MIT Open Access Articles

ARES: Adaptive, Reconfigurable, Erasure Coded, Atomic Storage

Massachusetts Institute of Technology

Ares: **Adaptive, Reconfigurable, Erasure coded, Atomic Storage**[*]

NICOLAS NICOLAOU, Algolysis Ltd, Limassol, Cyprus,

VIVECK CADAMBE, Pennsylvania State University, US,

N. PRAKASH, Intel Corp.,

ANDRIA TRIGEORGI, University of Cyprus, Nicosia, Cyprus,

KISHORI M. KONWAR, MURIEL MEDARD, and NANCY LYNCH, Massachusetts Institute of Technology, USA,

Emulating a shared *atomic*, read/write storage system is a fundamental problem in distributed computing. Replicating atomic objects among a set of data hosts was the norm for traditional implementations (e.g., [11]) in order to guarantee the availability and accessibility of the data despite host failures. As replication is highly storage demanding, recent approaches suggested the use of erasure-codes to offer the same fault-tolerance while optimizing storage usage at the hosts. Initial works focused on a fix set of data hosts. To guarantee longevity and scalability, a storage service should be able to dynamically mask hosts failures by allowing new hosts to join, and failed host to be removed without service interruptions. This work presents the first erasure-code based atomic algorithm, called Ares, which allows the set of hosts to be modified in the course of an execution. Ares is composed of three main components: (i) a *reconfiguration protocol*, (ii) a *read/write protocol*, and (iii) a set of *data access primitives*. The design of Ares is modular and is such to accommodate the usage of various erasure-code parameters on a per-configuration basis. We provide bounds on the latency of read/write operations, and analyze the storage and communication costs of the Ares algorithm.

1 INTRODUCTION

Distributed Storage Systems (DSSes) store large amounts of data in an affordable manner. Cloud vendors deploy hundreds to thousands of commodity machines, networked together to act as a single giant storage system. Component failures of commodity devices, and network delays are the norm, therefore, ensuring consistent data-access and availability at the same time is challenging. Vendors often solve availability by replicating data across multiple servers. These services use carefully constructed algorithms that ensure that these copies are consistent, especially when they can be accessed concurrently by different operations. The problem of keeping copies consistent becomes even more challenging when failed servers need to be replaced or new servers are added, without interrupting the service. Any type of service interruption in a heavily used DSS usually translates to immense revenue loss.

---

[*]A preliminary version of this work appeared in ICDCS'19 [44]

Authors' addresses: Nicolas Nicolaou, nicolas@algolysis.comAlgolysis Ltd, Limassol, Cyprus,,; Viveck Cadambe, vxc12@engr.psu.eduPennsylvania State University, US,,; N. Prakash, prakashn@mit.eduIntel Corp.,,; Andria Trigeorgi, aatrige01@cs.ucy.ac.cyUniversity of Cyprus, Nicosia, Cyprus,,; Kishori M. Konwar, kishori@csail.mit.edu; Muriel Medard, medard@mit.edu; Nancy Lynch, lynch@csail.mit.eduMassachusetts Institute of Technology, USA,,

The goal of this work is to provide an algorithm for implementing strongly consistent (i.e., atomic/linearizable), fault-tolerant distributed read/write storage, with low storage and communication footprint, and the ability to reconfigure the set of data hosts without service interruptions.

Replication-based Atomic Storage. A long stream of work used replication of data across multiple servers to implement atomic (linearizable) read/write objects in message-passing, asynchronous environments where servers (data hosts) may crash fail [10, 11, 21–23, 25, 26, 40]. A notable replication-based algorithm appears in the work by Attiya, Bar-Noy and Dolev [11] (we refer to as the ABD algorithm) which implemented non-blocking atomic read/write data storage via logical timestamps paired with values to order read/write operations. Replication based strategies, however, incur high storage and communication costs; for example, to store 1,000,000 objects each of size 1MB (a total size of 1TB) across a 3 server system, the ABD algorithm replicates the objects in all the 3 servers, which blows up the worst-case *storage cost* to 3TB. Additionally, every write or read operation may need to transmit up to 3MB of data (while retrieving an object value of size 1MB), incurring high *communication cost.*

Erasure Code-based Atomic Storage. Erasure Coded-based DSSes are extremely beneficial to save storage and communication costs while maintaining similar fault-tolerance levels as in replication based DSSes [16]. Mechanisms using an $[n, k]$ erasure code splits a value v of size, say 1 unit, into k elements, each of size $\frac{1}{k}$ units, creates *n coded elements* of the same size, and stores one coded element per server, for a total storage cost of $\frac{n}{k}$ units. So the $[n = 3, k = 2]$ code in the previous example will reduce the storage cost to 1.5TB and the communication cost to 1.5MB (improving also operation latency). Maximum Distance Separable (MDS) codes have the property that value v can be reconstructed from any k out of these n coded elements; note that replication is a special case of MDS codes with $k = 1$. In addition to the potential cost-savings, the suitability of erasure-codes for DSSes is amplified with the emergence of highly optimized erasure coding libraries, that reduce encoding/decoding overheads [3, 12, 46]. In fact, an exciting recent body of systems and optimization works [7, 33, 46, 49, 52–54, 58] have demonstrated that for several data stores, the use of erasure coding results in lower latencies than replication based approaches. This is achieved by allowing the system to carefully tune erasure coding parameters, data placement strategies, and other system parameters that improve workload characteristics – such as load and spatial distribution. A complementary body of work has proposed novel non-blocking algorithms that use erasure coding to provide an atomic storage over asynchronous message passing models [13, 15, 16, 20, 34, 35, 56]. Since erasure code-based algorithms, unlike their replication-based counter parts, incur the additional burden of synchronizing the access of multiple pieces of coded-elements from the *same version* of the data object, these algorithms are quite complex.

Reconfigurable Atomic Storage. *Configuration* refers to the set of storage servers that are collectively used to host the data and implement the DSS. *Reconfiguration* is the process of adding or removing servers in a DSS. In practice, reconfigurations are often desirable by system administrators [9], for a wide range of purposes, especially during system maintenance. As the set of storage servers becomes older and unreliable they are replaced with new ones to ensure data-durability. Furthermore, to scale the storage service to increased or decreased load, larger (or smaller) configurations may be needed to be deployed. Therefore, in order to carry out such reconfiguration steps, in addition to the usual read and write operations, an operation called reconfig is invoked by reconfiguration clients. Performing reconfiguration of a system, without service interruption, is a very challenging task and an active area of research. RAMBO [39] and DynaStore [8] are two of the handful of algorithms [17, 24, 27, 32, 47, 48] that allows reconfiguration on live systems; all these algorithms are replication-based.

A related body of work appeared for *erasure coded scaling*, although there exists important differences that distinguish the two problems. In particular works like [30, 50, 51, 55, 57] consider RAID-based systems with synchronous network communication and local computation. Synchrony allows processes to make assumptions on the time of message delivery, and in turn help them to infer whether a communicating party has failed or not. On an asynchronous system, similar to the one we consider in this work, messages may be delivered with

arbitrary delays. Therefore, it is impossible to distinguish whether a message from a source is in transit or the source has crashed before sending a message. This uncertainty makes it impossible to detect failed from operating nodes, and thus challenging to design algorithms to guarantee atomicity (strong consistency) and completion of reads and writes.

Despite the attractive prospects of creating strongly consistent DSSes with low storage and communication costs, so far, no algorithmic framework for reconfigurable atomic DSS employed erasure coding for fault-tolerance, or provided any analysis of bandwidth and storage costs. Our paper fills this vital gap in algorithms literature, through the development of novel reconfigurable approach for atomic storage that use *erasure codes* for fault-tolerance. From a practical viewpoint, our work may be interpreted as a bridge between the systems optimization works [7, 33, 46, 49, 52–54, 58] and non-blocking erasure coded based consistent storage [13, 15, 16, 20, 34, 35, 56]. Specifically, the uses of our *reconfigurable* algorithm would potentially enable a data storage service to dynamically shift between different erasure coding based parameters and placement strategies, as the demand characteristics (such as load and spatial distribution) change, without service interruption.

Our Contributions. We develop a *reconfigurable, erasure-coded, atomic* or *strongly consistent* [29, 38] read/write storage algorithm, called Ares. Motivated by many practical systems, Ares assumes clients and servers are separate processes * that communicate via logical point-to-point channels.

In contrast to the replication-based reconfigurable algorithms [8, 17, 24, 27, 32, 39, 47, 48], where a configuration essentially corresponds to the set of servers that stores the data, the same concept for erasure coding need to be much more involved. In particular, in erasure coding, even if the same set of n servers are used, a change in the value of k defines a new configuration. Furthermore, several erasure coding based algorithms [15, 20] have additional parameters that tune how many older versions each server store, which in turn influences the concurrency level allowed. Tuning of such parameters can also fall under the purview of reconfiguration.

To accommodate these various reconfiguration requirements, Ares takes a modular approach. In particular, Ares uses a set of primitives, called *data-access primitives* (DAPs). A different implementation of the DAP primitives may be specified in each configuration. Ares uses DAPs as a "black box" to: (i) transfer the object state from one configuration to the next during reconfig operations, and (ii) invoke read/write operations on a single configuration. Given the DAP implementation for each configuration we show that Ares correctly implements a *reconfigurable, atomic* read/write storage.

Algorithm	#rounds /write	#rounds /read	Reconfig.	Repl. or EC	Storage cost	read bandwidth	write bandwidth
CASGC [14]	3	2	No	EC	$(\delta + 1)\frac{n}{k}$	$\frac{n}{k}$	$\frac{n}{k}$
SODA [34]	2	2	No	EC	$\frac{n}{k}$	$(\delta + 1)\frac{n}{k}$	$\frac{n^2}{k}$
ORCAS-A [20]	3	≥ 2	No	EC	n	n	n
ORCAS-B [20]	3	3	No	EC	∞	∞	∞
ABD [11]	2	2	No	Repl.	n	$2n$	n
RAMBO [39]	2	2	Yes	Repl.	$\geq n$	$\geq n$	$\geq n$
Dynastore [8]	≥ 4	≥ 4	Yes	Repl.	$\geq n$	$\geq n$	$\geq n$
SmartMerge [32]	2	2	Yes	Repl.	$\geq n$	$\geq n$	$\geq n$
Ares (this paper)	2	2	Yes	EC	$(\delta + 1)\frac{n}{k}$	$(\delta + 1)\frac{n}{k}$	$\frac{n}{k}$

Table 1. Comparison of Ares with previous algorithms emulating atomic Read/Write Memory for replication (Repl.) and erasure-code based (EC) algorithms. δ is the maximum number of concurrent writes with any read during the course of an execution of the algorithm. In practice, $\delta < 4$ [16].

