

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/279458648>

# Lightweight Asynchronous Snapshots for Distributed Dataflows

Article · June 2015

Source: arXiv

---

CITATIONS

36

---

READS

371

5 authors, including:



**Paris Carbone**

KTH Royal Institute of Technology

10 PUBLICATIONS 415 CITATIONS

SEE PROFILE



**Seif Haridi**

KTH Royal Institute of Technology

227 PUBLICATIONS 4,088 CITATIONS

SEE PROFILE



**Kostas Tzoumas**

Technische Universität Berlin

21 PUBLICATIONS 1,020 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Mobile Security and Dependability [View project](#)



Streamline, E2E-clouds [View project](#)

# Lightweight Asynchronous Snapshots for Distributed Dataflows

Paris Carbone<sup>1</sup>

Gyula Fóra<sup>2</sup>

Stephan Ewen<sup>3</sup>

Seif Haridi<sup>1,2</sup>

Kostas Tzoumas<sup>3</sup>

<sup>1</sup>KTH Royal Institute of Technology - {parisc,haridi}@kth.se

<sup>2</sup>Swedish Institute of Computer Science - {gyfora, seif}@sics.se

<sup>3</sup>Data Artisans GmbH - {stephan, kostas}@data-artisans.com

## Abstract

Distributed stateful stream processing enables the deployment and execution of large scale continuous computations in the cloud, targeting both low latency and high throughput. One of the most fundamental challenges of this paradigm is providing processing guarantees under potential failures. Existing approaches rely on periodic global state snapshots that can be used for failure recovery. Those approaches suffer from two main drawbacks. First, they often stall the overall computation which impacts ingestion. Second, they eagerly persist all records in transit along with the operation states which results in larger snapshots than required. In this work we propose Asynchronous Barrier Snapshotting (ABS), a lightweight algorithm suited for modern dataflow execution engines that minimises space requirements. ABS persists only operator states on acyclic execution topologies while keeping a minimal record log on cyclic dataflows. We implemented ABS on Apache Flink, a distributed analytics engine that supports stateful stream processing. Our evaluation shows that our algorithm does not have a heavy impact on the execution, maintaining linear scalability and performing well with frequent snapshots.

**Keywords** fault tolerance, distributed computing, stream processing, dataflow, cloud computing, state management

## 1. Introduction

Distributed dataflow processing is an emerging paradigm for data intensive computing which allows continuous computations on data in high volumes, targeting low end-to-end latency while guaranteeing high throughput. Several time-critical applications could benefit

from dataflow processing systems such as Apache Flink [1] and Naiad [11], especially in the domains of real-time analysis (e.g. predictive analytics and complex event processing). Fault tolerance is of paramount importance in such systems, as failures cannot be afforded in most real-world use cases. Currently known approaches that guarantee exactly-once semantics on stateful processing systems rely on global, consistent snapshots of the execution state. However, there are two main drawbacks that make their application inefficient for real-time stream processing. Synchronous snapshotting techniques stop the overall execution of a distributed computation in order to obtain a consistent view of the overall state. Furthermore, to our knowledge all of the existing algorithms for distributed snapshots include records that are in transit in channels or unprocessed messages throughout the execution graph as part of the snapshot state. Most often this includes state that is larger than required.

In this work, we focus on providing lightweight snapshotting, specifically targeted at distributed stateful dataflow systems, with low impact on performance. Our solution provides asynchronous state snapshots with low space costs that contain only operator states in acyclic execution topologies. Additionally, we cover the case of cyclic execution graphs by applying downstream backup on selected parts of the topology while keeping the snapshot state to minimum. Our technique does not halt the streaming operation and it only introduces a small runtime overhead. The contributions of this paper can be summarised as follows:

- We propose and implement an asynchronous snapshotting algorithm that achieves minimal snapshots on acyclic execution graphs.

