarded or not. This setup enables us to study a wide range of parameter settings. To further strengthen the evaluation, we have also applied our pipelined fission solution to a real-world application called *LogWatch*, which we detail later in this section.

5.3.3. Operator selectivity

Fig. 13 plots throughput (y-axis) as a function of the mean operator selectivity (x-axis) for the optimal, pipelined fission, and sequential solutions using SPL. We see that all approaches achieve lower throughput as the operator selectivity increases. Pipelined fission solution provides practically the same performance as the optimal solution for selectivities beyond 0.7 and is within 15% of the optimal for selectivities as small as 0.3. When the selectivity gets very low, the variance in results significantly increases as very few tuples make it through the pipeline.




Fig. 13. Impact of selectivity (using SPL).




Fig. 16. The impact of the operator kind (using SPL).




Fig. 14. The impact of operator cost (using SPL).




Fig. 15. The impact of the number of operators (using SPL).

5.3.4. Operator cost mean

Fig. 14 plots throughput (y-axis) as a function of the mean operator cost (x-axis) for the optimal, pipelined fission, and sequential solutions using SPL. It is not surprising that the throughput decreases as the mean operator cost increases, for all approaches. More interestingly, our approach again performs as good as the optimal approach throughout the entire cost range, for which we are still within 10% of the optimal.

5.3.5. Number of operators

Fig. 15 plots throughput (y-axis) as a function of the number of operators (x-axis) for the optimal, pipelined fission, and sequential solutions using SPL. As expected, as the number of operators in a topology increases, throughput of a topology decreases for all approaches. Even for high number of operators, the throughput achieved by pipelined fission solution is as good as the optimal one.

5.3.6. Operator kind

Fig. 16 plots throughput (y-axis) as a function of the fraction of stateless operators. The change in operator kind does not affect the sequential solution as in Fig. 16. On the other hand, as the percentage of stateless operators increases, the throughput achieved increases. The reason is that stateless operators can benefit from both data and pipeline parallelism. As it can be seen from the figure, our pipelined fission solution again performs close to the optimal solution, providing the same throughput when the stateless fraction is 0.7 or more, and within 10% of the optimal for smaller fractions.

5.3.7. The LogWatch application

The LogWatch application is an SPL application that monitors Linux login audits, looking for logins that were preceded by many failed login attempts from the same host. Such logins are flagged as *break-ins*. The input is synthetic data based on real data collected from a public-facing server which experienced a break-in attempt about every 2 s for a 12 h period. The real data has been modified with several fake break-ins, and the 12 h period is cycled through 2000 times. The application was first used in [15], but has been modified for this work. The modification involves reworking a branch in the application flow to put it into chain topology. Fig. 17 lists the operators (in left-to-right order) in the LogWatch application, including their kind with respect to state, selectivity (between 0 and 1), and cost (in terms of microseconds/tuple). Note that while the kinds of the operators *Range* and *Breakins* are both partitioned stateful, they are partitioned on different attributes (denoted by *r* (host) and *u* (user) in the table).

Fig. 18 plots the throughput of the LogWatch application as a function of the number of cores used, for the sequential, pipelined fission, and optimal approaches. We observe that the pipelined fission and the optimal approaches perform similarly. This is consistent with the results from our synthetic application experiments. We also observe that our approach has good scalability. Going from 1 core to 12 cores (max available on our machine), the throughput reaches 9.7× the throughput of the sequential application.

Operator	kind	selectivity	cost
RawLines	stateful	N/A	1.52
Lines	stateless	1	3.87
ParsedLines	stateless	1	15.15
RawEvents	stateless	0.87	9.39
Events	stateless	1	13.81
Range	partitioned (r)	0.20	9.04
Cutoff	stateless	0.99	24.82
RealTime	stateless	1	27.04
Breakins	partitioned (u)	0.00027	17.33
Results	stateful	N/A	493.17

Fig. 17. Operators in the LogWatch application (in left-to-right order, cost in micro seconds per tuple).