The DAP primitives provide Ares a much broader view of the notion of a configuration as compared to replication-based algorithms. Specifically, the DAP primitives may be parameterized, following the parameters of protocols used for their implementation (e.g., erasure coding parameters, set of servers, quorum design, concurrency level, etc.). While transitioning from one configuration to another, our modular construction allows

---

*In practice, these processes can be on the same node or different nodes.

ARES to reconfigure between different sets of servers, quorum configurations, and erasure coding parameters. In principle, ARES even allows to reconfigure between completely different protocols as long as they can be interpreted/expressed in terms of the primitives; though in this paper, we only present one implementation of the DAP primitives to keep the scope of the paper reasonable. From a technical point of view, our modular structure makes the atomicity proof of a complex algorithm (like ARES) easier.

An important consideration in the design choice of ARES is to ensure that we gain/retain the advantages that come with erasure codes – cost of data storage and communication is low – while having the flexibility to reconfigure the system. Towards this end, we present an erasure-coded implementation of DAPs which satisfy the necessary properties, and are used by ARES to yield the first reconfigurable, *erasure-coded*, read/write atomic storage implementation, where read and write operations complete in *two-rounds*. We provide the atomicity property and latency analysis for any operation in ARES, along with the storage and communication costs resulting from the erasure-coded DAP implementation. In particular, we specify lower and upper bounds on the communication latency between the service participants, and we provide the necessary conditions to guarantee the termination of each read/write operation while concurrent with reconfig operations.

Table 1 compares ARES with a few well-known erasure-coded and replication-based (static and reconfigurable) atomic memory algorithms. From the table we observe that ARES is the only algorithm to combine a dynamic behavior with the use of erasure codes, while reducing the storage and communication costs associated with the read or write operations. Moreover, in ARES the number of rounds per write and read is at least as good as in any of the remaining algorithms.

We developed a *proof-of-concept* (PoC) implementation of ARES, and deployed it over a set of distributed devices in the experimental testbed Emulab [2]. The most important take home message from our experimental results is to show that it is possible to implement our algorithm according to the specifications and produces a correct execution and remains available during reconfiguration. Although, the correctness of the algorithm is shown analytically, the experimental validation corroborates the correctness. For this purpose, we have chosen simple parametarization (e.g., uniform selection of read/write invocation intervals), and picked ABD [11] as a benchmark algorithm which, despite being proposed more than 25 years ago, is the fundamental algorithm for emulating replicated quorum-based atomic shared memory. For instance, it is adopted in commercial/open-source implementations like Cassandra [36][†], and is being used as a standard benchmark algorithm (as can be seen in other recent works [19]). However, to demonstrate a real-world application we would need to compare with more algoritrhms and utilize a wide range of read/write distributions, and this is planned as a separate work.

Document Structure. Section 2 presents the model assumptions and Section 3, the DAP primitives. In Section 4, we present the implementation of the reconfiguration and read/write protocols in ARES using the DAPs. In Section 5, we present an erasure-coded implementation of a set of DAPs, which can be used in every configuration of the ARES algorithm. Section 7 provides operation latency and cost analysis, and Section 8 the DAP flexibility. Section 9 presents an experimental evaluation of the proposed algorithms. We conclude our work in Section 10. Due to lack of space omitted proofs can be found in [43].

2 MODEL AND DEFINITIONS

A shared atomic storage, consisting of any number of individual objects, can be emulated by composing individual atomic memory objects. Therefore, herein we aim in implementing a single atomic *read/write* memory object. A read/write object takes a value from a set $\mathcal{V}$. We assume a system consisting of four distinct sets of processes: a set $\mathcal{W}$ of writers, a set $\mathcal{R}$ of readers, a set $\mathcal{G}$ of reconfiguration clients, and a set $\mathcal{S}$ of servers. Let $\mathcal{I} = \mathcal{W} \cup \mathcal{R} \cup \mathcal{G}$ be the set of clients. Servers host data elements (replicas or encoded data fragments). Each writer is allowed to

---

[†]Cassandra [36] offers tuneable consistency, it uses protocol that is essentially ABD [11] for what they refer to as level 3 consistency (i.e., atomicity).

modify the value of a shared object, and each reader is allowed to obtain the value of that object. Reconfiguration clients attempt to introduce new configuration of servers to the system in order to mask transient errors and to ensure the longevity of the service. Processes communicate via *messages* through *asynchronous*, and *reliable* channels.

Configurations. A *configuration*, with a unique identifier from a set C, is a data type that describes the finite set of servers that are used to implement the atomic storage service. In our setting, each configuration is also used to describe the way the servers are grouped into sets, called *quorums*, s.t. each pair of quorums intersect, the consensus instance that is used as an external service to determine the next configuration, and a set of data access primitives that specify the interaction of the clients and servers in the configuration (see Section 3).

More formally, a configuration, $c \in C$, consists of: (*i*) $c.Servers \subseteq S$: a set of server identifiers; (*ii*) $c.Quorums$: the set of quorums on $c.Servers$, s.t. $\forall Q_1, Q_2 \in c.Quorums, Q_1, Q_2 \subseteq c.Servers$ and $Q_1 \cap Q_2 \neq \emptyset$; (*iii*) $DAP(c)$: the set of primitives (operations at level lower than reads or writes) that clients in I may invoke on $c.Servers$; and (*iv*) $c.Con$: a consensus instance with the values from C, implemented and running on top of the servers in $c.Servers$. We refer to a server $s \in c.Servers$ as a *member* of configuration c. The consensus instance $c.Con$ in each configuration c is used as a service that returns the identifier of the configuration that follows c.

Executions. An algorithm A is a collection of processes, where process A_p is assigned to process $p \in I \cup S$. The *state*, of a process A_p is determined over a set of state variables, and the state σ of A is a vector that contains the state of each process. Each process A_p implements a set of actions. When an action α occurs it causes the state of A_p to change, say from some state σ_p to some different state σ_p'. We call the triple $\langle \sigma_p, \alpha, \sigma_p' \rangle$ a *step* of A_p. Algorithm A performs a step, when some process A_p performs a step. An action α is *enabled* in a state σ if $\exists$ a step $\langle \sigma, \alpha, \sigma' \rangle$ to some state σ'. An *execution* is an alternating sequence of states and actions of A starting with the initial state and ending in a state. An execution ξ is *fair* if enabled actions perform a step infinitely often. In the rest of the paper we consider executions that are fair and well-formed. A process p *crashes* in an execution if it stops taking steps; otherwise p is *correct* or *non-faulty*. We assume a function $c.\mathcal{F}$ to describe the failure model of a configuration c.

Reconfigurable Atomic Read/Write Objects. A reconfigurable atomic object supports three operations: read(), write(v) and reconfig(c). A read() operation returns the value of the atomic object, write(v) attempts to modify the value of the object to $v \in \mathcal{V}$, and the reconfig(c) that attempts to install a new configuration $c \in C$. We assume *well-formed* executions where each client may invoke one operation (read(), write(v) or reconfig(c)) at a time.

An implementation of a read/write or a reconfig operation contains an *invocation* action (such as a call to a procedure) and a *response* action (such as a return from the procedure). An operation π is *complete* in an execution, if it contains both the invocation and the *matching* response actions for π; otherwise π is *incomplete*. We say that an operation π *precedes* an operation π' in an execution ξ, denoted by $\pi \to \pi'$, if the response step of π appears before the invocation step of π' in ξ. Two operations are *concurrent* if neither precedes the other. An implementation A of a read/write object satisfies the atomicity (linearizability [29]) property if the following holds [38]. Let the set Π contain all complete read/write operations in any well-formed execution of A. Then there exists an irreflexive partial ordering $\prec$ satisfying the following:

A1. For any operations π_1 and π_2 in Π, if $\pi_1 \to \pi_2$, then it cannot be the case that $\pi_2 \prec \pi_1$.
A2. If $\pi \in \Pi$ is a write operation and $\pi' \in \Pi$ is any read/write operation, then either $\pi \prec \pi'$ or $\pi' \prec \pi$.
A3. The value returned by a read operation is the value written by the last preceding write operation according to $\prec$ (or the initial value if there is no such write).

Storage and Communication Costs. We are interested in the *complexity* of each read and write operation. The complexity of each operation π invoked by a process p, is measured with respect to three metrics, during

the interval between the invocation and the response of π: (i) *number of communication round*, accounting the number of messages exchanged during π, (ii) *storage efficiency* (storage cost), accounting the maximum storage requirements for any single object at the servers during π, and (iii) *message bit complexity* (communication cost) which measures the size of the messages used during π.

We define the total storage cost as the size of the data stored across all servers, at any point during the execution of the algorithm. The communication cost associated with a read or write operation is the size of the total data that gets transmitted in the messages sent as part of the operation. We assume that metadata, such as version number, process ID, etc. used by various operations is of negligible size, and is hence ignored in the calculation of storage and communication cost. Further, we normalize both costs with respect to the size of the value v; in other words, we compute the costs under the assumption that v has size 1 unit.

Erasure Codes. We use an $[n, k]$ linear MDS code [31] over a finite field $\mathbb{F}_q$ to encode and store the value v among the n servers. An $[n, k]$ MDS code has the property that any k out of the n coded elements can be used to recover (decode) the value v. For encoding, v is divided into k elements $v_1, v_2, \ldots v_k$ with each element having size $\frac{1}{k}$ (assuming size of v is 1). The encoder takes the k elements as input and produces n coded elements $e_1, e_2, \ldots, e_n$ as output, i.e., $[e_1, \ldots, e_n] = \Phi([v_1, \ldots, v_k])$, where Φ denotes the encoder. For ease of notation, we simply write $\Phi(v)$ to mean $[e_1, \ldots, e_n]$. The vector $[e_1, \ldots, e_n]$ is referred to as the codeword corresponding to the value v. Each coded element c_i also has size $\frac{1}{k}$. In our scheme we store one coded element per server. We use Φ_i to denote the projection of Φ on to the i^{th} output component, i.e., $e_i = \Phi_i(v)$. Without loss of generality, we associate the coded element e_i with server i, $1 \leq i \leq n$.