- We describe and implement a generalisation of our algorithm that works on cyclic execution graphs.
- We show the benefits of our approach compared to the state-of-the-art using Apache Flink Streaming as a base system for comparisons.

The rest of the paper is organised as follows: Section 2 gives an overview of existing approaches for distributed global snapshots in stateful dataflow systems. Section 3 provides an overview of the Apache Flink processing and execution model followed by Section 4 where we describe our main approach to global snapshotting in detail. Our recovery scheme is described briefly in Section 5. Finally, Section 6 summarises our implementation followed by our evaluation in Section 7 and future work and conclusion in Section 8.

## 2. Related Work

Several recovery mechanisms have been proposed during the last decade for systems that do continuous processing [4, 11]. Systems that emulate continuous processing into stateless distributed batch computations such as Discretized Streams and Comet [6, 15] rely on state recomputation. On the other hand, stateful dataflow systems such as Naiad, SDGs, Piccolo and SEEP [3, 5, 11, 12], which are our main focus in this work, use checkpointing to obtain consistent snapshots of the global execution for failure recovery.

The problem of consistent global snapshots in distributed environments, as introduced by Chandy and Lamport [4], has been researched extensively throughout the last decades [4, 7, 8]. A global snapshot theoretically reflects the overall state of an execution, or a possible state at a specific instance of its operation. A simple but costly approach employed by Naiad [11] is to perform a synchronous snapshot in three steps: first halting the overall computation of the execution graph, then performing the snapshot and finally instructing each task to continue its operation once the global snapshot is complete. This approach has a high impact on both throughput and space requirements due to the need to block the whole computation, while also relying on upstream backup that logs emitted records at the producer side. Another popular approach, originally proposed by Chandy and Lamport [4], that is deployed in many systems today is to perform snapshots asynchronously while eagerly doing upstream backup [4, 5, 10]. This is achieved by distributing markers

throughout the execution graph that trigger the persistence of operator and channel state. This approach though still suffers from additional space requirements due to the need of an upstream backup and as a result higher recovery times caused by the reprocessing of backup records. Our approach extends the original asynchronous snapshotting idea of Chandy and Lamport, however, it considers no backup logging of records for acyclic graphs while also keeping very selective backup records on cyclic execution graphs.

## 3. Background: The Apache Flink System

Our current work is guided by the need for fault tolerance on Apache Flink Streaming, a distributed stream analytics system that is part of the Apache Flink Stack (former Stratosphere [2]). Apache Flink is architected around a generic runtime engine uniformly processing both batch and streaming jobs composed of stateful interconnected tasks. Analytics jobs in Flink are compiled into directed graphs of tasks. Data elements are fetched from external sources and routed through the task graph in a pipelined fashion. Tasks are continuously manipulating their internal state based on the received inputs and are generating new outputs.

### 3.1 The Streaming Programming Model

The Apache Flink API for stream processing allows the composition of complex streaming analytics jobs by exposing unbounded partitioned data streams (partially ordered sequences of records) as its core data abstraction, called *DataStreams*. *DataStreams* can be created from external sources (e.g. message queues, socket streams, custom generators) or by invoking operations on other *DataStreams*. *DataStreams* support several operators such as *map*, *filter* and *reduce* in the form of higher order functions that are applied incrementally per record and generate new *DataStreams*. Every operator can be parallelised by placing parallel instances to run on different partitions of the respective stream, thus, allowing the distributed execution of stream transformations.