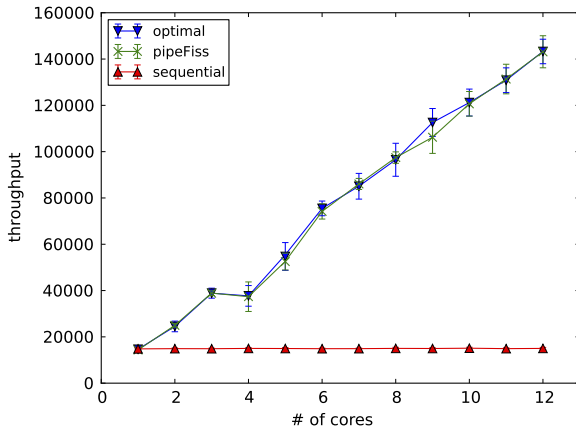


Fig. 18. The number of cores vs. LogWatch throughput.

6. Related work

Our work belongs to the general area of auto-parallelization. We first overview prior work in this area, and then focus on work related to the core subject of our paper: auto-parallelization in streaming systems.

6.1. Multi-threaded concurrency platforms

Determining parallelizable code regions and appropriately assigning those regions to computing units for execution are the two major issues that must be addressed by any automatic parallelization system.

Multi-threaded concurrency platforms, such as Cilk++ [24], OpenMP [30], and x10 [8], decouple expressing a program's innate parallelism from its execution configuration. OpenMP and Cilk++ are widely used language extensions for shared memory programs, which help express parallel execution in a program at development-time and take advantage of it at run-time.

Various platforms are proposed in the literature for automatically finding parallelizable program regions. One example is Kremlin [11], which is an auto-parallelization framework that complements OpenMP [30]. Kremlin recommends to programmers a list of regions for parallelization, which is ordered by achievable program speedup. The speedup is calculated based on an improved critical path analysis.

Cilkview [19] is a Cilk++ analyzer of program scalability in terms of number of cores. Cilkview performs system-level modeling of scheduling overheads (e.g., the bookkeeping costs to set up context and the overhead of cache misses), and predicts program speedup. Bounds on the speedup are presented to programmers for further analysis.

Autopin [26] is an auto-configuration framework for finding the best mapping between system cores and threads. Using profile runs, Autopin exhaustively probes all possible mappings and finds the best pinning configuration in terms of performance. Then, threads are re-pinned using the best mapping found.

Alchemist [42] is a dependence profiling technique based on post-dominance analysis and is used to detect candidate regions for parallel execution. It is based on the observation that a procedure with few dependencies with its continuation benefits more from parallelization.

There has been extensive research in the literature on compiler support for instruction-level or fine-grained pipelined parallelism [27]. In our work, we look at coarse-grained pipelining techniques that address the problem of decomposing an application into higher-level pieces that can execute in pipeline as well as data parallel. Relevant to our study is the work in [10], which provides compiler support for coarse-grained pipelined parallelism. To automate pipelining, it selects a set of candidate filter boundaries (a middleware interface exposed by DataCutter [6]), determines the communication volume for these boundaries, and performs decomposition and code generation in order to minimize the execution time. To select the best filters, communication costs across each filter boundary are estimated by static program analysis and a dynamic programming algorithm is used to find the optimal decomposition. In comparison, our work performs combined pipelining and fission and has support for partitioned stateful operators.

6.2. Pipelining/fusion in streaming systems

In most streaming systems operators that are fused together use the same thread, whereas nonfused operators can be run in parallel. The key problem is to divide a program into fused pieces that can be run in parallel, typically in a pipelined configuration. For instance, StreamIt [16], which is a language for creating streaming applications, uses fusion to coarsen the granularity of the graph to the target number of cores, based on cost estimates [18]. This is somewhat similar to our region configuration step, but is limited to stateless operators or operators that only have read-only sliding window state.

Aurora data stream management system uses fusion to minimize scheduling overhead [2]. Based on a similar model of streaming, SPADE [14] uses the COLA [25] fusion optimizer to combine operators as much as possible, until a single processing element fills the entire capacity of a core. A different approach is taken by Tang and Gedik [39], where the stream program initially runs as completely fused, and an *auto-pipeliner* is used to detect bottlenecks and inject new threads into the runtime system to improve throughput. With the exception of StreamIt, which we further cover shortly, these systems are limited to pipelined parallelism, and do not perform combined pipelining and fission.