Tags. We use logical tags to order operations. A tag τ is defined as a pair (z, w), where $z \in \mathbb{N}$ and $w \in \mathcal{W}$, an ID of a writer. Let $\mathcal{T}$ be the set of all tags. Notice that tags could be defined in any totally ordered domain and given that this domain is countably infinite, then there can be a direct mapping to the domain we assume. For any $\tau_1, \tau_2 \in \mathcal{T}$ we define $\tau_2 > \tau_1$ if (i) $\tau_2.z > \tau_1.z$ or (ii) $\tau_2.z = \tau_1.z$ and $\tau_2.w > \tau_1.w$.

For ease of reference, Table 2 presents the key notation used in this paper. Notice that some of the symbols shown are defined and used in following sections.

3 DATA ACCESS PRIMITIVES

In this section we introduce a set of primitives, we refer to as *data access primitives (DAP)*, which are invoked by the clients during read/write/reconfig operations and are defined for any configuration c in Ares. The DAPs allow us: (i) to describe Ares in a *modular* manner, and (ii) a cleaner reasoning about the correctness of Ares.

We define three data access primitives for each $c \in C$: (i) c.put-data($\langle \tau, v \rangle$), via which a client can ingest the tag value pair $\langle \tau, v \rangle$ in to the configuration c; (ii) c.get-data(), used to retrieve the most up to date tag and vlaue pair stored in the configuration c; and (iii) c.get-tag(), used to retrieve the most up to date tag for an object stored in a configuration c. More formally, assuming a tag τ from a set of totally ordered tags $\mathcal{T}$, a value v from a domain $\mathcal{V}$, and a configuration c from a set of identifiers C, the three primitives are defined as follows:

DEFINITION 1 (DATA ACCESS PRIMITIVES). *Given a configuration identifier $c \in C$, any non-faulty client process p may invoke the following data access primitives during an execution ξ, where c is added to specify the configuration specific implementation of the primitives:*

D1: c.get-tag() that returns a tag $\tau \in \mathcal{T}$;
D2: c.get-data() that returns a tag-value pair $(\tau, v) \in \mathcal{T} \times \mathcal{V}$,
D3: c.put-data($\langle \tau, v \rangle$) which accepts the tag-value pair $(\tau, v) \in \mathcal{T} \times \mathcal{V}$ as argument.

In order for the DAPs to be useful in designing the Ares algorithm we further require the following consistency properties. As we see later in Section 6, the safety property of Ares holds, given that these properties hold for the DAPs in each configuration.

$\mathcal{S}$	the set of server identifiers
$\mathcal{I}$	the set of client identifiers
$\mathcal{R}$	the set of reader identifiers in $\mathcal{I}$
$\mathcal{W}$	the set of writer identifiers in $\mathcal{I}$
$\mathcal{G}$	the set of reconfigurer identifiers in $\mathcal{I}$
$\mathcal{V}$	the set of values allowed to be written on the shared object
v	a value in $\mathcal{V}$
$\mathcal{T}$	the set of pairs in $\mathbb{N} \times \mathcal{W}$
τ	a pair $(z, w) \in \mathcal{T}$
C	the set of configuration identifiers
c	a configuration with identifier in C
$c.Servers$	the set of servers s.t. $c.Servers \subseteq \mathcal{S}$ in configuration c
$c.Quorums$	the set of subsets of servers s.t. $\forall Q \in c.Quorums, Q \subseteq c.Servers$ and $\forall Q_1, Q_2 \in c.Quorums, Q_1 \cap Q_2 \neq \emptyset$
σ	the state of an algorithm A
σ_p	the state of process $p \in \mathcal{I} \cup \mathcal{S}$ in state σ determined over a set of state variables
$p.var\vert_\sigma$	the value of the state variable var at process p in state σ
ξ	an execution of algorithm A which is a finite or infinite sequence of alternative states and actions beginning with the initial state of A
$\Phi([v_1, \dots, v_k])$ or $\Phi([v])$	the $[n, k]$ encoder function given k fragments of value v, $[v_1, \dots, v_k]$
e_i	the i^{th} encoded word, for $1 \leq i \leq n$, produced by $\Phi([v])$
$\mathcal{G}_L$	configuration sequence composed of pairs in $\{C \cup \{\bot\}\} \times \{F, P\}$, where F finalized and P pending, and initially contains $\langle c_0, F \rangle$

Table 2. List of Symbols used to describe our model of computation.

PROPERTY 1 (DAP CONSISTENCY PROPERTIES). *In an execution ξ we say that a DAP operation in an execution ξ is complete if both the invocation and the matching response step appear in ξ. If Π is the set of complete DAP operations in execution ξ then for any $\phi, \pi \in \Pi$:*

C1 *If ϕ is c.put-data($\langle \tau_\phi, v_\phi \rangle$), for $c \in C$, $\langle \tau_\phi, v_\phi \rangle \in \mathcal{T} \times \mathcal{V}$, and π is c.get-tag() (or c.get-data()) that returns $\tau_\pi \in \mathcal{T}$ (or $\langle \tau_\pi, v_\pi \rangle \in \mathcal{T} \times \mathcal{V}$) and ϕ completes before π is invoked in ξ, then $\tau_\pi \geq \tau_\phi$.*

C2 *If ϕ is a c.get-data() that returns $\langle \tau_\pi, v_\pi \rangle \in \mathcal{T} \times \mathcal{V}$, then there exists π such that π is a c.put-data($\langle \tau_\pi, v_\pi \rangle$) and ϕ did not complete before the invocation of π. If no such π exists in ξ, then (τ_π, v_π) is equal to (t_0, v_0).*

In Section 5 we show how to implement a set of DAPs, where erasure-codes are used to reduce storage and communication costs. Our DAP implementation satisfies Property 1.

As noted earlier, expressing Ares in terms of the DAPs allows one to achieve a modular design. Modularity enables the usage of different DAP implementation per configuration, during any execution of Ares, as long as the DAPs implemented in each configuration satisfy Property 1. For example, the DAPs in a configuration c may be implemented using replication, while the DAPs in the next configuration say c', may be implemented using erasure-codes. Thus, a system may use a scheme that offers higher fault tolerance (e.g. replication) when storage is not an issue, while switching to a more storage efficient (less fault-tolerant) scheme when storage gets limited.

In Section 8, we show that the presented DAPs are not only suitable for algorithm Ares, but can also be used to implement a large family of atomic read/write storage implementations. By describing an algorithm A according to a simple algorithmic template (see Algorithm 7), we show that A preserves safety (atomicity) if the used DAPs satisfy Property 1, and A preserves liveness (termination), if every invocation of the used DAPs terminates, under the failure model assumed.

4 THE ARES PROTOCOL

In this section, we describe Ares. In the presentation of Ares algorithm we decouple the reconfiguration service from the shared memory emulation, by utilizing the DAPs presented in Section 3. This allows Ares, to handle both the reorganization of the servers that host the data, as well as utilize a different atomic memory implementation per configuration. It is also important to note that Ares adopts a client-server architecture and separates the reader, writer and reconfiguration processes from the server processes that host the object data. More precisely, Ares algorithm comprises of three major components: (i) The reconfiguration protocol which consists of invoking, and subsequently installing new configuration via the reconfig operation by recon clients. (ii) The read/write protocol for executing the read and write operations invoked by readers and writers. (iii) The implementation of the DAPs for each installed configuration that respect Property 1 and which are used by the reconfig, read and write operations.

4.1 Implementation of the Reconfiguration Service.

In this section, we describe the reconfiguration service in Ares. The service relies on an underlying sequence of configurations (already proposed or installed by reconfig operations), in the form of a "distributed list", which we refer to as the *global configuration sequence (or list)* $\mathcal{G}_L$. Conceptually, $\mathcal{G}_L$ represents an ordered list of pairs $\langle cfg, status \rangle$, where cfg is a configuration identifier ($cfg \in C$), and a binary state variable $status \in \{F, P\}$ that denotes whether c is *finalized* (F) or is still *pending* (P). Initially, $\mathcal{G}_L$ contains a single element, say $\langle c_0, F \rangle$, which is known to every participant in the service.

To facilitate the creation of $\mathcal{G}_L$, each server in $c.Servers$ maintains a local variable $nextC \in \{C \cup \{\perp\}\} \times \{P, F\}$, which is used to point to the configuration that follows c in $\mathcal{G}_L$. Initially, at any server $nextC = \langle \perp, F \rangle$. Once $nextC$ is set to a value it is never altered. As we show below, at any point in the execution of Ares and in any configuration c, the $nextC$ variables of the non-faulty servers in c that are not equal to $\perp$ agree, i.e., $\{s.nextC : s \in c.Servers \wedge s.nextC \neq \perp\}$ is either empty of has only one element.

Clients discover the configuration that follows a $\langle c, * \rangle$ in the sequence by contacting a subset of servers in $c.Servers$ and collecting their $nextC$ variables. Every client in $\mathcal{I}$ maintains a local variable $cseq$ that is expected to be some subsequence of $\mathcal{G}_L$. Initially, at every client the value of $cseq$ is $\langle c_0, F \rangle$.

Reconfiguration clients may introduce new configurations, each associated with a unique configuration identifier from C. Multiple clients may concurrently attempt to introduce different configurations for same next link in $\mathcal{G}_L$. Ares uses consensus to resolve such conflicts: a subset of servers in $c.Servers$, in each configuration c, implements a distributed consensus service (such as Paxos [37], RAFT [45]), denoted by $c.Con$.

The reconfiguration service consists of two major components: (i) *sequence traversal*, responsible of discovering a recent configuration in $\mathcal{G}_L$, and (ii) *reconfiguration operation* that installs new configurations in $\mathcal{G}_L$.

Algorithm 1 Sequence traversal at each process $p \in \mathcal{I}$ of algorithm Ares.