The code example in 1 shows how to implement a simple incremental word count in Apache Flink. In this program words are read from a text file and the current count for each word is printed to the standard output. This is a stateful streaming program as sources need to be aware of their current file offset and counters need




Figure 1: The execution graph for incremental word count

to maintain the current count for each word as their internal state.

```

1 val env : StreamExecutionEnvironment = ...
2 env.setParallelism(2)
3
4 val wordStream = env.readTextFile(path)
5 val countStream = wordStream.groupBy(_).count
6 countStream.print
  
```

Example 1: Incremental Word Count

### 3.2 Distributed Dataflow Execution

When a user executes an application all *DataStream* operators compile into an execution graph that is in principle a directed graph  $G = (T, E)$ , similarly to Naiad [11] where vertices  $T$  represent tasks and edges  $E$  represent data channels between tasks. An execution graph is depicted in Fig. 1 for the incremental word count example. As shown, every instance of an operator is encapsulated on a respective task. Tasks can be further classified as sources when they have no input channels and sinks when no output channels are set. Furthermore,  $M$  denotes the set of all records transferred by tasks during their parallel execution. Each task  $t \in T$  encapsulates the independent execution of an operator instance and is composed of the following: (1) a set of input and output channels:  $I_t, O_t \subseteq E$ ; (2) an operator state  $s_t$  and (3) a user defined function (UDF)  $f_t$ . Data ingestion is pull-based : during its execution each task consumes input records, updates its operator state and generates new records according to its user defined function. More specifically, for each record  $r \in M$  received by a task  $t \in T$  a new state  $s'_t$  is produced along with a set of output records  $D \subseteq M$  according to its UDF  $f_t : s_t, r \mapsto \langle s'_t, D \rangle$ .

## 4. Asynchronous Barrier Snapshotting

In order to provide consistent results, distributed processing systems need to be resilient to task failures.

A way of providing this resilience is to periodically capture snapshots of the execution graph which can be used later to recover from failures. A snapshot is a global state of the execution graph, capturing all necessary information to restart the computation from that specific execution state.

### 4.1 Problem Definition

We define a global snapshot  $G^* = (T^*, E^*)$  of an execution graph  $G = (T, E)$  as a set of all task and edge states,  $T^*$  and  $E^*$  respectively. In more detail,  $T^*$  consists of all operator states  $s_t^* \in T^*, \forall t \in T$ , while  $E^*$  is a set of all channel states  $e^* \in E^*$  where  $e^*$  consists of records that are in transit on  $e$ .

We require that certain properties hold for each snapshot  $G^*$  in order to guarantee correct results after recovery such as *termination* and *feasibility* as described by Tel [14].

*Termination* guarantees that a snapshot algorithm eventually finishes in finite time after its initiation if all processes are alive. *Feasibility* expresses the meaningfulness of a snapshot, i.e. that during the snapshotting process no information has been lost regarding the computation. Formally, this implies that causal order [9] is maintained in the snapshot such that records delivered in tasks are also sent from the viewpoint of a snapshot.

### 4.2 ABS for Acyclic Dataflows

It is feasible to do snapshots without persisting channel states when the execution is divided into stages. Stages divide the injected data streams and all associated computations into a series of possible executions where all prior inputs and generated outputs have been fully processed. The set of operator states at the end of a stage reflects the whole execution history, therefore, it can be solely used for a snapshot. The core idea behind our algorithm is to create identical snapshots with staged snapshotting while keeping a continuous data ingestion.

In our approach, stages are emulated in a continuous dataflow execution by special barrier markers injected in the input data streams periodically that are pushed throughout the whole execution graph down to the sinks. Global snapshots are incrementally constructed as each task receives the barriers indicating execution stages. We further make the following assumptions for our algorithm:

