

# Stream-based State-Machine Replication

Laura Lawniczak and Tobias Distler

Friedrich-Alexander University Erlangen-Nürnberg (FAU)

Email: {lawniczak,distler}@cs.fau.de

**Abstract**—Developing state-machine replication protocols for practical use is a complex and labor-intensive process because of the myriad of essential tasks (e.g., deployment, communication, recovery) that need to be taken into account in an implementation. In this paper, we show how this problem can be addressed with stream-based replication, a novel approach that implements a replication protocol as application on top of a data-stream processing framework. With such framework already handling most essential tasks and furthermore providing means for debugging and monitoring, this technique has the key benefit of significantly minimizing overhead for both programmers as well as system operators. Our first stream-based protocol TARA tolerates crashes and comprises full-fledged mechanisms for request handling, checkpointing, and view changes. Still, TARA’s prototype implementation, which is based on Twitter’s Heron framework, consists of fewer than 1,500 lines of application-level code.

**Index Terms**—State-machine replication, fault tolerance, consensus, scalability, data-stream processing

## I. INTRODUCTION

State-machine replication protocols such as Paxos [1] or Raft [2] represent corner stones of many dependable services in production by enabling a system to tolerate crashes of participating processes. Unfortunately, implementing and operating these protocols usually is a difficult and time-consuming undertaking; not only due to the protocols’ inherent complexity [3], [4], but especially because there is a multitude of tasks that need to be taken care of in practice. Among other things, replicas for example must be installed and started on different servers, network connections have to be set up and maintained between nodes, exchanged messages need to be serialized and delivered to their intended receivers, and failed processes should be detected and recovered. All previously mentioned tasks have in common that they typically are not part of the replication-protocol logic and therefore further add to the complexity of the overall implementation. Nevertheless, since the tasks are essential for the execution of a replication protocol, they must be handled in some form or the other, for example by integrating external libraries (if possible) or by implementing them manually. Apart from complicating protocol development in general, this particularly makes it difficult to quickly create prototypes for testing new ideas.

In this paper, we present an approach that addresses these problems by implementing a state-machine replication protocol as application on top of a stream processing framework (e.g., Heron [5], Storm [6], or Flink [7]). Our choice of stream processing frameworks as underlying platform is motivated by several reasons: (1) With stream processing applications being widely used to analyze data, the frameworks are often already available and operational in many data centers. (2) Taking care

of tasks such as distributed deployment, communication, or the automated recovery of crashed application nodes, the frameworks provide many features that for replication protocols so far had to be specifically integrated. (3) Stream processing frameworks typically offer built-in support for scalability, which with our approach can be leveraged to improve the performance of replication protocols. (4) The frameworks are usually equipped with a logging infrastructure, means to collect runtime metrics (e.g., throughput), as well as graphical user interfaces, which each greatly facilitate the development and management of replication-protocol implementations.

Stream processing applications are designed as sets of processing nodes through which data tuples flow along the edges of directed acyclic graphs. This data-oriented perspective stands in sharp contrast with the replica-oriented perspective commonly used to specify state-machine replication protocols, where a small number of replicas repeatedly exchange messages in multiple phases [2], [4]. To show how to bridge this gap we present TARA, a *stream-based* replication protocol that has been specifically tailored to run in conjunction with stream processing frameworks. In order to be broadly applicable, TARA requires no modifications to the underlying platform and makes only weak assumptions about the services a framework provides with regard to deployment and communication. In particular, there is no need for the framework to implement consensus or replication-based fault tolerance at lower layers.

We implemented TARA based on Heron, a stream processing framework developed by Twitter for use in production. Thanks to leveraging Heron, TARA’s code base for request handling, checkpoint-based garbage collection, and view changes is about two thirds smaller than the implementation of the same tasks in the widely used replication library BFT-SMaRt [8].

In summary, this paper makes the following contributions: (1) It proposes an approach that facilitates the development and operation of replication protocols by implementing them as applications on top of stream processing frameworks. (2) It presents the design and implementation of TARA, the first stream-based replication protocol. (3) It uses TARA as an example to illustrate how stream-based protocols can leverage the underlying framework to achieve parallelism. (4) It evaluates TARA in the context of a coordination-service application.

The remainder is structured as follows: Section II introduces background on replication protocols and stream processing frameworks. Section III describes TARA with Section IV adding a parallelized version and Section V offering implementation details. After that, Section VI evaluates TARA. Finally, Sections VII and VIII discuss related work and conclude.

## II. BACKGROUND AND PROBLEM STATEMENT

This section provides background on replication protocols and stream processing frameworks and discusses the benefits and difficulties of integrating the former into the latter.

### A. Replication Protocols

State-machine replication protocols [1], [2] tolerate server crashes by modeling a system as a collection of replicas that each maintain an instance of the application state. To keep their state copies consistent, the replicas repeatedly execute a consensus protocol to agree on a common sequence in which to process newly incoming client requests. As illustrated in Figure 1, many protocols for this purpose assign replicas with different roles. A leader replica proposes a specific sequence number for a request, whereas its follower replicas are responsible for committing the sequence-number assignment.