6.3. Fission in streaming systems

StreamIt [17] performs both pipelining and fission. It addresses the safety question of fission by only replicating operators that are either stateless or whose operator state is a read-only sliding window. Assuming the same model of synchronous data flow (SDF), Kudlur and Mahlke present a solution for orchestrating the execution of stream programs on multicore platforms [28]. It uses an integrated unfolding and partitioning based integer linear programming solution for this purpose. As opposed to SDF based systems, our work targets data stream management systems that typically contain operators that are partitioned stateful and exhibit dynamic selectivity. Thus, rather than having a static schedule based execution model, we adopt a backpressure based runtime system. We model its throughput in order to formulate a pipelined

fission configuration that can provide optimal throughput. Work on *elastic operators* [32] also generalizes fission beyond the SDF setting to work on stateful operators with dynamic data rates. However, the work is limited to fission only and does not support pipelining.

A related problem is to perform fission dynamically, that is to adjust the width of the parallel region based on the changing runtime and workload conditions. SEDA achieves this via a thread-pool controller, which can adjust the number of threads to increase parallelism, while preserving locality [40]. MapReduce systems dynamically adjust the number of workers assigned to the map tasks [9]. Elastic operators [32] adjust the number of threads assigned to an operator by using a control loop. An extension of it [15] applies similar kind of control in a distributed setup, where replication is not limited to a single operator and replicas can run across different hosts. While our paper does not particularly deal with the adaptation aspect, its model based approach and efficient heuristic solver makes it perfectly suitable for runtime optimization based on feedback from a performance profiler.

A general overview of optimizations in streaming systems, including parallelization, is given in [23]. Overall, our work is distinguished from earlier work on streaming systems, as it is the only work that combines pipelining and fission in the context of partitioned stateful operators with dynamic selectivity.

7. Conclusion

We proposed a pipelined fission solution that can quickly locate a parallelization configuration for accelerating data stream processing applications, and can provide throughput close to the optimal. In order to achieve this aim, our algorithm incorporates stages that greedily perform data and pipeline parallelism with a varying fraction of resources dedicated to the two different parallelization approaches. Our model based experimental evaluation shows that our proposed algorithm is effective both in terms of running time and throughput under varying operator, system, and workload properties. Our evaluation using an industrial-strength stream processing engine showcases strong results as well, where our pipelined fission solution provides throughput that is very close (5%–10%) to that of the optimal.