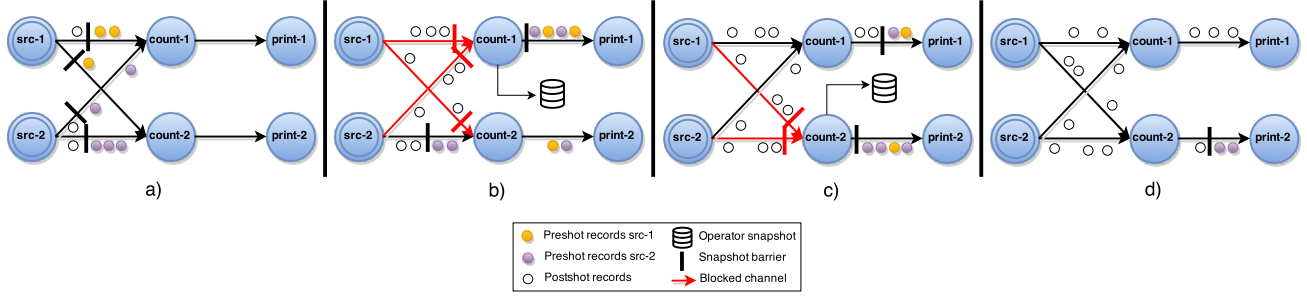


Figure 2: Asynchronous barrier snapshots for acyclic graphs

---

**Algorithm 1** Asynchronous Barrier Snapshotting for Acyclic Execution Graphs

---

```

1: upon event  $\langle \text{Init} \mid \text{input\_channels}, \text{out\_put\_channels}, \text{fun}, \text{init\_state} \rangle$  do
2:    $\text{state} := \text{init\_state}; \text{blocked\_inputs} := \emptyset;$ 
3:    $\text{inputs} := \text{input\_channels};$ 
4:    $\text{out\_puts} := \text{out\_put\_channels}; \text{udf} := \text{fun};$ 
5:
6: upon event  $\langle \text{receive} \mid \text{input}, \langle \text{barrier} \rangle \rangle$  do
7:   if  $\text{input} \neq \text{Nil}$  then
8:      $\text{blocked\_inputs} := \text{blocked\_inputs} \cup \{\text{input}\};$ 
9:     trigger  $\langle \text{block} \mid \text{input} \rangle;$ 
10:    if  $\text{blocked\_inputs} = \text{inputs}$  then
11:       $\text{blocked\_inputs} := \emptyset;$ 
12:      broadcast  $\langle \text{send} \mid \text{outputs}, \langle \text{barrier} \rangle \rangle;$ 
13:      trigger  $\langle \text{snapshot} \mid \text{state} \rangle;$ 
14:      for each  $\text{inputs}$  as  $\text{input}$ 
15:        trigger  $\langle \text{unblock} \mid \text{input} \rangle;$ 
16:
17:
18: upon event  $\langle \text{receive} \mid \text{input}, \text{msg} \rangle$  do
19:    $\{\text{state}', \text{out\_records}\} := \text{udf}(\text{msg}, \text{state});$ 
20:    $\text{state} := \text{state}';$ 
21:   for each  $\text{out\_records}$  as  $\{\text{out\_put}, \text{out\_record}\}$ 
22:     trigger  $\langle \text{send} \mid \text{out\_put}, \text{out\_record} \rangle;$ 
23:
24:

```

---

- Network channels are quasi-reliable, respect a FIFO delivery order and can be *blocked* and *unblocked*. When a channel is *blocked* all messages are buffered but not delivered until it gets *unblocked*.
- Tasks can trigger operations on their channel components such as *block*, *unblock* and *send* messages.

*Broadcasting* messages is also supported on all output channels.

- Messages injected in source tasks (i.e. stage barriers) are resolved into a “Nil” input channel.

**The ABS Algorithm 1 proceeds as follows (depicted in Fig. 2):** A central coordinator periodically injects stage barriers to all the sources. When a source receives a barrier it takes a snapshot of its current state, then broadcasts the barrier to all its outputs (Fig.2(a)). When a non-source task receives a barrier from one of its inputs, it blocks that input until it receives a barrier from all inputs (line 9, Fig.2(b)). When barriers have been received from all the inputs, the task takes a snapshot of its current state and broadcasts the barrier to its outputs (lines 12-13, Fig.2(c)). Then, the task unblocks its input channels to continue its computation (line 15, Fig.2(d)). The complete global snapshot  $G^* = (T^*, E^*)$  will consist solely of all operator states  $T^*$  where  $E^* = \emptyset$ .