---

    procedure read-config(*seq*)
2:    $\mu = \max(\{j : seq[j].status = F\})$
       $cs \leftarrow seq[\mu]$
4:    **while** $cs \neq \bot$ **do**
       $cs \leftarrow$ get-next-config($cs.cfg$)
6:    **if** $cs \neq \bot$ **then**
       $\mu \leftarrow \mu + 1$
8:    $seq[\mu] \leftarrow cs$
       put-config($seq[\mu - 1].cfg, seq[\mu]$)
10:   **end while**
     return *seq*
12: **end procedure**

    procedure get-next-config(*c*)
14:   **send** (READ-CONFIG) to each $s \in c.Servers$

    until $\exists Q, Q \in c.Quorums$ s.t. rec_i receives $nextC_s$ from $\forall s \in Q$
16:   **if** $\exists s \in Q$ s.t. $nextC_s.status = F$ **then**
      return $nextC_s$
18:   **else if** $\exists s \in Q$ s.t. $nextC_s.status = P$ **then**
      return $nextC_s$
20:   **else**
      return $\bot$
22: **end procedure**

    procedure put-config($c, nextC$)
24:   **send** (WRITE-CONFIG, $nextC$) to each $s \in c.Servers$
     until $\exists Q, Q \in c.Quorums$ s.t. rec_i receives ACK from $\forall s \in Q$
26: **end procedure**

Sequence Traversal. Any read/write/reconfig operation π utilizes the sequence traversal mechanism to discover the latest state of the global configuration sequence, as well as to ensure that such a state is discoverable by any subsequent operation π'. See Fig. 1 for an example execution in the case of a reconfig operation. In a high level, a client starts by collecting the $nextC$ variables from a quorum of servers in a configuration c, such that $\langle c, F \rangle$ is the last finalized configuration in that client's local $cseq$ variable (or c_0 if no other finalized configuration exists). If any server s returns a $nextC$ variable such that $nextC.cfg \neq \bot$, then the client (*i*) adds $nextC$ in its local $cseq$, (*ii*) propagates $nextC$ in a quorum of servers in $c.Servers$, and (*iii*) repeats this process in the configuration $nextC.cfg$. The client terminates when all servers reply with $nextC.cfg = \bot$. More precisely, the sequence parsing consists of three actions (see Alg. 1):

get-next-config(c): The action get-next-config returns the configuration that follows c in $\mathcal{G}_L$. During get-next-config(c), a client sends READ-CONFIG messages to all the servers in $c.Servers$, and waits for replies containing $nextC$ from a quorum in $c.Quorums$. If there exists a reply with $nextC.cfg \neq \bot$ the action returns $nextC$; otherwise it returns $\bot$.

put-config(c, c'): The put-config(c, c') action propagates c' to a quorum of servers in $c.Servers$. During the action, the client sends (WRITE-CONFIG, c') messages, to the servers in $c.Servers$ and waits for each server s in some quorum $Q \in c.Quorums$ to respond.

read-config(seq): A read-config(seq) sequentially traverses the installed configurations in order to discover the latest state of the sequence $\mathcal{G}_L$. At invocation, the client starts with the last finalized configuration $\langle c, F \rangle$ in the given seq (Line A1:2), and enters a loop to traverse $\mathcal{G}_L$ by invoking get-next-config(), which returns the next configuration, assigned to $\widehat{c}$. While $\widehat{c} \neq \bot$, then: (a) $\widehat{c}$ is appended at the end of the sequence seq; (b) a put-config action is invoked to inform a quorum of servers in $c.Servers$ to update the value of their $nextC$ variable to $\widehat{c}$. If $\widehat{c} = \bot$ the loop terminates and the action read-config returns seq.

Reconfiguration operation. A reconfiguration operation reconfig(c), $c \in C$, invoked by any reconfiguration client rec_i, attempts to append c to $\mathcal{G}_L$. The set of server processes in c are not a part of any other configuration different from c. In a high level, rec_i first executes a sequence traversal to discover the latest state of $\mathcal{G}_L$. Then it attempts to add the new configuration c, at the end of the discovered sequence by proposing c in the consensus instance of the last configuration in the sequence. The client accepts and appends the decision of the consensus

---

Algorithm 2 Reconfiguration protocol of algorithm ARES.

---

    at each reconfigurer rec_i

2: **State Variables:**
    $cseq[]s.t.cseq[j] \in C \times \{F, P\}$

4: **Initialization:**
    $cseq[0] = \langle c_0, F \rangle$

6: **operation** reconfig(c)
    if $c \neq \perp$ **then**
8:      $cseq \leftarrow$ read-config($cseq$)
      $cseq \leftarrow$ add-config($cseq, c$)
10:    update-config($cseq$)
      $cseq \leftarrow$ finalize-config($cseq$)
12: **end operation**

    procedure add-config(seq, c)
14:    $v \leftarrow |seq|$
      $c' \leftarrow seq[v].cfg$
16:    $d \leftarrow c'.Con.propose(c)$
      $seq[v + 1] \leftarrow \langle d, P \rangle$
18:    put-config($c', \langle d, P \rangle$)

    return seq
20: **end procedure**

    procedure update-config(seq)
22:    $\mu \leftarrow \max(\{j : seq[j].status = F\})$
      $v \leftarrow |seq|$
24:    $M \leftarrow \emptyset$
    for $i = \mu : v$ **do**
26:      $\langle t, v \rangle \leftarrow seq[i].cfg$.get-data()
      $M \leftarrow M \cup \{\langle \tau, v \rangle\}$
28:    $\langle \tau, v \rangle \leftarrow \max_t\{\langle t, v \rangle : \langle t, v \rangle \in M\}$
      $seq[v].cfg$.put-data($\langle \tau, v \rangle$)
30: **end procedure**

    procedure finalize-config(seq)
32:    $v = |seq|$
      $seq[v].status \leftarrow F$
34:    put-config($seq[v - 1].cfg, seq[v]$)
    return seq
36: **end procedure**

Algorithm 3 Server protocol of algorithm ARES.

---

    at each server s_i in configuration c_k

2: **State Variables:**
    $\tau \in \mathbb{N} \times \mathcal{W}$, initially, $\langle 0, \perp \rangle$

4: $v \in V$, initially, $\perp$
    $nextC \in C \times \{P, F\}$, initially $\langle \perp, P \rangle$

6: **Upon receive** (READ-CONFIG) s_i, c_k **from** q
    send $nextC$ to q

8: **end receive**

    Upon receive (WRITE-CONFIG, $cfgT_{in}$) s_i, c_k **from** q
10:    **if** $nextC.cfg = \perp \lor nextC.status = P$ **then**
      $nextC \leftarrow cfgT_{in}$
12:    send ACK to q
    end receive

---

instance (that might be different than c). Then it attempts to transfer the latest value of the read/write object to the latest installed configuration. Once the information is transferred, rec_i finalizes the last configuration in its local sequence and propagates the finalized tuple to a quorum of servers in that configuration. The operation consists of four phases, executed consecutively by rec_i (see Alg. 2):

read-config(seq): The phase read-config(seq) at rec_i, reads the recent global configuration sequence as described in the sequence traversal.

add-config(seq, c): The add-config(seq, c) attempts to append a new configuration c to the end of seq (client's view of $\mathcal{G}_L$). Suppose the last configuration in seq is c' (with status either F or P), then in order to decide the most recent configuration, rec_i invokes $c'.Con.propose(c)$, on the consensus object associated with configuration c'. Let $d \in C$ be the configuration identifier decided by the consensus service. If $d \neq c$, this implies that another (possibly concurrent) reconfiguration operation, invoked by a reconfigurer $rec_j \neq rec_i$, proposed and succeeded d as the configuration to follow c'. In this case, rec_i adopts d as it own propose configuration, by adding $\langle d, P \rangle$ to
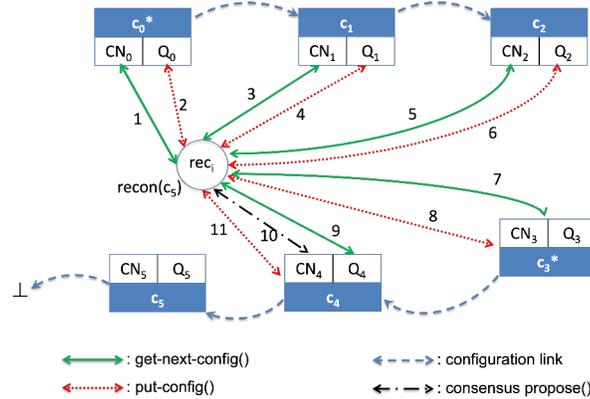
Fig. 1. Illustration of an execution of the reconfiguration steps.

the end of its local $cseq$ (entirely ignoring c), using the operation put-config(c', $\langle d, P \rangle$), and returns the extended configuration seq.

update-config(seq): Let us denote by μ the index of the last configuration in the local sequence $cseq$, at rec_i, such that its corresponding status is F; and ν denote the last index of $cseq$. Next rec_i invokes update-config($cseq$), which gathers the tag-value pair corresponding to the maximum tag in each of the configurations in $cseq[i]$ for $\mu \leq i \leq \nu$, and transfers that pair to the configuration that was added by the add-config action. The get-data and put-data DAPs are used to transfer the value of the object to the new configuration, and they are implemented with respect to the configuration that is accessed. Suppose $\langle t_{max}, v_{max} \rangle$ is the tag value pair corresponding to the highest tag among the responses from all the $\nu - \mu + 1$ configurations. Then, $\langle t_{max}, v_{max} \rangle$ is written to the configuration d via the invocation of $cseq[\nu].cfg$.put-data($\langle \tau_{max}, v_{max} \rangle$).

finalize-config($cseq$): Once the tag-value pair is transferred, in the last phase of the reconfiguration operation, rec_i executes finalize-config($cseq$), to update the status of the last configuration in $cseq$, say $d = cseq[\nu].cfg$, to F. The reconfigurer rec_i informs a quorum of servers in the previous configuration $c = cseq[\nu - 1].cfg$, i.e. in some $Q \in c.Quorums$, about the change of status, by executing the put-config(c, $\langle d, F \rangle$) action.