References

- [1] D. Abadi, Y. Ahmad, M. Balazinska, U. Çetintemel, M. Cherniack, J.-H. Hwang, W. Lindner, A. Maskey, A. Rasin, E. Ryzkina, N. Tatbul, Y. Xing, S. Zdonik, The design of the Borealis stream processing engine, in: Innovative Data Systems Research Conference, CIDR, 2005.
- [2] D.J. Abadi, D. Carney, U. Çetintemel, M. Cherniack, C. Convey, S. Lee, M. Stonebraker, N. Tatbul, S. Zdonik, *Aurora: A new model and architecture for data stream management*, VLDB J. 12 (2) (2003) 120–139.
- [3] H. Andrade, B. Gedik, D. Turaga, *Fundamentals of Stream Processing: Application Design, Systems, Analytics*, first ed., Cambridge Press, Cambridge, UK, 2014, (Chapter 12.4)—The Semiconductor Process Control application.
- [4] H. Andrade, B. Gedik, K.-L. Wu, P.S. Yu, Processing high data rate streams in systems, J. Parallel Distrib. Comput. (JPDC) 71 (2011) 145–156. Special Issue on Data Intensive Computing.
- [5] A. Arasu, B. Babcock, S. Babu, M. Datar, K. Ito, R. Motwani, I. Nishizawa, U. Srivastava, D. Thomas, R. Varma, J. Widom, STREAM: The Stanford stream data manager, IEEE Data Eng. Bull., Vol. 26, no. 1.
- [6] M.D. Beynon, T.M. Kurç, Ü.V. Çatalyürek, C. Chang, A. Sussman, J.H. Saltz, Distributed processing of very large datasets with DataCutter, Parallel Comput. J. 27 (11) (2001) 1457–1478.
- [7] E. Bouillet, R. Kothari, V. Kumar, L. Mignet, S. Nathan, A. Ranganathan, D.S. Turaga, O. Udre, O. Verscheure, Processing 6 billion CDRs/day: from research to production (experience report), in: International Conference on Distributed Event Based Systems, DEBS, 2012.
- [8] P. Charles, C. Grothoff, V.A. Saraswat, C. Donawa, A. Kielstra, K. Ebcioğlu, C. von Praun, V. Sarkar, X10: An object-oriented approach to non-uniform cluster computing, in: International Conference on Object-Oriented Programming, Systems, Languages & Applications, OOPSLA, 2005.
- [9] J. Dean, S. Ghemawat, Mapreduce: Simplified data processing on large clusters, in: USENIX Symposium on Operating Systems Design and Implementation, OSDI, 2004, pp. 137–150, 2008.
- [10] W. Du, R. Ferreira, G. Agrawal, Compiler support for exploiting coarse-grained pipelined parallelism, in: Supercomputing Conference, SC, 2003, p. 8.
- [11] S. Garcia, D. Jeon, C.M. Louie, M.B. Taylor, Kremlin: Rethinking and rebooting gprof for the multicore age, in: International Conference on Programming Language Design and Implementation, PLDI, 2011.
- [12] B. Gedik, H. Andrade, A model-based framework for building extensible, high performance stream processing middleware and programming language for IBM InfoSphere Streams, Software: Practice and Experience 42 (11) (2012) 1363–1391.
- [13] B. Gedik, H. Andrade, K.-L. Wu, A code generation approach to optimizing high-performance distributed data stream processing, in: ACM International Conference on Information and Knowledge Management, CIKM, 2009.
- [14] B. Gedik, H. Andrade, K.-L. Wu, P.S. Yu, M. Doo, SPADE: The systems declarative stream processing engine, in: ACM International Conference on Management of Data, SIGMOD, 2008, pp. 1123–1134.
- [15] B. Gedik, S. Schneider, K.-L.W.M. Hirzel, Elastic scaling for data stream processing, IEEE Trans. Parallel Distrib. Syst. (TPDS), <http://dx.doi.org/10.1109/TPDS.2013.295>.
- [16] M.I. Gordon, W. Thies, S. Amarasinghe, Exploiting coarse-grained task, data, and pipeline parallelism in stream programs, in: International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS, 2006.
- [17] M.I. Gordon, W. Thies, S. Amarasinghe, Exploiting coarse-grained task, data, and pipeline parallelism in stream programs, in: International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS, 2006, pp. 151–162.
- [18] M.I. Gordon, W. Thies, M. Karczmarek, J. Lin, A.S. Meli, A.A. Lamb, C. Leger, J. Wong, H. Hoffmann, D. Maze, S. Amarasinghe, A stream compiler for communication-exposed architectures, in: International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS, 2002, pp. 291–303.
- [19] Y. He, C.E. Leiserson, W.M. Leiserson, The cilkview scalability analyzer, in: ACM Symposium on Parallelism in Algorithms and Architectures, SPAA, 2010.
- [20] M. Hirzel, H. Andrade, B. Gedik, G. Jacques-Silva, R. Khandekar, V. Kumar, M. Mendell, H. Nasgaard, S. Schneider, R. Soule, K.-L. Wu, Streams processing language: Analyzing big data in motion, IBM J. Res. Dev. 57 (2013) 7:1–7:11.
