

Hazelcast jet

Low-latency stream processing at the 99.99th percentile

Gencer, Can; Topolnik, Marko; Ďurina, Viliam; Demirci, Emin; Kahveci, Ensar B.; Gürbüz, Ali; Lukáš, Ondřej; Fragkoulis, Marios; Katsifodimos, Asterios; More Authors

DOI

[10.14778/3476311.3476387](https://doi.org/10.14778/3476311.3476387)

Publication date

2021

Document Version

Final published version

Published in

Proceedings of the VLDB Endowment

Citation (APA)

Gencer, C., Topolnik, M., Ďurina, V., Demirci, E., Kahveci, E. B., Gürbüz, A., Lukáš, O., Fragkoulis, M., Katsifodimos, A., & More Authors (2021). Hazelcast jet: Low-latency stream processing at the 99.99th percentile. *Proceedings of the VLDB Endowment*, 14(12), 3110-3121. <https://doi.org/10.14778/3476311.3476387>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Hazelcast Jet: Low-latency Stream Processing at the 99.99th Percentile

Can Gencer, Marko Topolnik, Viliam Ďurina, Emin Demirci, Ensar B. Kahveci, Ali Gürbüz, Ondřej Lukáš, József Bartók, Grzegorz Gierlach, František Hartman, Ufuk Yılmaz, Mehmet Doğan, Mohamed Mandouh, Marios Fragkoulis*, Asterios Katsifodimos*

Hazelcast Inc.

*TU Delft

firstname@hazelcast.com

ABSTRACT

Jet is an open-source, high-performance, distributed stream processor built at Hazelcast during the last five years. Jet was engineered with millisecond latency on the 99.99th percentile as its primary design goal. Originally Jet's purpose was to be an execution engine that performs complex business logic on top of streams generated by Hazelcast's In-memory Data Grid (IMDG): a set of in-memory, partitioned and replicated data structures. With time, Jet evolved into a full-fledged, scale-out stream processor that can handle out-of-order streams and provide exactly-once processing guarantees. Jet's end-to-end latency lies in the order of milliseconds, and its throughput in the order of millions of events per CPU-core. This paper presents the main design decisions we made in order to maximize the performance per CPU-core, alongside lessons learned, and an empirical performance evaluation.

PVLDB Reference Format:

Can Gencer, Marko Topolnik, Viliam Ďurina, Emin Demirci, Ensar B. Kahveci, Ali Gürbüz, Ondřej Lukáš, József Bartók, Grzegorz Gierlach, František Hartman, Ufuk Yılmaz, Mehmet Doğan, Mohamed Mandouh, Marios Fragkoulis, Asterios Katsifodimos. Hazelcast Jet: Low-latency Stream Processing at the 99.99th Percentile. PVLDB, 14(12): 3110 - 3121, 2021. doi:10.14778/3476311.3476387

1 INTRODUCTION

Stream processing has completed more than twenty years of existence. Throughout the years, the database community has focused on defining semantics for streaming queries as well as system designs and optimizations. During the last seven years, stream processing is becoming mainstream: it is used for real-time monitoring, dynamic car-trip pricing, credit card fraud detection, real-time analytics, maintaining materialized views, and many other use cases and applications.

The needs of these new applications are starting to push existing streaming systems to their limits and each application may require a different system design. IoT applications and the banking industry, for example, have very different infrastructure and

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 14, No. 12
ISSN 2150-8097.
DOI:10.14778/3476311.3476387

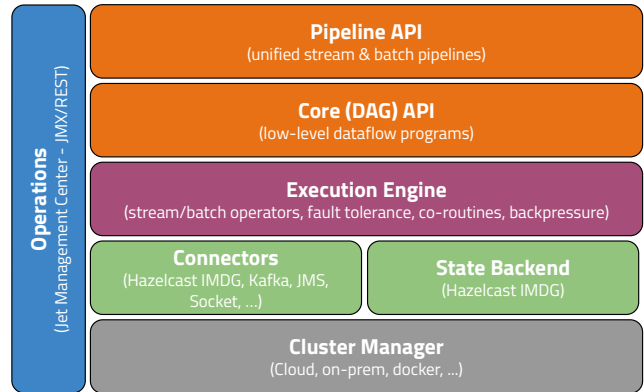


Figure 1: Jet's Software Stack

system design requirements. In an attempt to cover different application and deployment requirements, multiple streaming systems are being developed, each specializing in different mostly functional aspects. For instance, Apache Kafka [20] focuses on storing and integrating streams, Apache Flink [8] enables out-of-order [22] and exactly-once [7] processing with a very general programming model, and Heron [21] focuses on flexible container-based execution with at-least-once processing guarantees.

While most of these systems were designed with specific use cases in mind, their ability to adhere to strong SLAs has been an afterthought: the way that these systems are architected does not focus on low-latency deployments. When tested in practice, most existing scale-out stream processing systems fail to keep the 99.99th percentile latency at low levels (e.g. tens of milliseconds), even under ideal conditions (no failures, stragglers, etc.). In fact, even with resource over-provisioning, latency on the 99th percentile can easily reach seconds in state of the art stream processors [18]. We believe that this is the result of the design decisions behind these systems: in order to provide easy programming models and out-of-order processing [4] as well as flexibility in deploying streaming jobs in the cloud, existing streaming systems abstract away from the hardware and rely too much on the host language amenities such as JVM's default garbage collector.

At Hazelcast, around the end of 2015 we navigated the space of streaming systems and we concluded that none fulfilled our tight SLA requirements. In this paper we present Jet, a distributed streaming system that was designed with the latency at the 99.99th

*Work done while at Hazelcast.

```

Pipeline p = Pipeline.create();

p.readFrom(lineSource())
  .flatMap(e -> traverseArray(e.
    toLowerCase().split("\\W+")))
  .filter(word -> !word.isEmpty())
  .groupingKey(wholeItem())
  .aggregate(AggregateOperations.counting())
  .writeTo(someSink());

```

Listing 1: Word count in Jet’s Pipeline abstraction.

percentile as its primary focus. In the next paragraphs, we describe the motivation and discuss why we developed Jet.

Hazelcast’s In-memory Data Grid. Hazelcast IMDG is an open-source¹, distributed, in-memory object store supporting a wide variety of data structures and interfaces inspired by the standard Java libraries. IMDG’s data structures include Map, Queue, Ringbuffer, etc. For example, the Map interface, provides a distributed in-memory key-value store that Jet uses for storing its snapshots (Section 4.4). Hazelcast IMDG’s data structures are highly scalable and available: keys are stored in data partitions which are replicated across multiple machines in the cluster for fault resilience and parallelization. As IMDG members join and leave the cluster, partitions are re-balanced automatically. IMDG’s data structures have AP behavior [6]: they maintain linearizability when there’s no network partition and otherwise opt for availability.

From IMDG to Jet. With time, we noticed that more and more of IMDG’s users employed it for computation-heavy tasks. Although IMDG allows for very fast propagation of individual updates to replicas of its data structures, our clients wanted to perform more *global* operations: their tasks included joins among multiple distributed and replicated data structures, as well as aggregates that had to be updated whenever updates happened to different distributed IMDG data structures. To allow for such global queries, our team at Hazelcast first attempted to build a MapReduce-like prototype, only to realise after some time, that the requirements of our clients for very low latency were not satisfied by a batch-oriented architecture.

For this reason, that batch prototype was quickly discarded in favor of a full blown streaming dataflow engine that makes use of IMDG - both as a means to retrieve data but also to store data. To eliminate latency issues we decided that the partitioning of IMDG would have to *align* with the partitioning of our execution engine, to avoid costly and slow data re-partitioning. To this end, we built Jet², an open-source³ parallel dataflow processor that can execute custom business logic as well as SQL queries over data from an IMDG cluster.

Main Design Decisions. Jet was initially designed for continuous, low-latency computations over streams of changes, data structure updates, etc. The system design first focused on making operators as lightweight and performant as possible for scale-up computations on a single node [24]. At the same time, our primary development platform is the Java Virtual Machine (JVM) which is not ideal for low-latency systems. To this end, we had to make careful design

¹<https://github.com/hazelcast/hazelcast>

²<http://jet-start.sh>

³<https://github.com/hazelcast/hazelcast-jet>

decisions in order to: *i*) optimize the use of threads, *ii*) minimize the effects of the JVM and its garbage collector to the execution engine, and *iii*) minimize the system’s memory footprint. To be amenable to embedded deployments in constrained environments, our team has been working hard to make sure Jet does not have any external dependencies; Jet can be deployed as a single Jar file.