Server Protocol. Each server responds to requests from clients (Alg. 3). A server waits for two types of messages: READ-CONFIG and WRITE-CONFIG. When a READ-CONFIG message is received for a particular configuration c_k, then the server returns $nextC$ variables of the servers in $c_k.Servers$. A WRITE-CONFIG message attempts to update the $nextC$ variable of the server with a particular tuple $cfgT_{in}$. A server changes the value of its local $nextC.cfg$ in two cases: (i) $nextC.cfg = \bot$, or (ii) $nextC.status = P$.

Fig. 1 illustrates an example execution of a reconfiguration operation recon(c_5). In this example, the reconfigurer rec_i goes through a number of configuration queries (get-next-config) before it reaches configuration c_4 in which a quorum of servers replies with $nextC.cfg = \bot$. There it proposes c_5 to the consensus object of c_4 ($c_4.Con.propose(c_5)$ on arrow 10), and once c_5 is decided, recon(c_5) completes after executing finalize-config(c_5).

4.2 Implementation of Read and Write operations.

The read and write operations in Ares are expressed in terms of the DAP primitives (see Section 3). This provides the flexibility to Ares to use different implementation of DAP primitives in different configurations, without changing the basic structure of Ares. At a high-level, a write (or read) operation is executed where the client: (*i*) obtains the *latest configuration sequence* by using the read-config action of the reconfiguration service, (*ii*) queries the configurations, in $cseq$, starting from the last finalized configuration to the end of the discovered configuration sequence, for the latest $\langle tag, value \rangle$ pair with a help of get-tag (or get-data) operation as specified

---

Algorithm 4 Write and Read protocols at the clients for ARES.

<div style="display:flex">
<div>

    Write Operation:
2:    at each writer w_i
    State Variables:
4:    $cseq[] s.t. cseq[j] \in C \times \{F, P\}$
    Initialization:
6:    $cseq[0] = \langle c_0, F \rangle$

    operation write(val), $val \in V$
8:      $cseq \leftarrow$ read-config($cseq$)
      $\mu \leftarrow \max(\{i : cseq[i].status = F\})$
10:    $v \leftarrow |cseq|$
      for $i = \mu : v$ **do**
12:        $\tau_{max} \leftarrow \max(cseq[i].cfg.\text{get-tag}(), \tau_{max})$
        $\langle \tau, v \rangle \leftarrow \langle \langle \tau_{max}.ts + 1, \omega_i \rangle, val \rangle$
14:    $done \leftarrow false$
      while not $done$ **do**
16:      $cseq[v].cfg.\text{put-data}(\langle \tau, v \rangle)$
        $cseq \leftarrow$ read-config($cseq$)
18:      **if** $|cseq| = v$ **then**
          $done \leftarrow true$
20:      **else**
          $v \leftarrow |cseq|$
22:    **end while**
    end operation

</div>
<div>

24: **Read Operation:**
    at each reader r_i
26:    **State Variables:**
    $cseq[] s.t. cseq[j] \in C \times \{F, P\}$
28:    **Initialization:**
    $cseq[0] = \langle c_0, F \rangle$

30:    **operation** read( )
      $cseq \leftarrow$ read-config($cseq$)
32:    $\mu \leftarrow \max(\{j : cseq[j].status = F\})$
      $v \leftarrow |cseq|$
34:    **for** $i = \mu : v$ **do**
      $\langle \tau, v \rangle \leftarrow \max(cseq[i].cfg.\text{get-data}(), \langle \tau, v \rangle)$
36:    $done \leftarrow false$
      while not $done$ **do**
38:      $cseq[v].cfg.\text{put-data}(\langle \tau, v \rangle)$
        $cseq \leftarrow$ read-config($cseq$)
40:    **if** $|cseq| = v$ **then**
          $done \leftarrow true$
42:    **else**
          $v \leftarrow |cseq|$
44:    **end while**
    return v
46:    **end operation**

</div>
</div>

---

for each configuration, and (*iii*) repeatedly propagates a new $\langle tag', value' \rangle$ pair (the largest $\langle tag, value \rangle$ pair) with put-data in the last configuration of its local sequence, until no additional configuration is observed. In more detail, the algorithm of a read or write operation π is as follows (see Alg. 4):

A write (or read) operation is invoked at a client p when line Alg. 4:8 (resp. line Alg. 4:31) is executed. At first, p issues a read-config action to obtain the latest introduced configuration in $\mathcal{G}_L$, in both operations.

If π is a write p detects the last finalized entry in $cseq$, say μ, and performs a $cseq[j].conf.\text{get-tag}()$ action, for $\mu \le j \le |cseq|$ (line Alg. 4:9). Then p discovers the *maximum tag* among all the returned tags (τ_{max}), and it increments the maximum tag discovered (by incrementing the integer part of τ_{max}), generating a new tag, say τ_{new}. It assigns $\langle \tau, v \rangle$ to $\langle \tau_{new}, val \rangle$, where val is the value he wants to write (Line Alg. 4:13).

if π is a read, p detects the last finalized entry in $cseq$, say μ, and performs a $cseq[j].conf.\text{get-data}()$ action, for $\mu \le j \le |cseq|$ (line Alg. 4:32). Then p discovers the *maximum tag-value* pair ($\langle \tau_{max}, v_{max} \rangle$) among the replies, and assigns $\langle \tau, v \rangle$ to $\langle \tau_{max}, v_{max} \rangle$.

Once specifying the $\langle \tau, v \rangle$ to be propagated, both reads and writes execute the $cseq[v].cfg.\text{put-data}(\langle \tau, v \rangle)$ action, where $v = |cseq|$, followed by executing read-config action, to examine whether new configurations were introduced in $\mathcal{G}_L$. This is an essential step that ensures that any new value of the object is propagated in any recently introduced configuration. Omission to do so may lead an operation that reads from a newly established configuration to obtain an outdated value for the shared object, violating this way atomic consistency. Each operation repeats these steps until no new configuration is discovered (lines Alg. 4:15–21, or lines Alg. 4:37−43). Let $cseq'$ be the sequence returned by the read-config action. If $|cseq'| = |cseq|$ then no new configuration is

introduced, and the read/write operation terminates; otherwise, p sets $cseq$ to $cseq'$ and repeats the two actions. Note, in an execution of Ares, two consecutive read-config operations that return $cseq'$ and $cseq''$ respectively must hold that $cseq'$ is a prefix of $cseq''$, and hence $|cseq'| = |cseq''|$ only if $cseq' = cseq''$. Finally, if π is a read operation the value with the highest tag discovered is returned to the client.

Discussion Ares shares similarities with previous algorithms like RAMBO [28] and the framework in [48]. The reconfiguration technique used in Ares ensures the prefix property on the configuration sequence (resembling a blockchain data structure [42]) while the array structure in RAMBO allowed nodes to maintain an incomplete reconfiguration history. On the other hand, the DAP usage, exploits a different viewpoint compared to [48], allowing implementations of atomic read/write registers without relying on strong objects, like ranked registers [18]. Note that Ares is designed to capture a wide class of algorithms with different redundancy strategies. So while not directly implementing an EC-based atomic memory, it provides the "vehicle" without which dynamic EC-based implementations would not have been possible. Lastly, even though Ares is designed to support crash failures, as noted by [41], reconfiguration is more general and allows an algorithm to handle benign recoveries. That is, a recovered node that loses its state can be introduced as a new member of a new configuration. Stateful recoveries on the other hand are indistinguishable from long delays, thus can be handled effectively by an algorithm designed for the asynchronous model like Ares.

5 IMPLEMENTATION OF THE DAPS

In this section, we present an implementation of the DAPs, that satisfies the properties in Property 1, for a configuration c, with n servers using a $[n, k]$ MDS coding scheme for storage. Notice that the total number of servers in the system can be larger than n, however we can pick a subset of n servers to use for this particular key and instance of the algorithm. We store each coded element c_i, corresponding to an object at server s_i, where $i = 1, \cdots, n$. The implementations of DAP primitives used in Ares are shown in Alg. 5, and the servers' responses in Alg. 6.

Algorithm 5 DAP implementation for Ares.

---

at each process $p_i \in \mathcal{I}$

2: **procedure** c.get-tag()
    send (QUERY-TAG) to each $s \in c.Servers$
4:   **until** p_i receives $\langle t_s \rangle$ from $\left\lceil \frac{n+k}{2} \right\rceil$ servers in $c.Servers$
    $t_{max} \leftarrow \max(\{t_s : \text{ received } t_s \text{ from } s\})$
6:   **return** t_{max}
  end procedure

8: **procedure** c.get-data()
    send (QUERY-LIST) to each $s \in c.Servers$
10:   **until** p_i receives $List_s$ from each server $s \in \mathcal{S}_g$ s.t.
    $|\mathcal{S}_g| = \left\lceil \frac{n+k}{2} \right\rceil$ and $\mathcal{S}_g \subset c.Servers$
    $Tags_*^{\geq k} = $ set of tags that appears in k lists

12:   $Tags_{dec}^{\geq k} = $ set of tags that appears in k lists with values
    $t_{max}^* \leftarrow \max Tags_*^{\geq k}$
14:   $t_{max}^{dec} \leftarrow \max Tags_{dec}^{\geq k}$
    if $t_{max}^{dec} = t_{max}^*$ **then**
16:       $v \leftarrow $ decode value for t_{max}^{dec}
      return $\langle t_{max}^{dec}, v \rangle$
18: **end procedure**

  procedure c.put-data($\langle \tau, v \rangle$))
20:   $code\text{-}elems = [(\tau, e_1), \ldots, (\tau, e_n)], e_i = \Phi_i(v)$
    send (PUT-DATA, $\langle \tau, e_i \rangle$) to each $s_i \in c.Servers$
22:   **until** p_i receives ACK from $\left\lceil \frac{n+k}{2} \right\rceil$ servers in $c.Servers$
  end procedure

Each server s_i stores one state variable, $List$, which is a set of up to $(\delta + 1)$ (tag, coded-element) pairs. Initially the set at s_i contains a single element, $List = \{(t_0, \Phi_i(v_0)\}$. Below we describe the implementation of the DAPs.

Algorithm 6 The response protocols at any server $s_i \in \mathcal{S}$ in ARES for client requests.