**Proof Sketch**<sup>1</sup>: As mentioned earlier a snapshot algorithm should guarantee *termination* and *feasibility*.

*Termination* is guaranteed by the channel and acyclic execution graph properties. Reliability of the channels ensures that every barrier sent will eventually be received as long as the tasks are alive. Furthermore, as there is always a path from a source, every task in a DAG topology will eventually receive barriers from all its input channels and take a snapshot.

For *feasibility* it suffices to show that the operator states in a global snapshot reflect only the history of records processed up to the last stage. This is guaranteed by the FIFO ordering property of the channels and the blocking of input channels upon barriers ensuring that

<sup>1</sup>We omit the formal proof of the algorithms in this version of the paper due to space limitations






Figure 3: Asynchronous barrier snapshots for cyclic graphs

no post-shot records of a stage (records succeeding a barrier) are processed before a snapshot is taken.

### 4.3 ABS for Cyclic Dataflows

In the presence of directed cycles in the execution graph the ABS algorithm presented before would not terminate resulting into a deadlock, as tasks in a cycle would wait indefinitely to receive barriers from all their inputs. Additionally, records that are arbitrarily in transit within cycles would not be included in the snapshot, thus violating feasibility. It is therefore needed to consistently include all records generated within a cycle in the snapshot for feasibility and to put these records back in transit upon recovery. Our approach to deal with cyclic graphs extends the basic algorithm without inducing any additional channel blocking as seen in Algorithm 2. First, we identify *back-edges*  $L$  on loops in the execution graph by static analysis. From control flow graph theory a *back-edge* in a directed graph is an edge that points to a vertex that has already been visited during a depth-first search. The execution graph  $G(T, E \setminus L)$  is a DAG containing all tasks in the topology. From the perspective of this DAG the algorithm operates as before, nevertheless, we additionally apply downstream backup of records received from identified *back-edges* over the duration of a snapshot. This is achieved by each task  $t$  that is a consumer of back-edges  $L_t \subseteq I_t, L_t$  creating a backup log of all records received from  $L_t$  from the moment it forwards barriers until receiving them back from  $L_t$ . Barriers push all records in transit within loops into the downstream logs, so they are included once in the consistent snapshot.

**The ABS Algorithm 2 proceeds as follows (depicted in Fig. 3) in more detail:** Tasks with back-edge inputs create a local copy of their state once all their regular ( $e \notin L$ ) channels have delivered barriers (line 14, Fig.3 (b)). Moreover, from this point they log all records

delivered from their back-edges until they receive stage barriers from them (line 26). This allows, as seen in Fig. 3(c), for all pre-shot records that are in transit within loops to be included in the current snapshot. Mind that the final global snapshot  $G^* = (T^*, L^*)$  contains all task states  $T^*$  and only the back-edge records in transit  $L^* \subset E^*$ .

**Proof Sketch:** Again, we need to prove that *termination* and *feasibility* are guaranteed in this version of the algorithm.

As in 4.2 *termination* is guaranteed, because every task will eventually receive barriers from all its inputs (including the back-end channels) and complete its snapshot. By broadcasting the barrier as soon as receiving it from all regular inputs, we avoid the deadlock condition mentioned previously.

The FIFO ordering property still holds for back-edges and the following properties prove *feasibility*. (1) Each task state included in the snapshot is a state copy of the respective task taken before processing any post-shot events from barriers received on regular inputs. (2) The downstream log included in the snapshot is complete and contains all pending post-shot records prior to barriers received on back-edges due to FIFO guarantees.