A single node running Jet has been shown to aggregate 10 million events per second with a 99.99th percentile latency below 10 milliseconds⁴. Jet can be deployed with minimal memory and computational requirements, while it can support hundreds of concurrent jobs within the same JVM (Section 7). In the rest of the paper we outline the design of Jet, its main contributions in systems-engineering as well as the lessons learned from building it.

This rest of the paper is structured as follows. Section 2 provides an overview of Jet’s core components: its programming model, execution engine, and state backend. Section 3 elaborates Jet’s execution model followed by Section 4, which presents Hazelcast’s in-memory grid and how it supports state management, fault tolerance, and reconfiguration. Section 5 describes how Jet’s components come together to achieve single-digit millisecond latency at the 99.99th percentile. Section 6 narrates some of the interesting use cases that Jet supports and Section 7 complements that description with a thorough experiment including performance microbenchmarks, fault tolerance measurements, and job multiplexing scenarios on queries of the NEXMark benchmark. Section 8 presents related work and how Jet resembles but also differs from existing systems. Notably, Section 9 introduces our vision for Jet as a serverless platform for running event-driven applications, such as microservices and stateful applications. Finally, Section 10 concludes this paper.

2 SYSTEM OVERVIEW

Jet adopts the standard streaming data flow model, in which computations form a directed acyclic graph of operators (*vertices*) that apply transformations on data streams (*edges*). Jet can operate equally well as an embedded stream processor with a very low footprint, as well as a scale-out system that spans many nodes.

Jet offers three main entry points for users (Figure 1): *i*) the Jet Management Center – a web UI and REST API from where users can manage and monitor Jet jobs, *ii*) the Pipeline API – a high-level API that can be used to create streaming pipelines and *iii*) the Core API – a lower-level API that can be used to specify low-level execution strategies. Both APIs are used to create distributed, stateful, and fault-tolerant dataflow programs. We detail those below.

2.1 The Pipeline API

The Pipeline API is designed to be the primary API that Jet users face. It offers transformations such as `map`, `filter`, `aggregate`, etc. that users can use to build complex data flows. It very much resembles Java streams in that, it is a fluent API and it is type-safe: the operators’ input and output types are checked at compile-time. You can control some of the lower-level aspects at this level, too, like operator parallelism and traffic rebalancing. The example in Listing 1, shows the Java API of the typical “Word Count” program.

⁴<https://jet-start.sh/blog/2020/08/05/gc-tuning-for-jet>

```

Pipeline p = Pipeline.create();
BatchStage<Person> persons = p.readFrom("persons");
BatchStage<Tuple2<Integer, Long>> countByAge = persons
    .groupingKey(Person::age)
    .aggregate(counting());
StreamStage<Order> orders = p.readFrom(kafka(...));
StreamStage<Entry<Order, Long>> ordersWithAgeCounts = orders
    .hashJoin(countByAge, joiningMapEntries(Order::ageOfBuyer))
    .writeTo(someSink());

```

Listing 2: Hybrid batch & streaming pipeline.

Unifying Batch & Streaming in a Single API. The Pipeline API models computations as *stages*; a stage represents what is known in different dataflow systems as an *operator*. Stages can be either *streaming* or *batch* and they can be mixed and matched to build hybrid stream-batch processing data flows. The difference between the two is rather simple: streaming stages assume that their inputs are infinite, while batch stages assume that their inputs are finite. For instance, as listed in Listing 2, one can build a pipeline that performs a hashjoin between a batch “build side” stage and a streaming “probing side” stage. The batch side will pull all the inputs of the batch stage performing complex batch filters, aggregates, etc. when the pipeline initializes, and then the stream will simply probe the hashtable for each incoming event on the streaming side.

In short, the Pipeline API serves as syntactic sugar that simplifies building scalable Jet workflows: pipelines are actually translated to parallel, distributed DAGs of operators at the Core API which we detail in the following text.

2.2 The Core API

The Core API exposes all the capabilities of the dataflow execution engine. It can be used to build DAGs of operators with a finer level of control than the Pipeline API, but at the cost of increased complexity and the lack of sanity checks such as type safety. Even though it is possible, this API is not intended to create DAGs by hand. Instead, it offers the infrastructure to build high-level DSLs and APIs that describe dataflows. It serves as a form of intermediate representation and requires familiarity with concepts like partitioning schemes, distributed vs. local edges, watermarks, etc.

The Core API can be used to fine-tune the execution properties of streaming DAGs. For instance, one can define the size of the queue between two vertices, a custom partitioning strategy of a distributed edge, or create custom code on how to react to an incoming watermark.

2.3 Execution Engine

The execution engine contains implementations of very efficient operators for partitioning, window aggregation [31, 33], joins, as well as the base source and sink operators, etc. Jet implements a unique way of sharing the computational resources, called tasklets (Section 3.2), a form of co-routines that are cooperatively scheduled in a fixed pool of worker threads as shown in Figure 4. The execution engine also performs state management (Section 4) and fault tolerance Section 4.4 that we present later in this paper.

2.4 State Backend: Hazelcast IMDG

Unlike most streaming systems that store their snapshots in stable object storage like Amazon’s S3, Jet uses IMDG for storing snapshots in a partitioned and replicated manner. The partitioning of the state stored in IMDG matches the partitioning of Jet. As a result IMDG’s replicated partitions allow Jet to *i)* retain its low latency by recovering very quickly from a node holding a state replica in case of failure, and *ii)* to be able to elastically scale-out when the workload increases. We detail these in Section 4.4.

3 EXECUTION MODEL

Jet’s execution model involves the deployment of stateful dataflow graphs that we explain in Section 3.1 and the execution of computations presented in Section 3.2. Section 3.3 clarifies how Jet handles backpressure.

3.1 Deploying Stateful Dataflow Graphs

The logical dataflow graph from the Core API is parallelized as in most streaming systems [5, 8], i.e., massively parallel operators like maps and filters simply receive arbitrary parts of the stream, while *keyed* operators such as joins, and aggregates execute in a partitioned manner. However, Jet does not follow the typical operator-per-core model. Instead, it deploys the *complete* dataflow graph on every available CPU core as seen in Figure 3. Since the source and sink connectors depend on 3rd-party APIs, and some of these rely on blocking calls that cannot be made to work cooperatively, Jet must start dedicated threads for them. After empirical experimentation we decided to deploy at most two such operators per node. For instance, Figure 3 depicts one source and one sink on each node.

Deploying the complete dataflow graph on every single CPU core has several advantages with respect to performance. The most important is that it allows to keep data exchange local to the machine as much as possible, avoiding expensive data transfers across the network. Second it allows for a coroutine-based execution that we discuss in the next section.

Jet further minimizes communication costs between operators in three ways. First, it fuses (a.k.a. operator chaining) consecutive stateless operators [15]. Second, it keeps data exchange local whenever possible, even on partitioned data exchanges, where it applies a two-stage approach (local partial results followed by global combining). Finally, when an operator needs to pass events to a remote operator, Jet deploys an exchange operator [14] on the sender and receiver sides that takes over the data partitioning and exchange.

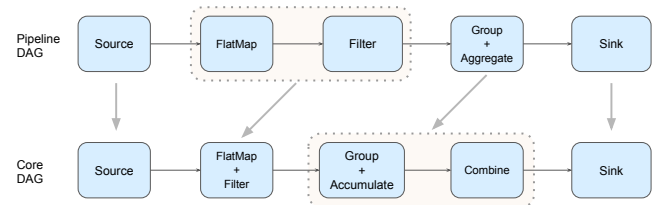


Figure 2: Logical Pipeline of a Jet job (top) and a Core DAG in which Jet applied operator fusion (bottom).

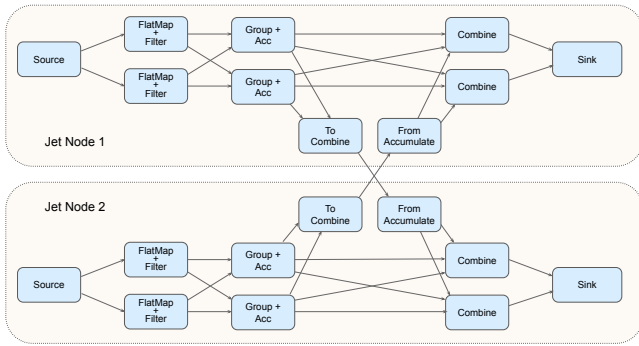


Figure 3: Dataflow deployed with local parallelism of 2.