---

at each server $s_i \in \mathcal{S}$ in configuration c_k

2: **State Variables:**
$List \subseteq \mathcal{T} \times C_s$, initially $\{(t_0, \Phi_i(v_0))\}$

   Upon receive (QUERY-TAG) $_{s_i, c_k}$ **from** q
4:    $\tau_{max} = \max_{(t,c) \in List} t$
   Send τ_{max} to q
6: **end receive**

   Upon receive (QUERY-LIST) $_{s_i, c_k}$ **from** q
8:    Send $List$ to q

   end receive
10:
   Upon receive (PUT-DATA, $\langle \tau, e_i \rangle$) $_{s_i, c_k}$ **from** q
12:    $List \leftarrow List \cup \{\langle \tau, e_i \rangle\}$
   if $|List| > \delta + 1$ **then**
14:      $\tau_{min} \leftarrow \min\{t : \langle t, * \rangle \in List\}$
   /* remove the coded value and retain the tag */
     $List \leftarrow List \setminus \{\langle \tau, e \rangle : \tau = \tau_{min} \wedge \langle \tau, e \rangle \in List\}$
16:      $List \leftarrow List \cup \{(\tau_{min}, \bot)\}$
   Send ACK to q
18: **end receive**

---

c.get-tag(): A client, during the execution of a c.get-tag() primitive, queries all the servers in $c.Servers$ for the highest tags in their $Lists$, and awaits responses from $\lceil \frac{n+k}{2} \rceil$ servers. A server upon receiving the GET-TAG request, responds to the client with the highest tag, as $\tau_{max} \equiv \max_{(t,c) \in List} t$. Once the client receives the tags from $\lceil \frac{n+k}{2} \rceil$ servers, it selects the highest tag t and returns it .

c.put-data($\langle t_w, v \rangle$): During the execution of the primitive c.put-data($\langle t_w, v \rangle$), a client sends the pair $(t_w, \Phi_i(v))$ to each server $s_i \in c.Servers$. When a server s_i receives a message (PUT-DATA, t_w, c_i) , it adds the pair in its local $List$, trims the pairs with the smallest tags exceeding the length $(\delta + 1)$ of the $List$, and replies with an ack to the client. In particular, s_i replaces the coded-elements of the older tags with $\bot$, and maintains only the coded-elements associated with the $(\delta + 1)$ highest tags in the $List$ (see Line Alg. 6:16). The client completes the primitive operation after getting acks from $\lceil \frac{n+k}{2} \rceil$ servers.

c.get-data(): A client, during the execution of a c.get-data() primitive, queries all the servers in $c.Servers$ for their local variable $List$, and awaits responses from $\lceil \frac{n+k}{2} \rceil$ servers. Once the client receives $Lists$ from $\lceil \frac{n+k}{2} \rceil$ servers, it selects the highest tag t, such that: (i) its corresponding value v is decodable from the coded elements in the lists; and (ii) t is the highest tag seen from the responses of at least k $Lists$ (see lines Alg. 5:11-14) and returns the pair (t, v). Note that in the case where anyone of the above conditions is not satisfied the corresponding read operation does not complete.

5.1 Correctness of the DAPs

To proof the correctness of the proposed DAPs, we need to show that they are both *safe*, i.e. ensure the necessary Property 1, and *live*, i.e. they allow each operation to terminate. We first proceed to prove that for any given execution ξ containing operations of the proposed implementation, then examining any pair of operations in ξ satisfy the DAP consistency properties (i.e. Property 1). That is, the tag returned by a get-tag() operation is larger than the value written by any preceding put-data() operation, and the value returned by a get-data() operation is either written by a put-data() operation or is the initial value of the object. Next, assuming that there cannot be more that δ put-data() operations concurrent with a single get-data() operation, we show that each operation in our implementation terminates. Otherwise a get-data() operation is at risk of not being able to discover a decodable value and thus fail to terminate and return a value.

Safety (Property 1). In this section we are concerned with only one configuration c, consisting of a set of servers $c.Servers$. We assume that at most $f \le \frac{n-k}{2}$ servers from $c.Servers$ may crash. Lemma 2 states that the DAP

implementation satisfies the consistency properties Property 1 which will be used to imply the atomicity of the Ares algorithm.

THEOREM 2 (SAFETY). *Let* Π *a set of complete DAP operations of Algorithm 5 in a configuration* $c \in C$, c.get-tag, c.get-data *and* c.put-data, *of an execution* ξ. *Then, every pair of operations* $\phi, \pi \in \Pi$ *satisfy Property 1.*

PROOF. As mentioned above we are concerned with only configuration c, and therefore, in our proofs it suffices to examine only one configuration. Let ξ be some execution of Ares, then we consider two cases for π for proving property $C1$: π is a get-tag, or π is a get-data primitive.

Case (*a*): ϕ is c.put-data($\langle \tau_\phi, v_\phi \rangle$) and π is a c.get-tag() returns $\tau_\pi \in \mathcal{T}$. Let c_ϕ and c_π denote the clients that invokes ϕ and π in ξ. Let $S_\phi \subset S$ denote the set of $\lceil \frac{n+k}{2} \rceil$ servers that responds to c_ϕ, during ϕ. Denote by S_π the set of $\lceil \frac{n+k}{2} \rceil$ servers that responds to c_π, during π. Let T_1 be a point in execution ξ after the completion of ϕ and before the invocation of π. Because π is invoked after T_1, therefore, at T_1 each of the servers in S_ϕ contains t_ϕ in its *List* variable. Note that, once a tag is added to *List*, it is never removed. Therefore, during π, any server in $S_\phi \cap S_\pi$ responds with *List* containing t_ϕ to c_π. Note that since $|S_\phi| = |S_\pi| = \lceil \frac{n+k}{2} \rceil$ implies $|S_\phi \cap S_\pi| \geq k$, and hence t_{max}^{dec} at c_π, during π is at least as large as t_ϕ, i.e., $t_\pi \geq t_\phi$. Therefore, it suffices to prove our claim with respect to the tags and the decodability of its corresponding value.

Case (*b*): ϕ is c.put-data($\langle \tau_\phi, v_\phi \rangle$) and π is a c.get-data() returns $\langle \tau_\pi, v_\pi \rangle \in \mathcal{T} \times \mathcal{V}$. As above, let c_ϕ and c_π be the clients that invokes ϕ and π. Let S_ϕ and S_π be the set of servers that responds to c_ϕ and c_π, respectively. Arguing as above, $|S_\phi \cap S_\pi| \geq k$ and every server in $S_\phi \cap S_\pi$ sends t_ϕ in response to c_ϕ, during π, in their *List*'s and hence $t_\phi \in Tags_*^{\geq k}$. Now, because π completes in ξ, hence we have $t_{max}^* = t_{max}^{dec}$. Note that $\max Tags_*^{\geq k} \geq \max Tags_{dec}^{\geq k}$ so $t_\pi \geq \max Tags_{dec}^{\geq k} = \max Tags_*^{\geq k} \geq t_\phi$. Note that each tag is always associated with its corresponding value v_π, or the corresponding coded elements $\Phi_s(v_\pi)$ for $s \in S$.

Next, we prove the $C2$ property of DAP for the Ares algorithm. Note that the initial values of the *List* variable in each servers s in S is $\{(t_0, \Phi_s(v_\pi))\}$. Moreover, from an inspection of the steps of the algorithm, new tags in the *List* variable of any servers of any servers is introduced via put-data operation. Since t_π is returned by a get-tag or get-data operation then it must be that either $t_\pi = t_0$ or $t_\pi > t_0$. In the case where $t_\pi = t_0$ then we have nothing to prove. If $t_\pi > t_0$ then there must be a put-data(t_π, v_π) operation ϕ. To show that for every π it cannot be that ϕ completes before π, we adopt by a contradiction. Suppose for every π, ϕ completes before π begins, then clearly t_π cannot be returned ϕ, a contradiction. $\square$

Liveness. To reason about the liveness of the proposed DAPs, we define a concurrency parameter δ which captures all the put-data operations that overlap with the get-data, until the time the client has all data needed to attempt decoding a value. However, we ignore those put-data operations that might have started in the past, and never completed yet, if their tags are less than that of any put-data that completed before the get-data started. This allows us to ignore put-data operations due to failed clients, while counting concurrency, as long as the failed put-data operations are followed by a successful put-data that completed before the get-data started. In order to define the amount of concurrency in our specific implementation of the DAPs presented in this section the following definition captures the put-data operations that overlap with the get-data, until the client has all data required to decode the value.

DEFINITION 3 (VALID get-data OPERATIONS). *A* get-data *operation* π *from a process* p *is valid if* p *does not crash until the reception of* $\lceil \frac{n+k}{2} \rceil$ *responses during the get-data phase.*

DEFINITION 4 (put-data CONCURRENT WITH A VALID get-data). *Consider a valid* get-data *operation* π *from a process* p. *Let* T_1 *denote the point of initiation of* π. *For* π, *let* T_2 *denote the earliest point of time during the execution when* p *receives all the* $\lceil \frac{n+k}{2} \rceil$ *responses. Consider the set* $\Sigma = \{\phi : \phi$ *is any* put-data *operation that completes before*

π is initiated}, and let $\phi^* = \arg\max_{\phi \in \Sigma} tag(\phi)$. Next, consider the set $\Lambda = \{\lambda : \lambda$ is any put-data operation that starts before T_2 such that $tag(\lambda) > tag(\phi^*)\}$. We define the number of put-data concurrent with the valid get-data π to be the cardinality of the set Λ.

Termination (and hence liveness) of the DAPs is guaranteed in an execution ξ, provided that a process no more than f servers in $c.Servers$ crash, and no more that δ put-data may be concurrent at any point in ξ. If the failure model is satisfied, then any operation invoked by a non-faulty client will collect the necessary replies independently of the progress of any other client process in the system. Preserving δ on the other hand, ensures that any operation will be able to decode a written value. These are captured in the following theorem:

THEOREM 5 (LIVENESS). *Let ξ be well-formed and fair execution of DAPs, with an $[n, k]$ MDS code, where n is the number of servers out of which no more than $\frac{n-k}{2}$ may crash, and δ be the maximum number of put-data operations concurrent with any valid get-data operation. Then any get-data and put-data operation π invoked by a process p terminates in ξ if p does not crash between the invocation and response steps of π.*

PROOF. Note that in the read and write operation the get-tag and put-data operations initiated by any non-faulty client always complete. Therefore, the liveness property with respect to any write operation is clear because it uses only get-tag and put-data operations of the DAP. So, we focus on proving the liveness property of any read operation π, specifically, the get-data operation completes. Let ξ be and execution of ARES and let c_ω and c_π be the clients that invokes the write operation ω and read operation π, respectively.