## 5. Failure Recovery

While not being the main focus of this work, a working failure recovery scheme motivates the application of our snapshotting approach. Thus, we provide a brief explanation here regarding its operation. There are several failure recovery schemes that work with consistent snapshots. In its simplest form the whole execution graph can be restarted from the last global snapshot as such: every task  $t$  (1) retrieves from persistent storage its associated state for the snapshot  $s_t$  and sets it as its initial state, (2) recovers its backup log and processes

---

**Algorithm 2** Asynchronous Barrier Snapshotting for Cyclic Execution Graphs
 

---

```

1: upon event  $\langle \text{Init} \mid \text{input\_channels},$ 
    $\text{backedge\_channels}, \text{output\_channels}, \text{fun},$ 
    $\text{init\_state} \rangle$  do
2:    $\text{state} := \text{init\_state}; \text{marked} := \emptyset;$ 
3:    $\text{inputs} := \text{input\_channels}; \text{logging} := \text{False};$ 
4:    $\text{outputs} := \text{output\_channels}; \text{udf} := \text{fun};$ 
5:    $\text{loop\_inputs} := \text{backedge\_channels};$ 
6:    $\text{state\_copy} := \text{Nil}; \text{backup\_log} := [];$ 
7:
8: upon event  $\langle \text{receive} \mid \text{input}, \langle \text{barrier} \rangle \rangle$  do
9:    $\text{marked} := \text{marked} \cup \{\text{input}\};$ 
10:   $\text{regular} := \text{inputs} \setminus \text{loop\_inputs};$ 
11:  if  $\text{input} \neq \text{Nil}$  AND  $\text{input} \notin \text{loop\_inputs}$  then
12:    trigger  $\langle \text{block} \mid \text{input} \rangle;$ 
13:    if  $\neg \text{logging}$  AND  $\text{marked} = \text{regular}$  then
14:       $\text{state\_copy} := \text{state}; \text{logging} := \text{True};$ 
15:      broadcast  $\langle \text{send} \mid \text{outputs}, \langle \text{barrier} \rangle \rangle;$ 
16:      for each  $\text{inputs}$  as  $\text{input}$ 
17:        trigger  $\langle \text{unblock} \mid \text{input} \rangle;$ 
18:
19:    if  $\text{marked} = \text{input\_channels}$  then
20:      trigger  $\langle \text{snapshot} \mid \{\text{state\_copy},$ 
    $\text{backup\_log}\} \rangle;$ 
21:       $\text{marked} := \emptyset; \text{logging} := \text{False};$ 
22:       $\text{state\_copy} := \text{Nil}; \text{backup\_log} := [];$ 
23:
24: upon event  $\langle \text{receive} \mid \text{input}, \text{msg} \rangle$  do
25:   if  $\text{logging}$  AND  $\text{node} \in \text{loop\_inputs}$  then
26:      $\text{backup\_log} := \text{backup\_log} :: [\text{input}];$ 
27:      $\{\text{state}', \text{out\_records}\} := \text{udf}(\text{msg}, \text{state});$ 
28:      $\text{state} := \text{state}';$ 
29:     for each  $\text{out\_records}$  as  $\{\text{output}, \text{out\_record}\}$ 
30:       trigger  $\langle \text{send} \mid \text{output}, \text{out\_record} \rangle;$ 
31:
32:

```

---

all contained records, (3) starts ingesting records from its input channels.

A partial graph recovery scheme is also possible, similarly to TimeStream [13], by rescheduling only upstream task dependencies (tasks that hold output channels to the failed tasks) and their respective upstream tasks up to the sources. An example recovery plan is shown in Fig. 4. In order to offer *exactly-once* semantics, duplicate records should be ignored in all down-




Figure 4: Upstream Dependency Recovery Scheme

stream nodes to avoid recomputation. To achieve this we can follow a similar scheme to SDGs [5] and mark records with sequence numbers from the sources, thus, every downstream node can discard records with sequence numbers less than what they have processed already.

## 6. Implementation

We contributed the implementation of the ABS algorithm to Apache Flink in order to provide exactly-once processing semantics for the streaming runtime. In our current implementation blocked channels store all incoming records on disk instead of keeping them in memory to increase scalability. While this technique ensures robustness, it increases the runtime impact of the ABS algorithm.