In summary, Jet leverages both pipeline parallelism and data parallelism within each node but also across nodes. Together with coroutine-based execution (Section 3) Jet manages extremely low-latency stream processing.

3.2 Tasklets & Cooperative Threads

Jet’s execution model is based on *tasklets*, very small computational units that share an execution thread by taking turns in using a CPU core. As seen in Figure 4, Jet deploys as many JVM threads as there are CPU cores. Instead of allocating one thread for each parallel tasklet, a thread takes over the execution of a number of tasklets. On each thread, Jet runs a loop that executes its tasklets in a round-robin fashion.

Tasklets, are a concept similar to *coroutines* [17]: they are small computation units that can be suspended and resumed at the programming language level. More specifically, a tasklet voluntarily yields control to the Jet framework after executing for a very short period of time, typically under 1 millisecond. Then Jet can schedule another tasklet to run on the same thread. This design achieves high CPU utilization by avoiding the costly operating system context switches and allowing execution threads to remain on the same CPU core for longer time periods, thereby preserving the CPU cache lines. A tasklet is guaranteed to perform meaningful work (e.g., perform an aggregate or a join of two events) in each very short execution period because it has no dependencies to the outside world. Jet owes much of its performance capacity to tasklets.

Cooperative Threads. A collection of tasklets collectively occupy a *cooperative* thread [1]. Tasklets start execution when their turn comes, and stop execution on their own accord. Tasklets allow Jet to sustain high load and extremely high numbers of tasks on modest resources: Jet can host tens of thousands of tasklets on a single execution thread, allowing for very large numbers of operators in a single node. For the same reason, Jet lends itself really well to multiplexing different jobs in a single deployment: one can run thousands of concurrent jobs on the same set of resources with very little overhead. This is because when a tasklet has no work to do – typically because there is no input available – it backs off from the thread and gives its place to another tasklet, without the need for context switching.

Although very performant, the tasklets execution model entails an important restriction: a tasklet should never block, because that

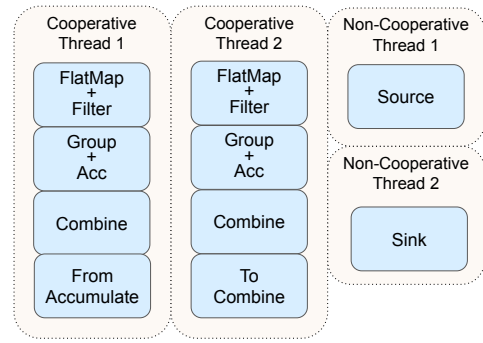


Figure 4: Tasklet scheduling in Jet

would jeopardize the execution progress of all tasklets sharing the execution thread. Although Jet could use a cooperative scheduler, in practice we have seen that simply iterating over all tasklets repeatedly works pretty well. However, Jet does schedule *blocking* computations on separate, *non-cooperative*, dedicated threads so that they do not interfere with the execution of tasklets. Non-cooperative computations are forced to frequently return control to Jet – at least every second. Jet needs to gain control of all execution threads in order to take state snapshots for fault tolerance strategy (Section 4).

Jet Processors. A *processor*, implements the custom logic (e.g., the UDF, or the operator logic such as joining two tuples) of a given DAG vertex in a Jet dataflow graph. Each processor includes an inbox of input records to be processed and an outbox of output records to be dispatched downstream. A tasklet manages the processor’s inbox and outbox, its state, and its inbound and outbound queues. A tasklet repeatedly calls the processor’s process method until the processor has consumed all the input records from the inbox, and then the tasklet refills the processor’s inbox with more input (possibly from a different input edge). For each output record in the outbox, Jet, according to the record’s key, finds the record’s downstream destination by computing the partition ID, and forwards that record to the respective downstream vertex.

Tasklets within the same node exchange data through shared-memory, single-producer-single-consumer queues that use wait-free algorithms. Specifically, a tasklet produces output to a queue that connects it with a specific consumer tasklet downstream. The consumer tasklet can retrieve its available input from the queue to fill in its processor’s inbox by pulling the records from these queues. The producer tasklet stops producing records if the consumer tasklet’s queue fills up. In essence, a queue instance serves the connection between a pair of tasklets, and on the receiver and producer sides Jet uses wait-free algorithms for predictable and performant data exchange.

3.3 Handling Backpressure

Jet employs a backpressure mechanism to limit the amount of items a source vertex sends to the destination/consumer when that consumer throttles. If this happens, streaming systems require a backpressure mechanism to signal back to the source to moderate its operation so that the whole pipeline stays in balance and operates

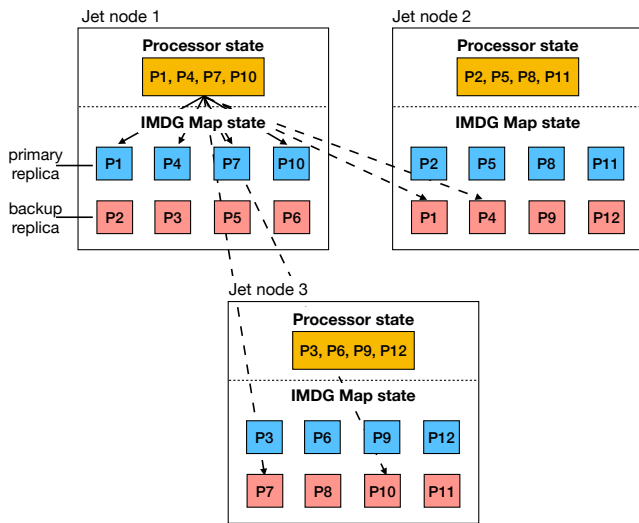


Figure 5: Snapshot partitioning and replication in IMDG.

at the speed of the slowest vertex. Local backpressure between tasklets inside the same Jet node is simple: Jet uses bounded queues and tasklets to back off as soon as all their output queues are full.

Backpressure is trickier over a network link. Jet uses a design very similar to the TCP/IP adaptive receive window: the producer must wait for an acknowledgment from the consumer specifying how many data items the producer can send. After processing item n , the receiver sends a message that the sender can send up to item $n + \text{receive_window}$. The consumer sends the acknowledgment message every 100ms: as long as the `receive_window` is large enough to hold the number of events processed within 100 milliseconds (plus network link latency), the receiver will always have data ready to be processed. Jet calculates the size of the `receive_window` based on the rate of event processing in a given tasklet and adaptively shrinks and expands the `receive_window` as the flow changes. In stable state the `receive_window` contains roughly 300 milliseconds' worth of data.

4 STATE MANAGEMENT

Jet stores state exclusively in memory, through the use of IMDG's IMap data structure. This section describes how Jet organizes computation state across a cluster (Section 4.1) and then presents the function of the in-memory grid and how it supports fault tolerance and reconfiguration (Section 4.2). Then, Sections 4.3 and 4.4 describe Jet's overall approach to elasticity and fault tolerance. Finally, Section 4.5 explains how Jet interacts with external source and sink systems to provide various levels of processing guarantees.

4.1 State Organization

State in Jet is a collection of key-value pairs, held in memory (Figure 5). Jet uses IMDG's IMap data structure to partition the key space into disjoint partitions that are allocated to parallel instances of a stateful vertex. Partitioned distributed edges connect keyed stateful computations such as aggregations, with all downstream vertices.

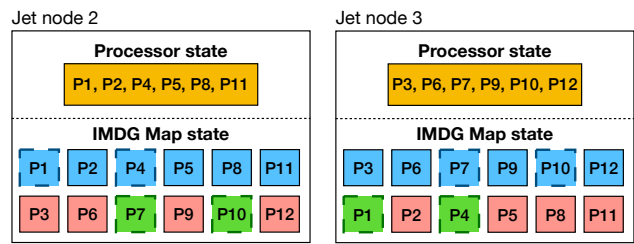


Figure 6: Snapshot re-balancing on fault recovery.

Stateful computations in Jet are supported by equally distributing disjoint partitions of the state to a stateful vertex's instances across all cluster nodes as shown in Figure 5. State is partitioned by a record key such that each vertex instance stores the state corresponding to a specific key space. Thus, each record is shipped through the edge that links with the vertex holding the partition that matches the record's key.

4.2 In-Memory Data Grid (IMDG)

One key feature that distinguishes Jet from other modern scale-out streaming systems is that, although fault-tolerant and highly available, Jet has zero dependency on disk storage. Instead, it uses the in-memory data grid, for state management and disaster recovery. The IMDG implements a concurrent, distributed, observable and queryable map data structure. Figure 5 depicts how IMDG supports a distributed dataflow graph.