Let S_ω be the set of $\lceil \frac{n+k}{2} \rceil$ servers that responds to c_ω, in the put-data operations, in ω. Let S_π be the set of $\lceil \frac{n+k}{2} \rceil$ servers that responds to c_π during the get-data step of π. Note that in ξ at the point execution T_1, just before the execution of π, none of the write operations in Λ is complete. Observe that, by algorithm design, the coded-elements corresponding to t_ω are garbage-collected from the *List* variable of a server only if more than δ higher tags are introduced by subsequent writes into the server. Since the number of concurrent writes $|\Lambda|$, s.t. $\delta > |\Lambda|$ the corresponding value of tag t_ω is not garbage collected in ξ, at least until execution point T_2 in any of the servers in S_ω.

Therefore, during the execution fragment between the execution points T_1 and T_2 of the execution ξ, the tag and coded-element pair is present in the *List* variable of every in S_ω that is active. As a result, the tag and coded-element pairs, $(t_\omega, \Phi_s(v_\omega))$ exists in the *List* received from any $s \in S_\omega \cap S_\pi$ during operation π. Note that since $|S_\omega| = |S_\pi| = \lceil \frac{n+k}{2} \rceil$ hence $|S_\omega \cap S_\pi| \geq k$ and hence $t_\omega \in Tags_{dec}^{\geq k}$, the set of decode-able tag, i.e., the value v_ω can be decoded by c_π in π, which demonstrates that $Tags_{dec}^{\geq k} \neq \emptyset$. Next we want to argue that $t_{max}^* = t_{max}^{dec}$ via a contradiction: we assume $\max Tags_*^{\geq k} > \max Tags_{dec}^{\geq k}$. Now, consider any tag t, which exists due to our assumption, such that, $t \in Tags_*^{\geq k}$, $t \notin Tags_{dec}^{\geq k}$ and $t > t_{max}^{dec}$. Let $S_\pi^k \subset S$ be any subset of k servers that responds with t_{max}^* in their *List* variables to c_π. Note that since $k > n/3$ hence $|S_\omega \cap S_\pi^k| \geq \lceil \frac{n+k}{2} \rceil + \lceil \frac{n+1}{3} \rceil \geq 1$, i.e., $S_\omega \cap S_\pi^k \neq \emptyset$. Then t must be in some servers in S_ω at T_2 and since $t > t_{max}^{dec} \geq t_\omega$. Now since $|\Lambda| < \delta$ hence $(t, \perp)$ cannot be in any server at T_2 because there are not enough concurrent write operations (i.e., writes in Λ) to garbage-collect the coded-elements corresponding to tag t, which also holds for tag t_{max}^*. In that case, t must be in $Tag_{dec}^{\geq k}$, a contradiction. □

6 CORRECTNESS OF ARES

In this section, we prove that ARES correctly implements an atomic, read/write, shared storage service. The correctness of ARES highly depends on the way the configuration sequence is constructed at each client process. Also, atomicity is ensured if the DAP implementation in each configuration c_i satisfies Property 1.

As a roadmap, we begin by showing that some critical properties are preserved by the reconfiguration service proposed in subsection 6.1. In particular, we show that the configuration sequence maintained in two processes is either the same or one is the prefix of the other. This in turn helps us to proof the correctness of ARES in

subsection 6.2 by showing that all the properties of atomicity (see Section 2) are satisfied, given the properties on the configuration sequence hold, and that the DAPs used in each configuration satisfy Property 1.

We proceed by first introducing some definitions and notation, that we use in the proofs that follow.

Notations and definitions. For a server s, we use the notation $s.var|_\sigma$ to refer to the value of the state variable var, in s, at a state σ of an execution ξ. If server s crashes at a state σ_f in an execution ξ then $s.var|_\sigma \triangleq s.var|_{\sigma_f}$ for any state variable var and for any state σ that appears after σ_f in ξ (i.e. the value of the variable remains unchanged).

We define as the tag of a configuration c the smallest tag among the maximum tags found in each quorum of c. This is essentially the smallest tag that an operation may witness when receiving replies from a single quorum in c. More formally:

DEFINITION 6 (TAG OF A CONFIGURATION). *Let $c \in C$ be a configuration, σ be a state in some execution ξ then we define the tag of c at state σ as $tag(c)|_\sigma \triangleq \min_{Q \in c.Quorums} \max_{s \in Q} (s.tag|_\sigma)$. We drop the suffix $|_\sigma$, and simply denote as $tag(c)$, when the state is clear from the context.*

Next we provide the notation to express the configuration sequence witnessed by a process p in a state σ (as $p.cseq|_\sigma$), the last finalized configuration in that sequence (as $\mu(\mathbf{c}_\sigma^p)$), and the length of that sequence (as $\nu(\mathbf{c}_\sigma^p)$). More formally:

DEFINITION 7. *Let σ be any point in an execution of ARES and suppose we use the notation $\mathbf{c}_\sigma^p$ for $p.cseq|_\sigma$, i.e., the cseq variable at process p at the state σ. Then we define as $\mu(\mathbf{c}_\sigma^p) \triangleq \max\{i : \mathbf{c}_\sigma^p[i].status = F\}$ and $\nu(\mathbf{c}_\sigma^p) \triangleq |\mathbf{c}_\sigma^p|$, where $|\mathbf{c}_\sigma^p|$ is the length of the configuration vector $\mathbf{c}_\sigma^p$.*

Last, we define the prefix operation on two configuration sequences.

DEFINITION 8 (PREFIX ORDER). *Let $\mathbf{x}$ and $\mathbf{y}$ be any two configuration sequences. We say that $\mathbf{x}$ is a prefix of $\mathbf{y}$, denoted by $\mathbf{x} \preceq_p \mathbf{y}$, if $\mathbf{x}[j].cfg = \mathbf{y}[j].cfg$, for all j such that $\mathbf{x}[j] \neq \perp$.*

Table 3 summarizes the new notation for ease of reference.

$\mathbf{c}_\sigma^p$	the value of the configuration sequence variable *cseq* at process p in state σ, i.e. a shorthand of $p.cseq|_\sigma$
$\mathbf{c}_\sigma^p[i]$	the i^{th} element in the configuration sequence $\mathbf{c}_\sigma^p$
$\mu(\mathbf{c}_\sigma^p)$	last finalized configuration in $\mathbf{c}_\sigma^p$
$\nu(\mathbf{c}_\sigma^p)$	the length of $\mathbf{c}_\sigma^p$, i.e. $|\mathbf{c}_\sigma^p|$

Table 3. Additional notation used in this section.

6.1 Reconfiguration Protocol Properties

In this section we analyze the properties that we can achieve through our reconfiguration algorithm. In high-level, we do show that the following properties are preserved:

   i **configuration uniqueness:** the configuration sequences in any two processes have identical configuration at any place i,

  ii **sequence prefix:** the configuration sequence witnessed by an operation is a prefix of the sequence witnessed by any succeeding operation, and

iii **sequence progress:** if the configuration with index i is finalized during an operation, then a configuration j, for $j \geq i$, will be finalized for a succeeding operation.

The first lemma shows that any two configuration sequences have the same configuration identifiers in the same indexes.

LEMMA 9. *For any reconfigurer r that invokes an* reconfig(c) *action in an execution ξ of the algorithm, If r chooses to install c in index k of its local $r.cseq$ vector, then r invokes the $Cons[k-1].propose(c)$ instance over configuration $r.cseq[k-1].cfg$.*

PROOF. It follows directly from the algorithm. □

LEMMA 10. *If a server s sets $s.nextC$ to $\langle c, F \rangle$ at some state σ in an execution ξ of the algorithm, then $s.nextC = \langle c, F \rangle$ for any state σ' that appears after σ in ξ.*

PROOF. Notice that a server s updates its $s.nextC$ variable for some specific configuration c_k in a state σ only when it receives a WRITE-CONF message. This is either the first WRITE-CONF message received at s for c_k (and thus $nextC = \bot$), or $s.nextC = \langle *, P \rangle$ and the message received contains a tuple $\langle c, F \rangle$. Once the tuple becomes equal to $\langle c, F \rangle$ then s does not satisfy the update condition for c_k, and hence in any state σ' after σ it does not change $\langle c, F \rangle$. □

LEMMA 11 (CONFIGURATION UNIQUENESS). *For any processes $p, q \in \mathcal{I}$ and any states σ_1, σ_2 in an execution ξ, it must hold that $\mathbf{c}^p_{\sigma_1}[i].cfg = \mathbf{c}^q_{\sigma_2}[i].cfg, \forall i$ s.t. $\mathbf{c}^p_{\sigma_1}[i].cfg, \mathbf{c}^q_{\sigma_2}[i].cfg \neq \bot$.*

PROOF. The lemma holds trivially for $\mathbf{c}^p_{\sigma_1}[0].cfg = \mathbf{c}^q_{\sigma_2}[0].cfg = c_0$. So in the rest of the proof we focus in the case where $i > 0$. Let us assume w.l.o.g. that σ_1 appears before σ_2 in ξ.

According to our algorithm a process p sets $p.cseq[i].cfg$ to a configuration identifier c in two cases: (i) either it received c as the result of the consensus instance in configuration $p.cseq[i-1].cfg$, or (ii) p receives $s.nextC.cfg = c$ from a server $s \in p.cseq[i-1].cfg.Servers$. Note here that (i) is possible only when p is a reconfigurer and attempts to install a new configuration. On the other hand (ii) may be executed by any process in any operation that invokes the read-config action. We are going to proof this lemma by induction on the configuration index.

Base case: The base case of the lemma is when $i = 1$. Let us first assume that p and q receive c_p and c_q, as the result of the consensus instance at $p.cseq[0].cfg$ and $q.cseq[0].cfg$ respectively. By Lemma 9, since both processes want to install a configuration in $i = 1$, then they have to run $Cons[0]$ instance over the configuration stored in their local $cseq[0].cfg$ variable. Since $p.cseq[0].cfg = q.cseq[0].cfg = c_0$ then both $Cons[0]$ instances run over the same configuration c_0 and thus by the agreement property the have to decide the same value, say c_1. Hence $c_p = c_q = c_1$ and $p.cseq[1].cfg = q.cseq[1].cfg = c_1$.