In addition to consensus, replication protocols typically comprise further mechanisms that are essential for the well-functioning of a replicated system. Among other things, this specifically includes sub-protocols for checkpointing and view changes. While periodic checkpoints allow replicas to garbage collect consensus messages, a view-change mechanism enables a system to elect a new leader in case the old one is no longer able to fulfill its duties (e.g., due to having crashed).

### B. Stream Processing Frameworks

Stream processing frameworks such as Heron [5], Storm [6], or Flink [7] are widely used in production, especially for scenarios in which new data (e.g., recently posted tweets [5]) needs to be quickly analyzed once it becomes available. As shown in Figure 2, stream processing applications are implemented as a set of nodes that are organized in a directed acyclic graph through which information flows from one or more source nodes to one or more sink nodes. Each node in the graph (if needed) maintains its own state and represents a different stage in the processing pipeline. Typical tasks include the aggregation or filtering of inputs, the analysis of data, and the combination of the outputs of multiple predecessor nodes.

Between nodes, information is forwarded in the form of data tuples (i.e., collections of key-value pairs) that only flow in one direction: from upstream nodes to downstream nodes. Communication with the outside world is usually handled via message queues (see Figure 2), for example provided by systems such as Kafka [9] or Kestrel [10]. If necessary, as it is the case for iterative computations [7], these queues can also be used to reinsert processed data into the application by relying on the same queue as both output queue and input queue.




Fig. 1. Basic architecture of replicated systems

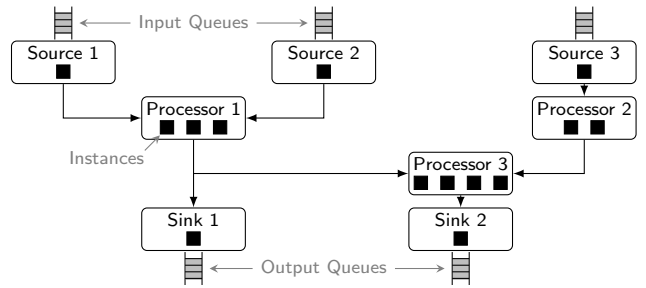


Fig. 2. Basic architecture of stream processing applications

From a system operator’s perspective, a stream processing framework offers several benefits with regard to executing applications. In the context of this paper, especially three aspects are of major importance: (1) **Deployment & Communication.** Provided with a logical graph of an application, the framework takes care of essential deployment tasks such as selecting servers, starting processing nodes, and setting up network connections. During execution, the framework then is responsible for serializing data tuples and routing them to their intended receiver nodes. In addition, frameworks commonly also comprise mechanisms for retransmitting tuples to tolerate network problems and for applying flow control to prevent nodes from being overwhelmed. (2) **Scalability.** To improve performance by introducing parallelism, the framework allows to partition the workload and execute multiple instances of a node, as illustrated by the black boxes in Figure 2. This way, a stream processing application is able to exploit multiple cores and servers even within the same processing stage. (3) **Automated Recovery.** Once an application is running, the framework continuously monitors the system for faults. In case it detects that a node instance has crashed, the framework automatically starts a new instance (possibly on a different server) and updates the routing configuration to ensure that the new instance is supplied with data.

### C. Problem Statement

Creating a replication-protocol implementation that is ready for use in practice is a difficult and time-consuming task since many deployment, scalability, and recovery aspects need to be considered that are usually not part of the theoretical protocol specification [3]. The current version of the widely used replication library BFT-SMaRt [8], for example, is the result of almost a decade of development, and even at this point it still does not support features such as the automated recovery of replicas. Our goal in this paper is to simplify the implementation of (existing and future) replication protocols by designing them as stream processing applications and thereby offloading most of the deployment and management tasks to the underlying framework. To be able to do so, we do not simply have to solve an engineering problem, but instead find a way to model replication protocols in the form of directed acyclic graphs, as required by stream processing frameworks. This is complicated by the fact that sub-protocols for consensus, garbage collection, and view changes are typ-

ically based on multiple phases of message exchange among the same group of replicas (see Section II-A) and heavily draw on the existing circular dependencies. As detailed in the following section, our solution to this problem is to abandon the replica-oriented perspective commonly applied in existing protocol specifications and move to a data-oriented perspective that primarily concentrates on how information needs to flow through the system. That is, instead of relying on a few large replicas with complex states, we model a replication protocol as a graph of small processing nodes with comparably simple states that each perform a different task.

### III. TARA

This section presents TARA, a replication protocol that has been tailored to run as an application on top of stream processing frameworks. In addition to the basic architecture, the section provides details on how TARA handles requests, garbage-collects consensus information, and performs view changes. In the following, we focus on giving an intuition of TARA’s core concepts. For the full protocol specification please refer to the extended version of the paper [11].

#### A. Overview

As illustrated in Figure 3, to meet the requirements of a stream processing application, the core of the TARA replication protocol is structured as a directed acyclic graph in which information is forwarded as streams of data tuples. Based on their responsibilities, the nodes in the graph can be divided into three main categories: request handling (yellow, see Section III-C), garbage collection (red, see Section III-D), and view change (blue, see Section III-E). The replicated service application is integrated with the *executor* stage.