A Jet processor operates directly on a local Map data structure (e.g., HashMap), state snapshots of which are stored in IMap. To optimize for locality, the partitioning of a Jet vertex matches the partitioning of the IMap data structure running on the same node. For each IMap partition, IMDG stores *i*) a *primary replica* on the same node and *ii*) one or more backup replicas on other nodes, according to the IMDG cluster configuration. The backup replicas of that state snapshot are distributed to other member nodes of the cluster for high availability. Figure 5 highlights how state objects of the processor instance in node 1 are stored in the local node's primary replicas matching the partitioning scheme used for both processing and state snapshot storage. Additionally, backup replicas of these state objects are sent to the 2nd and 3rd node.

When a member node in a Jet cluster fails, e.g. node 1 in Figure 5, the stored replicas of that node are lost together with the state of the Jet processor instances. In this case, IMDG has stored backup replicas of node 1, in nodes 2 and 3. To recover node 1, Jet instructs node 2 to promote partitions P1 and P4 from backup to primary as Figure 6 shows. Node 3 performs the same for partitions P7 and P10. Nodes 1 and 2 produce backup replicas for each other's newly added primary replicas. As a final step, Jet populates processor state from the corresponding primary replicas.

4.3 Elasticity and Reconfiguration

Beyond fault tolerance, IMDG is the key enabler for Jet's elasticity and reconfiguration functionality. When a new node enters a Jet cluster, Jet assigns replicas to that node in order to rebalance the

partitions. During the rebalancing phase, Jet minimizes data migration between the nodes employing consistent hashing [30]. For the job depicted in Figure 5, the new node is assigned the primary replica of P1 from node 1 and the backup replica of P2 again from node 1. Following the rebalancing phase, the job is restarted with processor states initialized from the local primary replicas of the latest state snapshot.

4.4 Fault Tolerance and Processing Guarantees

Jet’s approach to fault tolerance is based on the seminal work on distributed snapshots by Chandy and Lamport[11] that has been applied to stream processing [7, 8] and adopted by many streaming systems [8, 10, 16]. The approach involves periodic checkpoints that produce a consistent global snapshot of the system’s state.

At regular intervals, Jet instructs source vertices to take a state snapshot. Then, all processors belonging to source vertices save their state, emit a checkpoint barrier to the downstream processors through the data flow, and resume processing. When the checkpoint barrier reaches the sink processors, all processors have saved their state to stable storage and the state snapshot is complete. To recover from a failure, Jet will stop processing in all nodes and vertices, reload the latest state snapshots from IMDG recorded at the latest checkpoint, spawn a new instance to substitute the one that failed, and ask the input sources to replay the input data following the latest checkpoint.

The described protocol can offer *exactly-once* processing guarantees [7], i.e., each input will leave its effects on the system’s state exactly-once despite failures. Exactly-once consistency entails that no processing is performed while a checkpoint takes place. In the case of multiple input channels, once a checkpoint barrier arrives at one input channel, that channel needs to block and wait for all checkpoint barriers to arrive at the rest of the input channels of that vertex. For *at least-once* processing guarantees channels do not need to block, decreasing latency. Jet offers both levels of consistency as part of its configuration.

4.5 Assumptions and External Systems

Jet and most streaming systems [5, 8, 16, 27] rely on two assumptions for ensuring *exactly-once* guarantees. In essence, these assumptions complement the system’s fault tolerance approach. The assumptions regard the capabilities of systems functioning as data sources and data sinks to Jet.

Jet requires that source systems, which provide the input data, are *replayable* or *acknowledging*. A replayable source can replay the input data it produces from a specific offset. Jet needs this functionality when it is recovering from a failure in order to provide exactly-once processing guarantees as we described in Section 4.4. In fact, Jet’s source vertices take part in the checkpoints to store the source offset at which a checkpoint is taken so that a potential replay in case of a failure is precise.

If a source system is not replayable, but accepts acknowledgements that the data it stores can be safely deleted, then Jet can provide *exactly-once delivery* guarantee by acknowledging items only after they are processed by the entire pipeline and a successful snapshot has been taken. Since completing the snapshot and sending the acknowledgements isn’t atomic, it can happen that the

job fails before Jet acknowledges all the items. Because of this, Jet uses record IDs to deduplicate messages when the remote system re-sends unacknowledged messages after a recovery that had already been received by Jet prior to the failure. The record IDs are stored in a global state snapshot.

The exactly-once delivery guarantee can also be achieved on the output side of Jet if a sink system, which provides Jet’s output to a consuming application, functions as a *transactional sink* [7] or supports *idempotent writes*. A transactional sink withholds output and only makes it available to the outside world when a checkpoint is complete. In essence, it performs a two-phase commit on the output received by Jet. The commit-prepare phase executes when a checkpoint begins, with the second phase commit happening after the checkpoint is complete. At this point, the state corresponding to the output has been persisted.

Alternatively, the exactly-once delivery guarantee can be achieved if the sink system supports idempotent writes. Idempotent writes have the exact same effect irrespective of the number of times they are applied. Idempotent writes obviate the need for deduplication.

4.6 Fault Tolerance via Active Replication

From our experience with clients and Jet users we have noticed a very interesting trade-off: by implementing an algorithm for making snapshots in streaming computations such as the seminal Chandy-Lamport algorithm or one of its variants, the system has to pay the price of synchronizing barriers and blocking pipelines - especially to ensure exactly-once processing. During the Jet development journey we have made an important design decision: instead of focusing on creating low-latency snapshots, we opted for optimizing our system towards the efficient use of the available resources. In our experiments (Section 7) we show that Jet can easily handle a throughput of 2M events per second per CPU core, with very low overhead in latency. Thus, instead of running large deployments of a stream processor and requiring very efficient fault-tolerance mechanisms, we opted for enabling users to use less resources for a given workload, allowing them to run active-active deployments in which the job is executed twice (one active and one as active stand-by). The result is that in the absence of book-keeping and overhead for fault tolerance such a deployment can sustain failures, but it also performs extremely efficiently. In fact, this has proven quite a useful trade-off for our users and has simplified our system design. At the same time, we are considering alternative solutions like speculative replication [35].

5 MILLISECOND LATENCY ON THE JVM

The key competence of Jet is its capacity to perform ultra low-latency processing. Four design decisions drive this achievement. First, Jet schedules and executes threads at the programming language level leveraging the JVM in a way that resembles coroutines and green threads as we described in Section 3.2. Second, it minimizes the interference of concurrent garbage collection. Third, it utilizes the distributed IMDG to provide scalable and reliable state management in memory. Finally, Jet deploys a full dataflow graph instance to each member node of a cluster, thereby minimizing network connections for data transfers, which introduce latency.

Optimized data path. Jet optimizes the data path from source to sink: in fact, virtually all components involved in the datapath have been (re)implemented multiple times in the course of the last four years. Source vertices are local to each node and connect only with local vertices. Input data is received by each vertex's concurrent queues, which supply the data to the tasklet that will process it. Because tasklets do not depend on other computations, they can make progress in very short periods of time. When a tasklet's window of execution expires, it voluntarily yields control back to Jet, which selects the next tasklet to execute on the same thread without depending on operating system level schedulers. Periodically, Jet stores a state snapshot of the stateful vertex's state in the node's local IMDG space and replicates it to other nodes for high availability. Consequently, the datapath remains local to a node as much as possible passing through fast execution stages, which operate on data kept exclusively in-memory.

Garbage collection. A lot of attention and experimentation has been devoted to garbage collection in order to minimize its overhead to Jet's operation. Notably, garbage collection is recognized as one of the hidden performance enemies of stream processing, especially in terms of latency, due to its interference with the scheduling of computations.

As we described in Jet's execution model in Section 3, Jet employs a thread pool with as many threads as a member node's CPUs, however it can be configured to use a few cores less. The rationale for this choice is to allow garbage collection to be concurrent to the execution of tasklets by functioning as a background service without interfering with the scheduling of computation threads on a CPU. Dedicating a few threads to garbage collection and the remaining threads to computations optimizes the scheduling of threads to CPUs and ensures that each thread remains longer at a CPU. This design allowed Jet to bring its latency to under 10ms without impacting the throughput, as we show in the experiments.

6 USE CASES

Hazelcast clients and open-source users deploy Jet in many different ways; many of which we found surprising. Here we report on a subset of them that we believe are closest to the interest of the SIGMOD community, and were not typical (e.g., windowed aggregates) uses of streaming technology. We gathered these use cases from our internal engineering teams, our mailing lists, and Jet's Slack channels.