In order to distinguish operator state from data we introduced an explicit *OperatorState* interface which contains methods for updating and checkpointing the state. We provided *OperatorState* implementations for the stateful runtime operators supported by Apache Flink such as offset based sources or aggregations.

Snapshot coordination is implemented as an actor process on the job manager that keeps a global state for an execution graph of a single job. The coordinator periodically injects stage barriers to all sources of the execution graph. Upon reconfiguration, the last globally snapshotted state is restored in the operators from a distributed in-memory persistent storage.

## 7. Evaluation

The goal of our evaluation is to compare the runtime overhead of ABS with the globally synchronised snapshot algorithm employed in Naiad [11] and also to test the scalability of the algorithm to larger number of nodes.

### 7.1 Experiment Setup

The execution topology (Fig. 5) used for evaluation consists of 6 distinct operators with parallelism equal to the number of cluster nodes, which translates into a




Figure 5: Execution topology used for evaluation

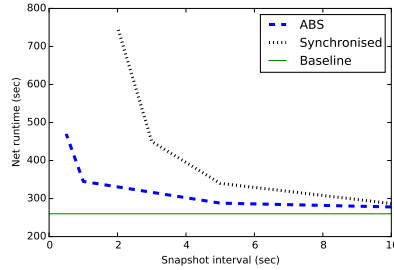


Figure 6: Runtime impact comparison on varying snapshot intervals

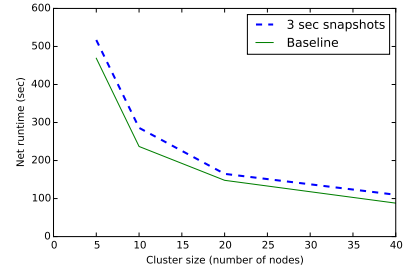


Figure 7: Runtime with ABS on different cluster sizes

6\*cluster size task vertices. The execution contains 3 full network shuffles in order to accentuate the possible impact of channel blocking in ABS. Sources generate a total of 1 billion records which are distributed uniformly among the source instances. The state of the operators in the topology were per-key aggregates and source offsets. The experiments were run on Amazon EC2 cluster using up to 40 m3.medium instances.

We measured the runtime overhead of our evaluation job running under different snapshotting schemes, namely ABS and synchronised snapshotting [11] with varying snapshot intervals. We implemented the synchronous snapshotting algorithm used in Naiad [11] on Apache Flink in order to have identical execution backend for the comparison. This experiment was run using a 10 node cluster. To evaluate the scalability of our algorithm, we processed a fixed amount of input records (1 billion) while increasing the parallelism of our topology from 5 up to 40 nodes.

## 7.2 Results

In Fig. 6 we depict the runtime impact of the two algorithms against the baseline (no fault tolerance). The large performance impact of synchronous snapshotting is especially visible when the interval of snapshotting is small. That is due to the fact that the system spends more time not processing any data, in order to obtain global snapshots. ABS has a much lower impact on the runtime as it runs continuously without blocking the overall execution, while maintaining a rather stable throughput rate. For larger snapshot intervals the impact of the synchronised algorithm is less significant since it operates in bursts (of 1-2 sec in our experiments) while letting the system run at its normal throughput during the rest of its execution. Nevertheless, bursts can often violate SLAs in terms of real-

time guarantees for many applications that are latency-critical, such as intrusion detection pipelines. Thus, such applications would further benefit by the performance of ABS. In Fig. 7 we compare the scalability of the topology running ABS with 3 second snapshot intervals against the baseline (without fault tolerance). It is clear that both the baseline job and ABS achieved linear scalability.