For crash tolerance, each TARA stage consists of multiple node instances, which in the following we refer to as replicas. Replicas belonging to the same stage are placed on separate servers, whereas replicas of different stages may be co-located. The number of replicas required per stage in Figure 3 is symbolized by the number of black boxes. Two boxes indicate that a stage comprises  $f + 1$  replicas to tolerate up to  $f$  replica crashes within the stage; three boxes represent  $2f + 1$  replicas.

The numbers in Figure 3 mark TARA’s main workflow which consists of ① receiving client commands that arrive through request input queues, ② ordering these commands using a consensus algorithm, ③ executing the commands in the service application, and ④ placing the corresponding results in reply output queues. Clients with access to the queues are able to directly submit their commands to them and collect the results right away. All other clients typically communicate with TARA through *front-end* components that represent server-side proxies and act on the clients’ behalf. If a result to a submitted command does not arrive within a configured amount of time (e.g., due to the contacted front end having crashed), a client retries the operation by sending the command to another front end. As detailed in Section III-C, we designed TARA to deal with command retransmissions by enabling executor replicas to detect and filter out duplicates.

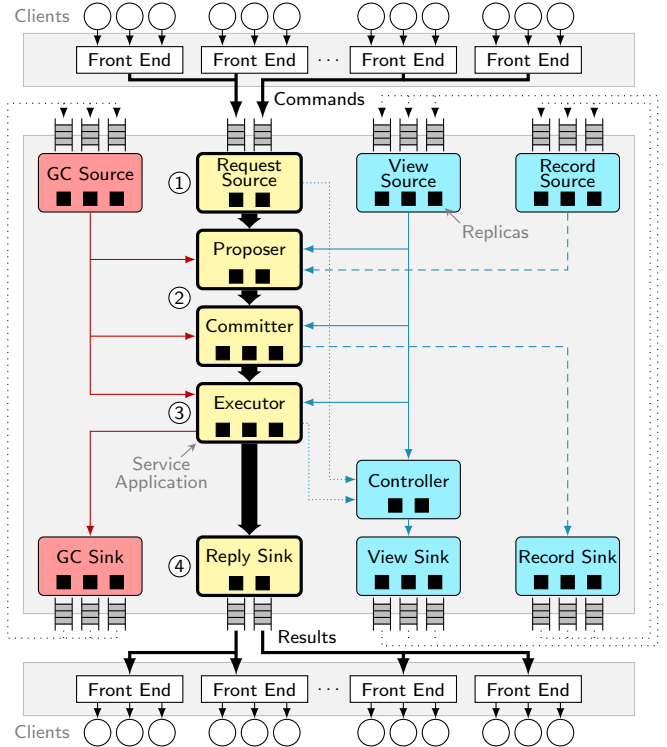


Fig. 3. Overview of the TARA replication protocol, comprising nodes for request handling (yellow), garbage collection (red), and view changes (blue).

In contrast to request handling, TARA’s mechanisms for garbage collection and view changes do not require external inputs, which is why their output queues are directly connected with their respective input queues to reinsert the emitted tuples back into the protocol (see Figure 3). Independent of their specific type, all queues in TARA are associated with a dedicated source or sink replica and therefore themselves do not have to provide any fault-tolerance guarantees. Specifically, we assume that a queue can crash as the result of a failure of its associated replica, and vice versa. Apart from using advanced queues such as Kafka [9], this assumption for example makes it also possible to implement a queue as a simple Web service that is integrated with its replica, as done in our prototype.

#### B. System Model

TARA makes only weak assumptions about the underlying stream processing framework and thus is compatible with a variety of existing systems (e.g., [5], [6], [7]). Specifically, it is not necessary for the framework to already provide consensus-based fault tolerance at a lower layer. At system startup, the framework needs to deploy all TARA replicas and ensure that the replicas of neighboring stages know each other. Once the system is running, the primary task of the framework is to handle the communication between replicas by routing the emitted tuples to their designated receivers. TARA allows the exchange of tuples to be performed using an asynchronous network but due to the FLP impossibility [12] needs a partially synchronous environment to guarantee eventual progress, as it is the case for other replication protocols [1], [2], [8].

TARA clients and replicas (where necessary) are equipped with retransmission mechanisms that are responsible for ensuring eventual delivery even if the framework itself does not offer reliable transmission of tuples. However, as discussed in Section V, TARA can exploit built-in features to improve efficiency in case a framework provides stronger semantics.

TARA tolerates up to  $f$  simultaneous replica crashes in each protocol stage. If the underlying framework comprises means to automatically recover replicas and their states after a failure, the protocol is able to reintegrate them and thereby self-heal.

### C. Request Handling

TARA relies on a Paxos-style consensus algorithm to ensure that its executors process client commands in the same order and thereby remain consistent. The algorithm is view based and comprises three phases. First, a *proposer* replica assigns monotonically increasing sequence numbers to newly incoming requests. In a second step, the requests are replicated to a set of *committers* that store and confirm the sequence-number assignments. Finally, executors apply the commands in the determined order and forward their results. For this algorithm to be safe, only one of the proposer replicas at a time may make new proposals. The active proposer  $p$  is determined by the current view  $v$  (e.g.,  $p = v \% (f + 1)$ ) and therefore loses its role as the result of a view change (see Section III-E).