Real-time Rule Execution. A user from the banking industry aims at performing transactions. In order to check if a transaction is fraudulent or not, they need to first inspect multiple terabytes of data stored in a Hazelcast IMDG gathering ML features of a client and their previous transactions, before they can execute custom business logic, and decide whether a transaction is fraudulent. Because of the complexity of the process, Jet is assigned a maximum of 2ms for executing the complete set of tens of business rules. Similarly, another user from the banking industry combines complex business rules executing against the streaming operator state kept inside Jet operators.

Internet of Things. Since Jet is very lightweight and has very few dependencies, we have seen a set of interesting IoT deployments.

Deployments in factories and vehicles perform typical metrics aggregation (speed, temperatures, etc.) used for monitoring purposes, with relaxed latency requirements (hundreds of ms latency).

View Maintenance. We observed that a lot of Jet users use it to maintain views: they subscribe to different IMDG objects and external databases and capture changes in data items in the form of Change Data Capture (CDC)⁵. Subsequently, they consume the CDC stream and build materialized views that are updated with every change that happens to the external data sources.

Stateful AI. Finally, one of the most interesting use cases of Jet is its deployment as a stateful backend for a chatbot: the chatbot performs lookups for information that are related to humans' questions in a chatbox. The chatbot is deployed as an automaton where Jet operators are states and edges represent transitions. On each interaction with the human, the chatbot updates its state and responds to users. Our client scaled the chatbot to thousands of messages per second in a limited amount of computational resources.

Oil Rig Drilling. The operating costs of a drilling rig are very high and any downtime throughout the drilling process can have a significant effect on the rig operator's bottom line. The rigs are equipped with a large number of sensors to detect small vibrations during the drilling process. Hazelcast Jet enables human operators to immediately act on the streaming data in near real-time to prevent costly, catastrophic failures in the drilling process. It enables managing physical resources better through high-frequency feedback on a per-well basis, which amounts to reducing the drilling time by as much as 20 percent, from a typical 15 days to 12 days thereby saving up to millions of dollars.

The data input to Jet are generated from sensors. At any time during rig operations, up to 70 channels of high-frequency data enter Jet at various frequencies. Jet records and analyzes the data or events that occur, and then applies proprietary algorithms to make very fine-tuned adjustments to the drilling process. An example of this is the real-time adjustment of the revolutions per minute (RPM) of the drilling string and bit. A drill bit is a type of drill that must be operated at a very precise level to prevent equipment failure and costly delays to the drilling process. Jet computes stateful aggregates over 10K messages/second maintaining latency under 10ms. The workload resembles query 6 of NEXMark queries that we include in the experiments.

Real-time Payments. Hazelcast Jet is used as the backbone framework of an instant payment processing application. Within this application, Jet operates as a processing pipeline for each step of the payment process. The payment management application acts as an orchestrator which analyzes XML payment instructions and forwards them to the respective card's issuing bank or card association for verification, and also carries out a series of anti-fraud measures against the transaction before settling the payment transaction.

The computation is a stateful complex workflow of 150 stages in total, which communicate with external systems, such as databases. Multiple Jet processing jobs (tens of jobs) are involved as pipeline components where Jet's DAG API is used for stream processing while its pipeline API is used for batch processing. Jet's API offers considerable flexibility in expressing sophisticated multi-stage

⁵<https://debezium.io>

computations and combining stream and batch processing stages. The pipeline components rely on the distributed IMaps of Hazelcast IMDG for transaction ingestion and messaging through high-performance connectors that enable low-latency operations. Finally, the quick recovery mechanisms of the Jet cluster provides high availability to the instant payments application protecting it against failures.

The payments workload is CPU-bound with current throughput at the input side averaging at 10K transactions per second. It is served by a 3-node Jet cluster with 8GB RAM and an 8-core CPU per node. In this setting, throughput scales linearly with CPU cores until it reaches 65K operations per node. Then adding another node is necessary to keep the latency flat. Jet keeps latency stable at 85ms on the input and 150 ms on the output depending on the number of workflow steps in a specific transaction flow. Each workflow updates in-memory business entities at a rate of 75 updates per transaction on average.

7 EXPERIMENTS

In this section, we present the experiments we performed with Jet. First, we describe our experimental methodology (Section 7.1). Second, we show experiment results with respect to throughput (Sections 7.3, 7.4) and latency (Sections 7.2 and 7.5). Then, we measure the overhead of the fault tolerance approach to Jet’s normal operation in terms of latency (Section 7.6). Finally, we demonstrate Jet’s capacity for running a big number of jobs on a small cluster with low latency (Section 7.7).

7.1 Experiments Methodology

We evaluate Jet v4.3 in an Amazon Cloud cluster comprising 1, 5 and 10 c5.4xlarge machines with 16 vCPUs and 32 GB RAM each. Jet runs on Oracle OpenJDK v15.0.1 with 12 cooperative threads per node and the G1 garbage collector is configured with a GC pause target of at most 5 milliseconds. It does most of the GC work concurrently.

Query Workload. The workload used in the experiments consists of queries 1, 2, 5, and 8 of the NEXMark benchmark [34] as described in the Apache Beam project.⁶ The benchmark defines queries over people participating in auctions.

- Query 1 is a simple map converting an amount from one currency to another.
- Query 2 is a simple filter that selects auctions based on auction numbers.
- Query 3 is a join and filter that selects sellers in particular US states.
- Query 4 is a join with custom window functions and aggregation that reports the average selling price for each auction category.
- Query 5 is a sliding window with an aggregation to report which auctions have seen the most bids in a given window.
- Query 6 illustrates a specialized combiner that reports the average selling price per seller for their last ten closed auctions.

⁶<https://beam.apache.org/documentation/sdks/java/testing/nexmark/>

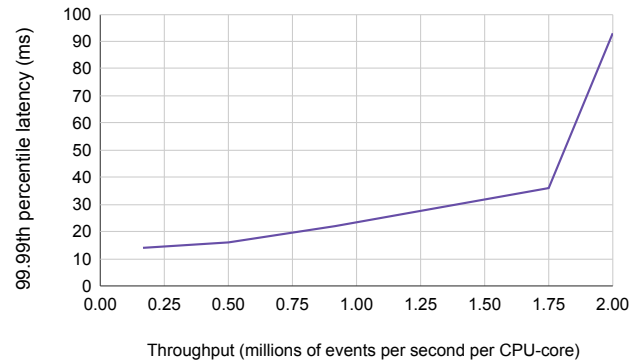


Figure 7: Throughput per CPU-core vs. Latency for Q5 on a single node (12 CPU cores) with 10ms window slide.

- Query 7 illustrates fanout using side input and selects the highest bids per period.
- Query 8 is a join of the stream of (new) users with the stream of auctions, reporting users that created an auction in the last period.
- Query 13 is a join between a stream of auctions and a bounded side-input.

Data Generator Configuration. We use the NEXMark data generator configured as follows. We define 10 thousand distinct keys that correspond to persons and auctions in the input dataset; we generate 1M records per second, by drawing keys randomly.

Window size and slide. In all windowed queries (Q5, 8, 13), unless stated otherwise, we use a 10-second sliding window join with a sliding step of 10ms. We choose to push Jet to its limits and show how it performs under pressure: triggering every 10ms is something that, to the best of our knowledge, no other scale-out stream processor can perform at the time of writing.

Metrics & Experiment Setup. We first let the JVM warm up by letting the streaming job run for 20 seconds, then we run the measurements over a period of 240 seconds. With 100 sliding window results per second, that gives us 24,000 data points. The primary metric of interest in our experiments is the processing *latency*. Specifically, the latency clock starts for each event at its predetermined time of occurrence. At that point the source tasklet is allowed to emit it, but any delay in actually emitting it is already affecting the reported latency number. If the event’s timestamp is such that it is the first event *after* a given window’s end time, it will trigger the aggregating stage to start emitting that window’s result. The clock stops when Jet has started emitting the window results.

In order to measure latency we control for input throughput, which we fix it at 1 million events per second. In the fault tolerance experiments of Section 7.6, we configure Jet to take a state snapshot every second and replicate the snapshots to another 1 member node of the cluster.