## 8. Future Work and Conclusion

In future work we are planning to explore the possibilities of further lowering the impact of ABS by decoupling snapshotting state and operational state. That allows for purely asynchronous state management as tasks can continuously process records while persisting snapshots. In such scheme there is also the need to synchronise pre-shot and post-shot records with the respective state which can be resolved by marking the records depending on the snapshot they belong to. As this approach would increase the computational, space and network I/O requirements of the algorithm, we plan to compare its performance to our current ABS implementation. Finally, we plan to investigate different recovery techniques that maintain exactly-once semantics while minimising the need for reconfiguration by operating on a per-task granularity.

In summary, we focused on the problem of performing periodic global snapshots on distributed dataflow systems. We introduced ABS, a new snapshotting technique that achieves good throughput. ABS is the first algorithm that considers the minimum state possible for acyclic execution topologies. Furthermore, we extend ABS to work with cyclic execution graphs by storing only the records that need to be reprocessed upon recovery. We implemented ABS on Apache Flink and evaluated our approach against synchronous snap-



shotting. At this early stage ABS shows good results, having low impact on the overall execution throughput with linear scalability.

## References

- [1] Apache flink. <https://flink.apache.org/>.
- [2] A. Alexandrov, R. Bergmann, S. Ewen, J.-C. Freytag, F. Hueske, A. Heise, O. Kao, M. Leich, U. Leser, V. Markl, et al. The stratosphere platform for big data analytics. *The VLDB Journal/The International Journal on Very Large Data Bases*, 23(6):939–964, 2014.
- [3] R. Castro Fernandez, M. Migliavacca, E. Kalyvianaki, and P. Pietzuch. Integrating scale out and fault tolerance in stream processing using operator state management. In *Proceedings of the 2013 ACM SIGMOD international conference on Management of data*, pages 725–736. ACM, 2013.
- [4] K. M. Chandy and L. Lamport. Distributed snapshots: determining global states of distributed systems. *ACM Transactions on Computer Systems (TOCS)*, 3(1):63–75, 1985.
- [5] R. C. Fernandez, M. Migliavacca, E. Kalyvianaki, and P. Pietzuch. Making state explicit for imperative big data processing. In *USENIX ATC*, 2014.
- [6] B. He, M. Yang, Z. Guo, R. Chen, B. Su, W. Lin, and L. Zhou. Comet: batched stream processing for data intensive distributed computing. In *Proceedings of the 1st ACM symposium on Cloud computing*, pages 63–74. ACM, 2010.
- [7] A. D. Kshemkalyani, M. Raynal, and M. Singhal. An introduction to snapshot algorithms in distributed computing. *Distributed systems engineering*, 2(4):224, 1995.
- [8] T. H. Lai and T. H. Yang. On distributed snapshots. *Information Processing Letters*, 25(3):153–158, 1987.
- [9] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.
- [10] Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein. Distributed graphlab: a framework for machine learning and data mining in the cloud. *Proceedings of the VLDB Endowment*, 5(8):716–727, 2012.
- [11] D. G. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, and M. Abadi. Naiad: a timely dataflow system. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pages 439–455. ACM, 2013.
- [12] R. Power and J. Li. Piccolo: Building fast, distributed programs with partitioned tables. In *OSDI*, volume 10, pages 1–14, 2010.
- [13] Z. Qian, Y. He, C. Su, Z. Wu, H. Zhu, T. Zhang, L. Zhou, Y. Yu, and Z. Zhang. Timestream: Reliable stream computation in the cloud. In *Proceedings of the 8th ACM European Conference on Computer Systems*, pages 1–14. ACM, 2013.
- [14] G. Tel. *Introduction to distributed algorithms*. Cambridge university press, 2000.
- [15] M. Zaharia, T. Das, H. Li, S. Shenker, and I. Stoica. Discretized streams: an efficient and fault-tolerant model for stream processing on large clusters. In *Proceedings of the 4th USENIX conference on Hot Topics in Cloud Computing*, pages 10–10. USENIX Association, 2012.