**Consensus Process.** To invoke an operation  $op$  at the service application replicated by TARA, a client  $c$  creates a  $\langle \text{COMMAND}, c, t, op \rangle$  and then (either directly or via a front end) inserts the command into a request input queue;  $t$  is a logical or physical timestamp selected to be greater than all timestamps of previous commands issued by this client. In combination, client ID  $c$  and timestamp  $t$  enable TARA to uniquely identify a command (e.g., to filter out duplicates, see below). When request-source replica  $r$  removes the command  $cmd$  from its input queue, the replica wraps it in a  $\langle \text{REQUEST}, r, q, cmd \rangle$  tuple and sends the request to the proposer stage;  $q$  is a monotonically increasing request number that enables TARA to track the request-handling progress.

Whenever the active proposer  $p$  of the current view  $v$  receives a request  $req$ , it assigns a new sequence number  $s$  to the request and multicasts a  $\langle \text{PROPOSE}, p, s, v, req \rangle$  tuple to all committers. A committer  $k$  only accepts the proposal if it is also currently in view  $v$ . In such case, the committer locally stores the sequence-number assignment for the request and attests this step in a  $\langle \text{COMMIT}, k, s, v, req \rangle$  tuple to all executors. An executor accepts a request as soon as it has obtained  $f + 1$  commits from different committers for the same sequence number, provided that the commits have all been issued for the view the executor is currently in. At this point, the request is committed and the consensus process complete.

**Command Execution.** Executors perform duplicate detection to prevent the same command from being processed more than once. For this purpose, they maintain a vector  $\vec{T}_{exec}$  in which for each client they store the highest executed command timestamp. Furthermore, executors comprise a cache with the latest

results to be able to respond to duplicate commands without having to reexecute them. Relying on these data structures, an executor processes committed requests in the order of their sequence numbers by performing the following steps. First, the executor extracts the client command from the committed request. Next, it compares the command’s timestamp  $t_{cmd}$  to the client’s locally stored timestamp  $t_{exec}$ . If  $t_{cmd} \leq t_{exec}$ , the executor retrieves the result from the cache. Otherwise, the executor updates  $t_{exec}$ , invokes the command’s operation, and caches the computed result. Finally, the executor sends the result (via reply sinks and possibly front end) to the client.

### D. Garbage Collection

To prevent replicas from running out of memory, each protocol stage in TARA only maintains state for a limited amount of consensus instances, represented by a fixed-size window of sequence numbers whose lower bound is defined by a stability threshold  $s_{stable}$ . If its window is full, a replica temporarily suspends its participation in the consensus process until it learns that  $s_{stable}$  has increased. Raising the stability threshold is the main responsibility of TARA’s garbage-collection mechanism and triggered by periodic executor checkpoints of the application state. The rationale behind this approach is that once the effects of an executed command are reflected by a checkpoint, there is no longer a need to store consensus information about the command. As a result, replicas in such case are allowed to move their windows forward and garbage-collect information from all previous consensus instances.

**Creating Checkpoints.** Executors in TARA periodically create a checkpoint before processing a command with consensus sequence number  $s \% CP = 0$ ;  $CP$  is a configurable system-wide constant that represents the checkpoint interval. A checkpoint includes all essential information that is necessary to recreate the executor’s state at this sequence number, which includes a snapshot of the service application, the vector  $\vec{T}_{exec}$  that is used for duplicate detection (see Section III-C), as well as the result cache. Once the checkpoint is complete, the executor stores it at a predefined location. Stream processing frameworks such as Heron [5] for these purposes typically offer nodes access to local and remote file systems, and even cloud-based storage services. Having stored the checkpoint, the executor  $e$  in a last step sends sequence number  $s$  in a  $\langle \text{CHECKPOINT}, e, s \rangle$  tuple to TARA’s garbage-collection sinks from where the tuple is forwarded to the garbage-collection sources via message queues (see Figure 3).

**Updating the Stability Threshold.** Garbage-collection source replicas maintain a vector  $\vec{S}_{cp}$  that for each executor contains the highest checkpoint sequence number learned from checkpoint notifications. They select the stability threshold  $s_{stable}$  to be the  $f + 1$  highest element in  $\vec{S}_{cp}$  as this guarantees that at least one copy of the corresponding checkpoint remains available, even if up to  $f$  copies are no longer accessible (e.g., due to storage-server crashes). Whenever the stability threshold increases, a source replica  $g$  emits a  $\langle \text{STABLE}, g, s_{stable} \rangle$  tuple that is sent to all proposer, committer, and executor replicas.

Non-source replicas select the stability threshold as the  $f+1$  highest value provided by different sources, which ensures that other replicas will eventually learn the same threshold, even if up to  $f$  sources crash in the meantime. Each time the stability threshold increases, a replica adjusts its local consensus window accordingly and discards all information pertaining to lower sequence numbers. If an executor is lagging behind, for example as a result of asynchrony in the network, a window shift may cause the executor to skip sequence numbers. In order to catch up, the executor in such case first loads another executor’s checkpoint for the stability threshold before continuing to process further commands.