7.2 Latency Under Normal Operation

In this experiment we want to investigate the effect of scaling out to multiple machines (from 1 to 20), while keeping the input rate constant. The results are depicted in Figure 8. Overall, we observe

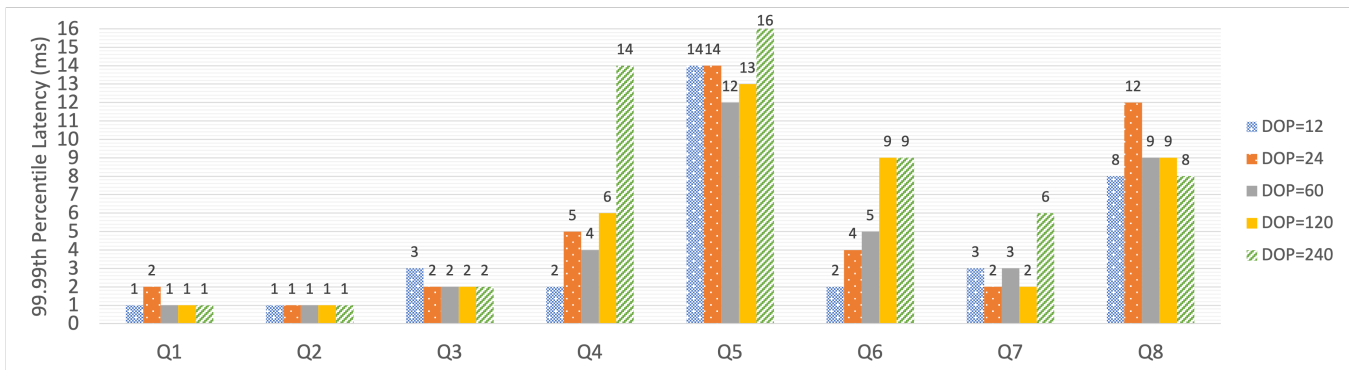


Figure 8: 99th percentile latency for all NEXMark queries for fixed input throughput of 1M events/s.

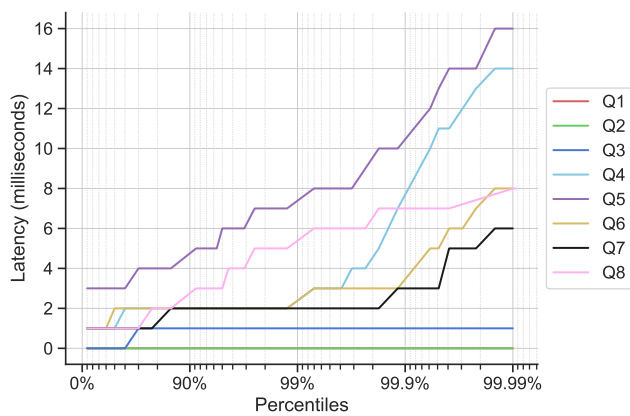


Figure 9: Distribution of latencies of all NEXMark queries for 1M events per second and DOP=240 (20 nodes).

that the latency at the 99.99th percentile never exceeds 16ms (in the case of DOP=240 on Q5). Moreover, we observe that Jet adds very little latency to simple queries (e.g., filters, etc.) while the most challenging one seems to be Q5 and Q8. Moreover, in Figure 9 we depict the distribution of latencies. We observe that 99.9th percentile latency is, in the worst case 10ms. The low latency, though, is to be expected as the input throughput was set well below the limits of what Jet can sustain. We do increase the throughput for Q5 in the next experiment, in order to see the effect to latency.

7.3 Throughput vs. Latency: 10ms Slide on Q5

In this experiment we want to push Jet to its limit with respect to throughput per CPU-core on a windowed aggregate of a tiny slide of 10ms on Query 5. We deployed Jet on a single node (c5.4xlarge, using 12 out of the 16 vCPUs) and used the size of the key set to vary the total input+output throughput. One sliding window consists of many key-value pairs, so the output throughput scales linearly with the key set size. We increased the total throughput gradually: from less than half a million items per second per CPU core, up to 2 million events per second per core. Our goal is to establish the maximum throughput that the system can sustain and the effect of throughput on the latency of the system. Figure 7 depicts the results. What we observe is that when throughput is

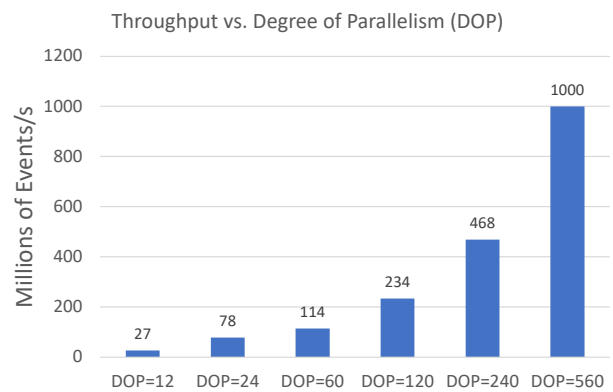


Figure 10: Throughput as we increase the cluster size from one VM (12 cores) to 45 VMs (560 cores), for Q5 with a sliding window of 500ms.

relatively low (around half a million events per second) Jet can sustain the system's latency very well: around 13ms on the 99.99th percentile. When pushed further, a single Jet processor on a single CPU core can sustain up to 2M events per second raising the 99.99th percentile latency up to 98ms. The take away message from this experiment is that for windowed aggregates, more than 1.75M events per second can put a burden on Jet's processors, increasing its latency quite considerably.

7.4 Throughput: 500ms Slide on Q5

In this experiment we wanted to evaluate Jet's ability to ingest massive data streams. To this end, we increased the slide of the windows in Q5 to 500ms, a very reasonable number for this type of queries. As seen in Figure 10 as we increase the cluster size from 12 cores on one VM to 45 VMs and a total of 560 cores, Jet manages to ingest up to 1 billion events per second (with 26ms latency 99th percentile latency). By further increasing the cluster size to 45 nodes, Jet In fact, this was possible due to the use of combiners as after some point the pre-aggregates on the keys (we are using 10K keys) reach a maximum number and the data exchanged is constant. At the same time, the 99.99th latency never exceeded 26ms. Note that this is possible due to the large window (500ms); in fact, in the previous experiment (Figure 7) with the 10ms sliding window,

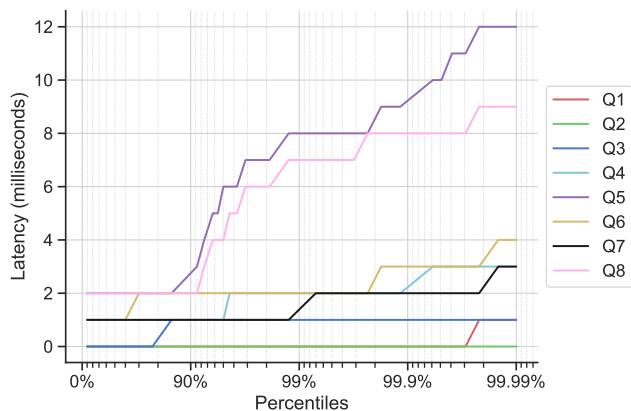


Figure 11: Latency for NEXMark queries, 5-node cluster.

we have seen a very big increase in the latency, as the throughput increased.

7.5 Latency: Windowed Aggregates & Joins

In this experiment we want to show the capacity of Jet to compute very low-latency aggregates and windowed joins. To this end, we have disabled fault tolerance: as explained in Section 4.6 a lot of Jet users prefer to execute jobs in an active-active fashion since Jet allows them to handle very high-throughput streams.

Figures 11 and 12 depict the latency for all queries on 5 and 10 nodes respectively. A general observation is that with simple queries that contain only `map` or `filter` Jet shows extremely low latency. This is to say that the 99.99th percentile stays at or below 1ms. At the same time, queries with joins exhibit 99.99th latencies in the order of 11-13ms. Bear in mind that the window triggers every 10ms, meaning that this is an extreme query example that exemplifies the capacity of Jet for low-latency processing. At the same time, more than 90% of the events exhibit a latency of 2ms or less for joins and equal or less than a millisecond for the windowed aggregates.

7.6 Latency: Fault Tolerance

Fault tolerance experiments suggest the overhead of Jet’s snapshot mechanism under normal operation. As seen in Figure 13 Jet’s latency at the 99.99th percentile when checkpoints are enabled is about 350ms. Latency remains very low for 70% of the events approximately, then spikes up to approximately 200ms at the 90%, and continues to rise sharply up to the 99%th percentile where it smoothens until it stabilizes at the 99.99th percentile.