- [21] M. Hirzel, H. Andrade, B. Gedik, V. Kumar, G. Losa, M. Mendell, H. Nasgaard, R. Soule, K.-L. Wu, SPL language spec., Tech. Rep. RC24897, IBM, 2009.
- [22] M. Hirzel, B. Gedik, Streams that compose using macros that oblige, in: ACM Workshop on Partial Evaluation and Program Manipulation, PEPM, 2012.
- [23] M. Hirzel, R. Soule, S. Schneider, B. Gedik, R. Grimm, A catalog of stream processing optimizations, ACM Comput. Surv., Vol. 46, no. 4.
- [24] Intel cilk++. <http://software.intel.com/en-us/articles/intel-cilk-plus/> (retrieved October, 2014).
- [25] R. Khandekar, K. Hildrum, S. Parekh, D. Rajan, J. Wolf, K.-L. Wu, H. Andrade, B. Gedik, COLA: Optimizing stream processing applications via graph partitioning, in: ACM/IFIP/USENIX Middleware Conference, Middleware, 2009.
- [26] T. Klug, M. Ott, J. Weidendorfer, C. Trinitis, Autopin: Automated optimization of thread-to-core pinning on multicore systems, Trans. High-Perform. Embedded Archit. Compil. (HiPEAC) 3 (2011) 219–235.
- [27] S.M. Krishnamurthy, A brief survey of papers on scheduling for pipelined processors, ACM SIGPLAN Not. 25 (7) (1990) 97–106.
- [28] M. Kudlur, S.A. Mahlke, Orchestrating the execution of stream programs on multicore platforms, in: International Conference on Programming Language Design and Implementation, PLDI, 2008, pp. 114–124.
- [29] S. Kulkarni, N. Bhagat, M. Fu, V. Kedigehalli, C. Kellogg, S. Mittal, J.M. Patel, K. Ramasamy, S. Taneja, Twitter heron: Stream processing at scale, in: ACM International Conference on Management of Data, SIGMOD, 2010, pp. 239–250.
- [30] Openmp. <http://www.openmp.org> (retrieved October, 2014).
- [31] Samza project. <http://samza.apache.org/> (retrieved Nov., 2014).
- [32] S. Schneider, H. Andrade, B. Gedik, A. Biem, K.-L. Wu, Elastic scaling of data parallel operators in stream processing, in: IEEE International Parallel and Distributed Processing Symposium, IPDPS, 2009.
- [33] S. Schneider, M. Hirzel, B. Gedik, K.-L. Wu, Safe data parallelism for general streaming, IEEE Trans. Comput. (TC), Vol. 64, no. 2.
- [34] M.A. Shah, J.M. Hellerstein, S. Chandrasekaran, M.J. Franklin, Flux: An adaptive partitioning operator for continuous query systems, in: IEEE International Conference on Data Engineering, ICDE, 2003.
- [35] D.M. Sow, A. Biem, M. Blount, M. Ebling, O. Verscheure, Body sensor data processing using stream computing.
- [36] Storm project. <http://storm-project.net/> (retrieved Nov., 2013).
- [37] StreamBase Systems. <http://www.streambase.com> (retrieved Nov., 2013).
- [38] S4 distributed stream computing platform. <http://www.s4.io/> (retrieved May, 2012).
- [39] Y. Tang, B. Gedik, Auto-pipelining for data stream processing, IEEE Trans. Parallel Distrib. Syst. (TPDS) 24 (12) (2013) 2344–2354.
- [40] M. Welsh, D. Culler, E. Brewer, SEDA: An architecture for well-conditioned, scalable Internet services, in: ACM Symposium on Operating Systems Principles (SOSP), 2001, pp. 230–243.

- [41] X.J. Zhang, H. Andrade, B. Gedik, R. King, J.F. Morar, S. Nathan, Y. Park, R. Pavuluri, E. Pring, R. Schnier, P. Selo, M. Spicer, V. Uhlig, C. Venkatramani, Implementing a high-volume, low-latency market data processing system on commodity hardware using IBM middleware, in: Workshop on High Performance Computational Finance, SC-WHPCF, 2009.
- [42] X. Zhang, A. Navabi, S. Jagannathan, Alchemist: A transparent dependence distance profiling infrastructure, in: International Symposium on Code Generation and Optimization, CGO, 2009, pp. 47–58.



Habibe Güldamla Özsema is a graduate student in the Department of Computer Engineering, Bilkent University, Turkey. Her research interest is in parallel stream processing systems.



Buğra Gedik is an Associate Professor in the Department of Computer Engineering, Bilkent University, Turkey. He holds a Ph.D. degree in Computer Science from Georgia Institute of Technology. His research interests are in data-intensive distributed systems.



Özcan Öztürk is an Associate Professor in the Department of Computer Engineering, Bilkent University, Turkey. He holds a Ph.D. degree in Computer Science from Pennsylvania State University. His research interests are in parallel systems, computer architecture, and compilers.