### E. View Change

TARA’s view-change mechanism enables the protocol to switch to another proposer in case the previously active proposer crashes. Decisions about whether a view change is needed or not are made by a set of *controller* replicas that continuously monitor the progress of the consensus process. If the consensus gets stuck, the controllers announce a new view, which is then installed by the replicas of other stages.

**Triggering a View Change.** The crash of the active proposer temporarily results in no new requests being proposed for ordering. Controllers in TARA are responsible for detecting such a scenario and for this purpose constantly compare (1) the number of incoming commands known to request sources with (2) the number of commands whose consensus process report the executors as complete. To provide the controllers with the necessary information, each request source  $r$  periodically emits a  $\langle \text{TARGET}, r, q \rangle$  tuple to all controllers, which includes the highest request number  $q$  the source has assigned to any request (see Section III-C). In a similar way, each executor  $e$  periodically reports the consensus progress by emitting an  $\langle \text{ACTUAL}, e, \vec{q} \rangle$  tuple;  $\vec{q}$  is a vector that for each request source contains the highest agreed request number. Relying on request numbers in the described way has the key benefit of enabling controllers to determine whether a request source has outstanding commands that so far have not been executed.

A controller triggers a view change if the number of completed requests does not increase on at least  $f+1$  executors for a configurable amount of time, even though the corresponding source has reported the existence of new commands. In such case, the controller  $x$  increments a local view counter and announces the new view  $v$  in a  $\langle \text{VIEW}, x, v \rangle$  tuple. From this point on, the system-wide publication of the new view via view sinks and sources follows the same principle as the distribution of the stability threshold, which was detailed in Section III-D.

**Entering a New View.** To ensure that the outcomes of already completed consensus instances remain stable across a view change, the newly appointed active proposer must learn about the requests that might have committed in previous views, and therefore possibly were processed by at least one executor. TARA solves this problem by requiring committers to maintain a record  $d = (s, v, req)$  for each sequence number  $s$  in their window; the record contains the most recent view  $v$  for which

a committer has received a proposal as well as the associated request  $req$ . Whenever a committer  $k$  learns about a higher view, it combines all of its records in a set  $\mathcal{D}$  and emits a  $\langle \text{RECORD}, k, \mathcal{D} \rangle$  tuple to TARA’s record sinks. The sinks rely on their output queues to forward the tuples to the record sources and finally to the active proposer of the new view.

Having been notified about a new view, the new proposer waits until it has obtained  $f+1$  record tuples from different committers. This guarantees that if a request had previously passed the consensus process (which requires confirmations from  $f+1$  committers, see Section III-C), at least one of the received record tuples must contain the sequence-number assignment for the request. For each sequence number included in record tuples, the new proposer selects the associated request with the highest view as new proposal. After this procedure is complete, the proposer is allowed to suggest new requests for all sequence numbers not covered by the record tuples.

## IV. PARALLELIZING TARA

This section discusses how to increase scalability in TARA by applying the concept of consensus-oriented parallelization [13]. The main idea of this approach is to (1) first distribute the responsibility for performing consensus across multiple partitions and then (2) afterwards deterministically merge the outcomes of different partitions into a single sequence of commands. Thanks to the underlying stream processing framework handling tasks such as deployment and communication, the integration of partitioning into TARA requires only minor modifications at the protocol level.

As shown in Figure 4, each partition comprises its own sets of proposer and committer replicas in order to be able to run the agreement process independently of other partitions. Since proposers are partition local, so are the view-change stages responsible for switching to another proposer replica in case the old one crashes. In contrast, to minimize overhead we enable all partitions to share the same garbage-collection replicas.

To ensure that all executors process requests in a consistent manner, they use a round-robin algorithm to compute a deterministic execution sequence number  $s_{exec} = s_i * P + i$ ,  $s_i \in \mathbb{N}$ ,  $i \in \{0, \dots, P-1\}$ , with  $s_i$  being the partition-local sequence number assigned to a command by a partition  $i$  and  $P$  denoting the number of partitions. Once a command has been processed in the application, an executor forwards the result to the reply sinks of the corresponding partition.




Fig. 4. Use of partitions to parallelize consensus in TARA.



## V. IMPLEMENTATION

Our TARA prototype implementation<sup>1</sup> is based on Heron [5], a stream processing framework that, for example, is currently used in production at Twitter. TARA offers similar client and application interfaces as existing replication libraries. A client or front end can asynchronously submit a request to any request source and will receive a result over the same channel. Since requests and results are represented as simple byte strings, both can have an arbitrary size and structure.

As typical for Heron, each TARA node (see Figure 3) runs in an independent process with its own Java virtual machine. This strongly improves the fault isolation in the system, as the failure of one node does not affect others, even when they run on the same server. It additionally allows for easy debugging, because each node can be accessed and analyzed individually.

In general, TARA makes it possible to distribute its nodes across a large number of servers. However, for better comparability with traditional replication libraries we co-locate replicas of different stages (e.g., a proposer, a committer, an executor) on the same server. Furthermore, for improved efficiency we integrate the reply sinks with their corresponding executors.

As discussed in Section III-B, the TARA protocol where necessary specifies its own retransmission mechanisms in order to ensure progress in the presence of an unreliable network. Implementing TARA on top of Heron, we are able to outsource some of the retransmission logic to the underlying framework by exploiting Heron’s built-in support for at-least-once delivery of tuples that originate in source nodes.