As discussed in Section 4.6, the Jet team has not focused on further optimizing the fault-tolerance mechanism, but rather focused on optimizing the datapath under normal operation. We do have plans on optimizing the datapath with fault-tolerance enabled in the future, especially focusing on at-least once processing guarantees.

7.7 Latency: Multiplexing Jobs

Jet’s lightweight execution model allows numerous jobs to execute on the same thread without a considerable performance lag. This is a feature of Jet that comes for free, given its design with tasklets and

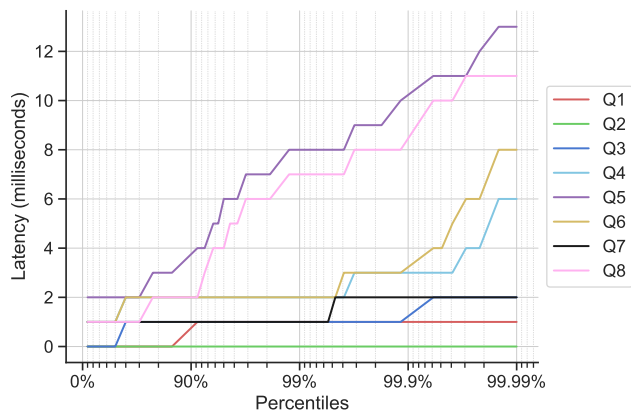


Figure 12: Latency for NEXMark queries, 10-node cluster.

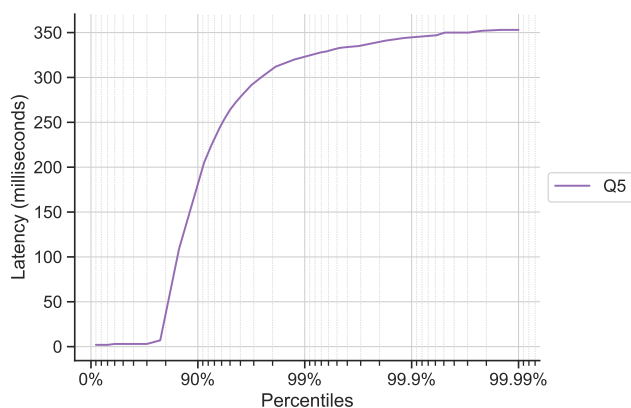


Figure 13: Latency in Query 5, with checkpointing enabled.

cooperative threads. In fact, we executed one hundred Query 5 jobs concurrently on a single node, to observe the effect of concurrent jobs to the latency of the system. We observed roughly 200ms 99.99th percentile latency, when running 100 concurrent jobs with an aggregate throughput of one million events per second.

8 RELATED WORK

A lot of streaming systems have made their debut in the previous years driving important advances. Some systems, such as Twitter’s Storm [32] and Heron [21], prioritize performance over consistency, while others like Google’s Millwheel [3] and Microsoft’s Timestream [28] and Streamscope [23] favor results with strict consistency guarantees. Academia has also developed novel streaming systems such as Apache Spark [5], Apache Flink [8], and Seep [12], which have paved new roads in fault tolerance, scalability, and performance. Notably, open source communities have contributed significantly to the success of such systems. At the same time, systems like Naiad [25] and IBM Streams [16] expand the traditional DAG streaming model of computation to include cycles for instance. Another class of systems such as Microsoft’s Trill [9] inspired a library implementation of a stream processor that can be easily embedded to an application. Finally, LinkedIn’s Samza [27] has been

built around the concept of a messaging system, i.e. Apache Kafka, which has comfortably found a place in the architecture of modern event-driven applications. A survey on the evolution of streaming systems with respect to out-of-order data management, state management, fault tolerance and high availability, and elasticity and reconfiguration is provided in reference [13].

Hazelcast Jet has certainly been inspired by existing streaming systems. The design of its programming model, out-of-order processing, fault tolerance, and elasticity are notable examples. On the other hand, because Jet is driven by the need for low-latency performance, it introduces two aspects that are novel in the stream processing domain. First, Jet involves an execution model based on user-level scheduling of cooperative threads that take turns on the same CPU, where each thread corresponds to an operator instance. With this design a single CPU typically runs an instance of the whole operator graph of a job and data transmissions over the network are minimized. Second, Jet performs state management completely in memory leveraging in-memory data replication to protect a job's checkpointed state from failures. These novel features enable Jet to provide reliably high performance that qualifies for strict service level agreements in production use cases.

9 EVENT-DRIVEN APPLICATIONS AS STATEFUL DATAFLOW GRAPHS

Jet aims at becoming an execution platform for event-driven applications [19], such as microservices [26] and stateful functions [2, 29], and offer a novel high-level programming model that allows users to convey the business logic of their application by means of standard functions. In this model, functions can simply call other functions and they can maintain state. Moreover, functions can be involved in a transaction with other functions while Jet guarantees that changes to the state are applied ensuring atomicity. Transactions can be marked by programmers with annotations and data flows and dependencies can be automatically derived with static analysis. Finally, Jet can take care of partitioning and replicating the state while keeping the operational state local to the function it belongs for fast access.

Jet's as a Backend for Stateful Functions. The inherent design of Jet offers a number of key features that make it a good candidate platform for running event-driven applications. First, Jet is particularly good at providing low-latency event processing, and keeps state local to computations. In addition, the IMDG can be used for fault tolerance, and especially high availability by ensuring that state is partitioned and replicated across the cluster. Furthermore, as we described in Section 3, Jet maintains a pool of threads that corresponds to the available CPUs of a member node. Since Jet can share a thread among numerous lightweight tasklets, it can overload a thread with tasklets scheduling only those that have pending requests and parking all others until a new request arrives for them. This design entails that Jet can deal with tens of thousands of tasklets on a single execution thread and take advantage of lightweight statistical multiplexing for e.g., stateful functions [2].

Open Problems. At the same time, Jet lacks a high-level *application* programming model and lacks support for transactions. In addition, dynamic reconfiguration, which is required for updating application components like microservices without disrupting the service, is

also missing. Finally, Jet requires a query interface that can perform ad-hoc queries or maintain views over distributed state much like a database system. This is just a small subset of a long list of open problems that we are currently tackling at Hazelcast.

10 CONCLUSIONS

In this paper we present Jet, a high performance distributed stream processor built from the grounds up with the goal of minimizing latency at the 99.99th percentile. Jet has been the favorite choice of customers for a big variety of use cases, such as banking transactions, internet of things, and stateful AI. We elaborate Jet's execution model and in-memory state backend, the IMDG, which are the key architectural components that make Jet special. In addition, we describe Jet as a full-fledged streaming system offering fault tolerance with varying consistency guarantees, elasticity and reconfiguration for scale-in scale-out actions, and out-of-order processing. Experiments with Jet show that it can support service-level agreements on less than 10ms latency at the 99.99th percentile. We envision the future of Jet as a serverless platform for running event-driven applications, like microservices and stateful functions.