## VI. EVALUATION

This section evaluates our prototype implementation of TARA, while relying on the widely used replication library BFT-SMaRt [8] as baseline. As the name suggests, BFT-SMaRt was originally designed to tolerate Byzantine faults, however for a fair comparison with TARA we only evaluate the library in its later added configuration for crash tolerance.

### A. Development Effort

In order to get an impression of how our approach simplifies the development of replication protocols, we analyze the size of TARA’s code base compared with BFT-SMaRt (version 1.2).

**Methodology.** To ensure meaningful results, we only consider functionality that is present in both implementations. For BFT-SMaRt, this for example means that we exclude all code that is solely required and executed if the system is configured to tolerate Byzantine faults. Specifically, our analysis focuses on the core parts of the two protocols, namely the mechanisms for request handling, garbage collection, and view change.

Both the BFT-SMaRt and TARA implementation are written in Java and use a similar coding style, which allows us to perform the analysis based on the number of code lines spent on a specific functionality. Of course, examining code-line counts does not necessarily reveal all the complexity that might be contained in an implementation, however it gives us

the big picture of what is necessary to develop a replication infrastructure. Since we are interested in the effort it takes a programmer to implement a protocol, we exclude code lines that are trivial (e.g., due to only comprising a closing bracket), usually automatically inserted by a programmer’s development environment (e.g., import statements), or have no impact on the running system (e.g., empty lines and comments).

**Results.** Table I presents the findings of our analysis. Apart from the main mechanisms, we also report numbers for two additional categories: data structures that cannot necessarily be attributed to one specific mechanism, and code parts handling general infrastructure tasks such as system startup or the communication between replicas. Overall, the results show that thanks to leveraging Heron as underlying platform, TARA’s implementation is only about a third of the size of BFT-SMaRt.

Specifically, our analysis enables us to make three key observations: (1) Attributing code parts to one of the five categories was not always straightforward in case of BFT-SMaRt as the implementation often intertwines multiple mechanisms within the same class, occasionally even within the same method. In contrast, TARA’s design cleanly separates the different mechanisms by implementing them in dedicated nodes (as illustrated in Figure 3). For nodes that need to participate in more than one mechanism (e.g., executors), this separation is still visible within the node implementation in the form of different handlers for different types of incoming tuples (e.g., commit tuples or garbage-collection tuples). Altogether, TARA’s architecture makes it not only easier to analyze the code, but also significantly simplifies development, maintenance, and debugging. (2) A considerable amount of protocol-level code in BFT-SMaRt is spent on the synchronization of workflows that are implemented in different threads. TARA, on the other hand, leaves most of the concurrency handling to the underlying Heron and consequently minimizes the need for programmers to deal with synchronization themselves, a task that is notoriously error-prone. (3) BFT-SMaRt itself comprises means for typical infrastructure duties such as establishing network connections, dispatching received messages, and handling communication failures, all of which are examples of functionality that in TARA is provided by Heron. Therefore, most code lines in TARA dedicated to infrastructure are used to configure the Heron platform and specify the node graph representing the protocol. As a consequence, in the infrastructure category TARA allows a code-size reduction of about 67% compared with BFT-SMaRt.

TABLE I  
CODE-BASE COMPARISON BETWEEN BFT-SMaRt AND TARA

	BFT-SMaRt	TARA	Difference
Request Handling	905 LoC	274 LoC	-70%
Garbage Collection	551 LoC	185 LoC	-66%
View Change	867 LoC	297 LoC	-66%
Data Structures	756 LoC	421 LoC	-44%
Infrastructure	918 LoC	299 LoC	-67%
<b>Total</b> (LoC: Lines of Code)	3,997 LoC	1,476 LoC	-63%

<sup>1</sup>The prototype is publicly available under <https://gitlab.cs.fau.de/i4refit/tara>




Fig. 5. Throughput versus latency results for non-batched consensus.




Fig. 6. Throughput versus latency for batched consensus (batch size: 5).




Fig. 7. TARA throughput for different numbers of partitions.




Fig. 8. Impact of proposer failures on the throughput of BFT-SMaRt and TARA.

## B. Performance

To evaluate the performance of TARA in comparison with BFT-SMaRt, we conduct experiments with a coordination service [14] which we integrate with both systems. For the replicas we use a cluster of three servers (Intel Xeon CPU E3-1275, 3.6 GHz, 16 GB RAM); the clients are hosted by up to five additional machines. We configure the clients to submit requests in a closed loop, meaning that each client waits with the transmission of a new request until it has received the result to its previous one. The performance results reported in the following represent the average of three runs each.

**Non-batched Consensus.** In our first experiment, we configure BFT-SMaRt and TARA to order one client request per consensus instance, since this setting is a stress test for the request-handling mechanism and hence ideal for assessing the efficiency of the replication protocol. As shown in Figure 5, when we increase the workload BFT-SMaRt quickly reaches saturation at a throughput of about 4,300 requests per second. This behavior is caused by the fact that BFT-SMaRt’s consensus sub-protocol strictly sequentializes the agreement process by executing at most one consensus instance at a time. Consequently, once the system reaches a point at which there is always a consensus instance active, it is no longer able to further increase throughput. This bottleneck in turn causes the latency to increase considerably, as also visible in Figure 5.

Compared with BFT-SMaRt TARA starts with a higher latency, which is mainly a result of the inter-process communication overhead in Heron. However, in contrast to BFT-SMaRt, TARA is able to keep response times low at higher throughputs due to processing consensus instances in a pipelined manner. That is, as long as there are free slots in the sequence-number window (see Section III-D), the proposer in TARA can submit new proposals without the need to wait for the previous consensus instances to complete. This pipelining allows TARA to handle up to 13,000 requests per second in this setting, which is about three times the throughput of BFT-SMaRt.

**Batched Consensus.** To reduce the agreement overhead per request, replication protocols commonly offer the possibility to order a batch of multiple requests within the same consensus instance [8]. For our second experiment, we implement this technique in TARA by enabling request sources to combine incoming client commands and forward them as a single batch request. Figure 6 presents the results of this experiment with a maximum batch size of 5 in comparison with BFT-SMaRt.

In essence, the performance numbers show the same picture as in the non-batched consensus case, only at higher absolute throughput of close to 40,000 requests per second. This clearly shows that TARA can benefit from batching.

Besides batching, other common protocol-level optimizations such as read optimization or tentative execution could be integrated into TARA as well. On the other hand, optimizations relying on a monolithic system layout (e.g., fast paths between steps on the same replica) are not applicable to TARA as each step by design resides in its own node.

## C. Partitions

Parallelizing consensus, as done by the partitioned variant of TARA (see Section IV), is especially relevant in use-case scenarios in which the agreement process constitutes the performance bottleneck of the system [15], for example due to the nodes involved being resource constrained. In our third experiment, we create such a setting by individually throttling the proposer so that overall throughput is now at about 800 requests per second. As shown in Figure 7, in such a setting the use of multiple partitions enables TARA to scale by parallelizing consensus across a larger number of (also resource-constrained) nodes. Notice that BFT-SMaRt does not support partitioned consensus, which is why for this experiment we only report measurement results for TARA.

## D. Fault Tolerance

In our fourth experiment, we expose the two systems to leader failures in order to evaluate their view-change mechanisms. As shown in Figure 8, when we deliberately crash the current leader replica (i.e., the active proposer node in TARA) after 60 seconds of uninterrupted service, both systems trigger a view change to assign the leader role to a different replica. While without leader, during the view change no new requests can be agreed on and consequently the throughput temporarily drops to zero before eventually returning to a higher level. When we repeat the procedure with the newly elected leader 120 seconds into the experiment, BFT-SMaRt becomes unavailable since a single remaining replica is insufficient to safely make progress. In contrast, TARA is able to continue request processing even after the second proposer crash, because Heron in the meantime has automatically restarted the first proposer replica after having detected its failure. This example scenario illustrates a key benefit of implementing replication protocols on top of stream

processing frameworks such as Heron: TARA can leverage Heron’s support for automated recovery, a feature that is commonly not part of traditional replication libraries.

## VII. RELATED WORK

Replication protocols are notoriously difficult to implement, which is why several previous works have aimed at easing their development. For this reason, for example, a large body of secondary literature exists whose main purpose is to give advice on how to implement replication protocols such as Paxos based on their specification (e.g., [3], [4], [16]). Ongaro et al. [2] even went one step further by designing a replication protocol from scratch and targeting understandability as most important property. In this paper, we have shown stream-based replication to be an approach that effectively minimizes the number of problems a programmer needs to worry about when implementing a replication protocol. Still, the resulting system is able to match (or in some cases even exceed) the efficiency of traditional replication libraries, as confirmed by our evaluation.

Replication is an essential concept when it comes to providing fault tolerance in stream processing systems [17], [18], [19], [20], however notice that previous works in this area had a fundamentally different focus than TARA. While other researchers aimed at providing replication-based fault tolerance to stream processing applications, TARA itself is a stream processing application that offers fault tolerance to arbitrary network-based services. That is, instead of integrating replication mechanisms into the framework and tailoring them to the specific characteristics of stream processing applications [17], [18], [19], [20], TARA implements state-machine replication on top of such a framework in a generic manner.

With stream processing systems playing an important role in production, over the years significant efforts have been made to improve different aspects of these frameworks. Among other things, this includes approaches to optimize the deployment of processing nodes on the machines available [21], techniques to minimize energy consumption based on the current workload [22], and mechanisms to retain high quality-of-service levels even in the presence of load spikes [23], [24]. Since TARA itself is a stream processing application, it is able to benefit from many existing (and potentially future) improvements to underlying stream processing frameworks. For example, apart from the fact that replicas of the same stage must not be executed on the same server, TARA imposes no restrictions on the mapping of processing nodes to machines, and consequently can profit from optimized placement algorithms.

## VIII. CONCLUSION

In this paper we presented stream-based state-machine replication, an approach that leverages stream processing frameworks to significantly simplify the development, deployment, and operation of general-purpose replication protocols. The analysis of our first stream-based protocol TARA shows that our method reduces code size by about two thirds compared with the BFT-SMaRt library. Even though the additional layer of the stream processing engine leads to an increase in latency,

TARA is still able to sustain throughputs of tens of thousands of requests per second. Additionally, the automated recovery mechanism of the stream processing engine allows TARA to automatically restart and reintegrate failed leaders without any necessary manual actions.