REFERENCES

- [1] Martin Abadi and Gordon Plotkin. 2009. A model of cooperative threads. In *Proceedings of the 36th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. 29–40.
- [2] Adil Akhter, Marios Fragkoulis, and Asterios Katsifodimos. 2019. Stateful Functions as a Service in Action. *Proc. VLDB Endow.* 12, 12 (Aug. 2019), 1890–1893. <https://doi.org/10.14778/3352063.3352092>
- [3] Tyler Akidau, Alex Balikov, Kaya Bekiroğlu, Slava Chernyak, Josh Haberman, Reuven Lax, Sam McVeety, Daniel Mills, Paul Nordstrom, and Sam Whittle. 2013. MillWheel: fault-tolerant stream processing at internet scale. *Proceedings of the VLDB Endowment* 6, 11 (2013), 1033–1044.
- [4] Tyler Akidau, Robert Bradshaw, Craig Chambers, Slava Chernyak, Rafael J Fernández-Moctezuma, Reuven Lax, Sam McVeety, Daniel Mills, Frances Perry, Eric Schmidt, et al. 2015. The dataflow model: a practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *Proceedings of the VLDB Endowment* 8, 12 (2015), 1792–1803.
- [5] Michael Armbrust, Tathagata Das, Joseph Torres, Burak Yavuz, Shixiong Zhu, Reynold Xin, Ali Ghodsi, Ion Stoica, and Matei Zaharia. 2018. Structured Streaming: A Declarative API for Real-Time Applications in Apache Spark. In *Proceedings of the 2018 International Conference on Management of Data (Houston, TX, USA) (SIGMOD '18)*. ACM, New York, NY, USA, 601–613. <https://doi.org/10.1145/3183713.3190664>
- [6] E. Brewer. 2012. CAP twelve years later: How the "rules" have changed. *Computer* 45, 2 (2012), 23–29. <https://doi.org/10.1109/MC.2012.37>
- [7] Paris Carbone, Stephan Ewen, Gyula Fóra, Seif Haridi, Stefan Richter, and Kostas Tzoumas. 2017. State Management in Apache Flink&Reg: Consistent Stateful Distributed Stream Processing. *Proc. VLDB Endow.* 10, 12 (Aug. 2017), 1718–1729. <https://doi.org/10.14778/3137765.3137777>
- [8] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache Flink TM : Stream and Batch Processing in a Single Engine. *IEEE Data Eng. Bull.* 38 (2015), 28–38.
- [9] Badrish Chandramouli, Jonathan Goldstein, Mike Barnett, Robert DeLine, Danyel Fisher, John C Platt, James F Terwilliger, and John Wernsing. 2014. Trill: a high-performance incremental query processor for diverse analytics. *Proceedings of the VLDB Endowment* 8, 4 (2014), 401–412.
- [10] Badrish Chandramouli, Jonathan Goldstein, Mike Barnett, and James F. Terwilliger. 2015. Trill: Engineering a Library for Diverse Analytics. *IEEE Data Eng. Bull.* 38 (2015), 51–60.
- [11] K Mani Chandy and Leslie Lamport. 1985. Distributed snapshots: determining global states of distributed systems. *ACM Transactions on Computer Systems (TOCS)* 3, 1 (1985), 63–75.
- [12] Raul Castro Fernandez, Matteo Migliavacca, Evangelia Kalyvianaki, and Peter Pietzuch. 2013. 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 (New York, New York, USA) (SIGMOD '13)*. ACM, New York, NY, USA, 725–736. <https://doi.org/10.1145/2463676.2465282>

- [13] Marios Fragkoulis, Paris Carbone, Vasiliki Kalavri, and Asterios Katsifodimos. 2020. A Survey on the Evolution of Stream Processing Systems. arXiv:2008.00842 [cs.DC]
- [14] Goetz Graefe. 1990. Encapsulation of parallelism in the Volcano query processing system. *ACM SIGMOD Record* 19, 2 (1990), 102–111.
- [15] Martin Hirzel, Robert Soulé, Scott Schneider, Buğra Gedik, and Robert Grimm. 2014. A catalog of stream processing optimizations. *ACM Computing Surveys (CSUR)* 46, 4 (2014), 1–34.
- [16] Gabriela Jacques-Silva, Fang Zheng, Daniel Debrunner, Kun-Lung Wu, Victor Dogaru, Eric Johnson, Michael Spicer, and Ahmet Erdem Sariyüce. 2016. Consistent Regions: Guaranteed Tuple Processing in IBM Streams. *Proc. VLDB Endow.* 9, 13 (Sept. 2016), 1341–1352. <https://doi.org/10.14778/3007263.3007272>
- [17] Gilles Kahn and David MacQueen. 1976. Coroutines and networks of parallel processes. (1976).
- [18] Jeyhun Karimov, Tilmann Rabl, Asterios Katsifodimos, Roman Samarev, Henri Heiskanen, and Volker Markl. 2018. Benchmarking distributed stream data processing systems. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. IEEE, 1507–1518.
- [19] A. Katsifodimos and M. Fragkoulis. 2019. Operational Stream Processing: Towards Scalable and Consistent Event-Driven Applications. In *EDBT*.
- [20] Jay Kreps, Neha Narkhede, Jun Rao, et al. 2011. Kafka: A distributed messaging system for log processing. In *Proceedings of the NetDB*, Vol. 11. 1–7.
- [21] Sanjeev Kulkarni, Nikunj Bhagat, Maosong Fu, Vikas Kedigehalli, Christopher Kellogg, Sailesh Mittal, Jignesh M. Patel, Karthik Ramasamy, and Siddarth Taneja. 2015. Twitter Heron: Stream Processing at Scale. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Melbourne, Victoria, Australia) (SIGMOD '15)*. ACM, New York, NY, USA, 239–250. <https://doi.org/10.1145/2723372.2742788>
- [22] Jin Li, Kristin Tufte, Vladislav Shkapenyuk, Vassilis Papadimos, Theodore Johnson, and David Maier. 2008. Out-of-order processing: a new architecture for high-performance stream systems. *Proceedings of the VLDB Endowment* 1, 1 (2008), 274–288.
- [23] Wei Lin, Haochuan Fan, Zhengping Qian, Junwei Xu, Sen Yang, Jingren Zhou, and Lidong Zhou. 2016. STREAMSCOPE: Continuous Reliable Distributed Processing of Big Data Streams. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation (Santa Clara, CA) (NSDI'16)*. USENIX Association, Berkeley, CA, USA, 439–453. <http://dl.acm.org/citation.cfm?id=2930611.2930640>
- [24] Frank McSherry, Michael Isard, and Derek G Murray. 2015. Scalability! But at what {COST}?. In *15th Workshop on Hot Topics in Operating Systems (HotOS {XV})*.
- [25] Derek G Murray, Frank McSherry, Rebecca Isaacs, Michael Isard, Paul Barham, and Martín Abadi. 2013. Naiad: a timely dataflow system. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM, 439–455.
- [26] Sam Newman. 2015. *Building Microservices* (1st ed.). O'Reilly Media, Inc.
- [27] Shadi A. Noghabi, Kartik Paramasivam, Yi Pan, Navina Ramesh, Jon Bringham, Indranil Gupta, and Roy H. Campbell. 2017. Samza: Stateful Scalable Stream Processing at LinkedIn. *Proc. VLDB Endow.* 10, 12 (Aug. 2017), 1634–1645. <https://doi.org/10.14778/3137765.3137770>
- [28] Zhengping Qian, Yong He, Chunzhi Su, Zhuojie Wu, Hongyu Zhu, Taizhi Zhang, Lidong Zhou, Yuan Yu, and Zheng Zhang. 2013. TimeStream: Reliable Stream Computation in the Cloud. In *Proceedings of the 8th ACM European Conference on Computer Systems (Prague, Czech Republic) (EuroSys '13)*. ACM, New York, NY, USA, 1–14. <https://doi.org/10.1145/2465351.2465353>
- [29] Vikram Sreekanti, Chenggang Wu, Xiayue Charles Lin, Johann Schleier-Smith, Joseph E. Gonzalez, Joseph M. Hellerstein, and Alexey Tumanov. 2020. Cloudburst: Stateful Functions-as-a-Service. 13, 12 (2020). <https://doi.org/10.14778/3407790.3407836>
- [30] Ion Stoica, Robert Morris, David Karger, M Frans Kaashoek, and Hari Balakrishnan. 2001. Chord: A scalable peer-to-peer lookup service for internet applications. *ACM SIGCOMM Computer Communication Review* 31, 4 (2001), 149–160.
- [31] Kanat Tangwongsan, Martin Hirzel, and Scott Schneider. 2017. Low-latency sliding-window aggregation in worst-case constant time. In *Proceedings of the 11th ACM international conference on distributed and event-based systems*. 66–77.
- [32] Ankit Toshniwal, Siddarth Taneja, Amit Shukla, Karthik Ramasamy, Jignesh M. Patel, Sanjeev Kulkarni, Jason Jackson, Krishna Gade, Maosong Fu, Jake Donham, Nikunj Bhagat, Sailesh Mittal, and Dmitriy Ryaboy. 2014. Storm@Twitter. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (Snowbird, Utah, USA) (SIGMOD '14)*. ACM, New York, NY, USA, 147–156. <https://doi.org/10.1145/2588555.2595641>
- [33] Jonas Traub, Philipp M Grulich, Alejandro Rodríguez Cuéllar, Sebastian Breß, Asterios Katsifodimos, Tilmann Rabl, and Volker Markl. 2019. Efficient Window Aggregation with General Stream Slicing. In *EDBT*. 97–108.
- [34] P Tucker, K. Tufte, V. Papadimos, and D. Maier. 2002. *NEXMark—A Benchmark for Queries over Data Streams*. Technical Report. Technical report, OGI School of Science & Engineering at OHSU.
- [35] Ke Wang, Avriilia Floratou, Ashvin Agrawal, and Daniel Musgrave. 2020. Spur: Mitigating Slow Instances in Large-Scale Streaming Pipelines. In *Proceedings of the SIGMOD 2020 International Conference on Management of Data*. ACM, 2271–2285.