*Acknowledgments:* This work was partially supported by the German Research Council (DFG) under grant no. DI 2097/1-2 (“REFIT”).

## REFERENCES

- [1] L. Lamport, “The part-time parliament,” *ACM Trans. on Computer Systems*, vol. 16, no. 2, pp. 133–169, 1998.
- [2] D. Ongaro and J. Ousterhout, “In search of an understandable consensus algorithm,” in *Proc. of USENIX ATC ’14*, 2014, pp. 305–320.
- [3] T. D. Chandra, R. Griesemer, and J. Redstone, “Paxos made live: An engineering perspective,” in *Proc. of PODC ’07*, 2007, pp. 398–407.
- [4] J. Kirsch and Y. Amir, “Paxos for system builders: An overview,” in *Proc. of LADIS ’08*, 2008, pp. 14–18.
- [5] S. Kulkarni, N. Bhagat, M. Fu, V. Kedighalli, C. Kellogg, S. Mittal, J. M. Patel, K. Ramasamy, and S. Taneja, “Twitter Heron: Stream processing at scale,” in *Proc. of SIGMOD ’15*, 2015, pp. 239–250.
- [6] A. Toshniwal, S. Taneja, A. Shukla, K. Ramasamy, J. M. Patel, S. Kulkarni, J. Jackson, K. Gade, M. Fu, J. Donham, N. Bhagat, S. Mittal, and D. Ryaboy, “Storm @Twitter,” in *Proc. of SIGMOD ’14*, 2014, pp. 147–156.
- [7] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, and K. Tzoumas, “Apache Flink: Stream and batch processing in a single engine,” *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering*, vol. 36, no. 4, 2015.
- [8] A. Bessani, J. Sousa, and E. E. P. Alchieri, “State machine replication for the masses with BFT-SMaRt,” in *Proc. of DSN ’14*, 2014.
- [9] J. Kreps, N. Narkhede, and J. Rao, “Kafka: A distributed messaging system for log processing,” in *Proc. of NetDB ’11*, 2011, pp. 1–7.
- [10] “Kestrel,” <https://github.com/twitter-archive/kestrel>.
- [11] L. Lawniczak and T. Distler, “Stream-based state-machine replication (extended version),” 2021, arXiv:2106.13019 [cs.DC].
- [12] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process,” *Journal of the ACM*, vol. 32, no. 2, pp. 374–382, 1985.
- [13] J. Behl, T. Distler, and R. Kapitza, “Scalable BFT for multi-cores: Actor-based decomposition and consensus-oriented parallelization,” in *Proc. of HotDep ’14*, 2014, pp. 49–54.
- [14] T. Distler, C. Cachin, and R. Kapitza, “Resource-efficient Byzantine fault tolerance,” *IEEE Trans. on Computers*, vol. 65, no. 9, pp. 2807–2819, 2016.
- [15] J. Behl, T. Distler, and R. Kapitza, “Consensus-oriented parallelization: How to earn your first million,” in *Proc. of Middleware ’15*, 2015, pp. 173–184.
- [16] T. Distler, “Byzantine fault-tolerant state-machine replication from a systems perspective,” *ACM Computing Surveys*, vol. 54, no. 1, 2021.
- [17] M. Balazinska, H. Balakrishnan, S. Madden, and M. Stonebraker, “Fault-tolerance in the Borealis distributed stream processing system,” in *Proc. of SIGMOD ’05*, 2005, pp. 13–24.
- [18] A. Martin, T. Knauth, S. Creutz, D. Becker, S. Weigert, C. Fetzer, and A. Brito, “Low-overhead fault tolerance for high-throughput data processing systems,” in *Proc. of SRDS ’11*, 2011, pp. 689–699.
- [19] X. Liu, A. Harwood, S. Karunasekera, B. Rubinstein, and R. Buyya, “E-Storm: Replication-based state management in distributed stream processing systems,” in *Proc. of ICPP ’17*, 2017, pp. 571–580.
- [20] A. Martin, A. Brito, and C. Fetzer, “Low cost synchronization for actively replicated data streams,” in *Proc. of LADC ’19*, 2019.
- [21] L. Aniello, R. Baldoni, and L. Querzoni, “Adaptive online scheduling in Storm,” in *Proc. of DEBS ’13*, 2013, p. 207–218.
- [22] C. Eibel, C. Gulden, W. Schröder-Preikschat, and T. Distler, “Strome: Energy-aware data-stream processing,” in *Proc. of DAIS ’18*, 2018, pp. 40–57.
- [23] N. Tatbul, U. Çetintemel, and S. Zdonik, “Staying FIT: Efficient load shedding techniques for distributed stream processing,” in *Proc. of VLDB ’07*, 2007, pp. 159–170.
- [24] N. Rivetti, Y. Busnel, and L. Querzoni, “Load-aware shedding in stream processing systems,” in *Trans. on Large-Scale Data-and Knowledge-Centered Systems XLVI*, 2020, pp. 121–153.