Let us examine the case now where p or q assign a configuration c they received from some server $s \in c_0.Servers$. According to the algorithm only the configuration that has been decided by the consensus instance on c_0 is propagated to the servers in $c_0.Servers$. If c_1 is the decided configuration, then $\forall s \in c_0.Servers$ such that $s.nextC(c_0) \neq \bot$, it holds that $s.nextC(C_0) = \langle c_1, * \rangle$. So if p or q set $p.cseq[1].cfg$ or $q.cseq[1].cfg$ to some received configuration, then $p.cseq[1].cfg = q.cseq[1].cfg = c_1$ in this case as well.

Hypothesis: We assume that $\mathbf{c}^p_{\sigma_1}[k] = \mathbf{c}^q_{\sigma_2}[k]$ for some $k, k \geq 1$.

Induction Step: We need to show that the lemma holds for $i = k + 1$. If both processes retrieve $p.cseq[k+1].cfg$ and $q.cseq[k+1].cfg$ through consensus, then both p and q run consensus over the previous configuration. Since according to our hypothesis $\mathbf{c}^p_{\sigma_1}[k] = \mathbf{c}^q_{\sigma_2}[k]$ then both process will receive the same decided value, say c_{k+1}, and hence $p.cseq[k+1].cfg = q.cseq[k+1].cfg = c_{k+1}$. Similar to the base case, a server in $c_k.Servers$ only receives the configuration c_{k+1} decided by the consensus instance run over c_k. So processes p and q can only receive c_{k+1}

from some server in $c_k.Servers$ so they can only assign $p.cseq[k+1].cfg = q.cseq[k+1].cfg = c_{k+1}$ at Line A2:8. That completes the proof. □

Lemma 11 showed that any two operations store the same configuration in any cell k of their $cseq$ variable. It is not known however if the two processes discover the same number of configuration ids. In the following lemmas we will show that if a process learns about a configuration in a cell k then it also learns about all configuration ids for every index i, such that $0 \le i \le k-1$.

LEMMA 12. *In any execution ξ of the algorithm , If for any process $p \in \mathcal{I}$, $\mathbf{c}^p_\sigma[i] \ne \bot$ in some state σ in ξ, then $\mathbf{c}^p_{\sigma'}[i] \ne \bot$ in any state σ' that appears after σ in ξ.*

PROOF. A value is assigned to $\mathbf{c}^p_*[i]$ either after the invocation of a consensus instance, or while executing the read-config action. Since any configuration proposed for installation cannot be $\bot$ (A2:7), and since there is at least one configuration proposed in the consensus instance (the one from p), then by the validity of the consensus service the decision will be a configuration $c \ne \bot$. Thus, in this case $\mathbf{c}^p_*[i]$ cannot be $\bot$. Also in the read-config procedure, $\mathbf{c}^p_*[i]$ is assigned to a value different than $\bot$ according to Line A2:8. Hence, if $\mathbf{c}^p_\sigma[i] \ne \bot$ at state σ then it cannot become $\bot$ in any state σ' after σ in execution ξ. □

LEMMA 13. *Let σ_1 be some state in an execution ξ of the algorithm. Then for any process p, if $k = max\{i : \mathbf{c}^p_{\sigma_1}[i] \ne \bot\}$, then $\mathbf{c}^p_{\sigma_1}[j] \ne \bot$, for $0 \le j < k$.*

PROOF. Let us assume to derive contradiction that there exists $j < k$ such that $\mathbf{c}^p_{\sigma_1}[j] = \bot$ and $\mathbf{c}^p_{\sigma_1}[j+1] \ne \bot$. Consider first that $j = k-1$ and that σ_1 is the state immediately after the assignment of a value to $\mathbf{c}^p_{\sigma_1}[k]$, say c_k. Since $\mathbf{c}^p_{\sigma_1}[k] \ne \bot$, then p assigned c_k to $\mathbf{c}^p_{\sigma_1}[k]$ in one of the following cases: (i) c_k was the result of the consensus instance, or (ii) p received c_k from a server during a read-config action. The first case is trivially impossible as according to Lemma 9 p decides for k when it runs consensus over configuration $\mathbf{c}^p_{\sigma_1}[k-1].cfg$. Since this is equal to $\bot$, then we cannot run consensus over a non-existent set of processes. In the second case, p assigns $\mathbf{c}^p_{\sigma_1}[k] = c_k$ in Line A1:8. The value c_k was however obtained when p invoked get-next-config on configuration $\mathbf{c}^p_{\sigma_1}[k-1].cfg$. In that action, p sends READ-CONFIG messages to the servers in $\mathbf{c}^p_{\sigma_1}[k-1].cfg.Servers$ and waits until a quorum of servers replies. Since we assigned $\mathbf{c}^p_{\sigma_1}[k] = c_k$ it means that get-next-config terminated at some state σ' before σ_1 in ξ, and thus: (a) a quorum of servers in $\mathbf{c}^p_{\sigma'}[k-1].cfg.Servers$ replied, and (b) there exists a server s among those that replied with c_k. According to our assumption however, $\mathbf{c}^p_{\sigma_1}[k-1] = \bot$ at σ_1. So if state σ' is before σ_1 in ξ, then by Lemma 12, it follows that $\mathbf{c}^p_{\sigma'}[k-1] = \bot$. This however implies that p communicated with an empty configuration, and thus no server replied to p. This however contradicts the assumption that a server replied with c_k to p.

Since any process traverses the configuration sequence starting from the initial configuration c_0, then with a simple induction and similar reasoning we can show that $\mathbf{c}^p_{\sigma_1}[j] \ne \bot$, for $0 \le j \le k-1$. □

We can now move to an important lemma that shows that any read-config action returns an extension of the configuration sequence returned by any previous read-config action. First, we show that the last finalized configuration observed by any read-config action is at least as recent as the finalized configuration observed by any subsequent read-config action.

LEMMA 14. *If at a state σ of an execution ξ of the algorithm $\mu(\mathbf{c}^p_\sigma) = k$ for some process p, then for any element $0 \le j < k$, $\exists Q \in \mathbf{c}^p_\sigma[j].cfg.Quorums$ such that $\forall s \in Q, s.nextC(\mathbf{c}^p_\sigma[j].cfg) = \mathbf{c}^p_\sigma[j+1]$.*

PROOF. This lemma follows directly from the algorithm. Notice that whenever a process assigns a value to an element of its local configuration (Lines A1:8 and A2:17), it then propagates this value to a quorum of the previous configuration (Lines A1:9 and A2:18). So if a process p assigned c_j to an element $\mathbf{c}^p_{\sigma'}[j]$ in some state σ'

in ξ, then p may assign a value to the $j+1$ element of $\mathbf{c}_{\sigma''}^p[j+1]$ only after put-config($\mathbf{c}_{\sigma'}^p[j-1].cfg, \mathbf{c}_{\sigma'}^p[j]$) occurs. During put-config action, p propagates $\mathbf{c}_{\sigma'}^p[j]$ in a quorum $Q \in \mathbf{c}_{\sigma'}^p[j-1].cfg.Quorums$. Hence, if $\mathbf{c}_\sigma^p[k] \neq \bot$, then p propagated each $\mathbf{c}_{\sigma'}^p[j]$, for $0 < j \leq k$ to a quorum of servers $Q \in \mathbf{c}_{\sigma'}^p[j-1].cfg.Quorums$. And this completes the proof. □

LEMMA 15 (SEQUENCE PREFIX). *Let π_1 and π_2 two completed* read-config *actions invoked by processes $p_1, p_2 \in \mathcal{I}$ respectively, such that $\pi_1 \to \pi_2$ in an execution ξ. Let σ_1 be the state after the response step of π_1 and σ_2 the state after the response step of π_2. Then $\mathbf{c}_{\sigma_1}^{p_1} \preceq_p \mathbf{c}_{\sigma_2}^{p_2}$.*

PROOF. Let $v_1 = v(\mathbf{c}_{\sigma_1}^{p_1})$ and $v_2 = v(\mathbf{c}_{\sigma_2}^{p_2})$. By Lemma 11 for any i such that $\mathbf{c}_{\sigma_1}^{p_1}[i] \neq \bot$ and $\mathbf{c}_{\sigma_2}^{p_2}[i] \neq \bot$, then $\mathbf{c}_{\sigma_1}^{p_1}[i].cfg = \mathbf{c}_{\sigma_2}^{p_2}[i].cfg$. Also from Lemma 13 we know that for $0 \leq j \leq v_1, \mathbf{c}_{\sigma_1}^{p_1}[j] \neq \bot$, and $0 \leq j \leq v_2, \mathbf{c}_{\sigma_2}^{p_2}[j] \neq \bot$. So if we can show that $v_1 \leq v_2$ then the lemma follows.

Let $\mu = \mu(\mathbf{c}_{\sigma'}^{p_2})$ be the last finalized element which p_2 established in the beginning of the read-config action π_2 (Line A2:2) at some state σ' before σ_2. It is easy to see that $\mu \leq v_2$. If $v_1 \leq \mu$ then $v_1 \leq v_2$ and the lemma follows. Thus, it remains to examine the case where $\mu < v_1$. Notice that since $\pi_1 \to \pi_2$ then σ_1 appears before σ' in execution ξ. By Lemma 14, we know that by σ_1, $\exists Q \in \mathbf{c}_{\sigma_1}^{p_1}[j].cfg.Quorums$, for $0 \leq j < v_1$, such that $\forall s \in Q, s.nextC = \mathbf{c}_{\sigma_1}^{p_1}[j+1]$. Since $\mu < v_1$, then it must be the case that $\exists Q \in \mathbf{c}_{\sigma_1}^{p_1}[\mu].cfg.Quorums$ such that $\forall s \in Q, s.nextC = \mathbf{c}_{\sigma_1}^{p_1}[\mu+1]$. But by Lemma 11, we know that $\mathbf{c}_{\sigma_1}^{p_1}[\mu].cfg = \mathbf{c}_{\sigma'}^{p_2}[\mu].cfg$. Let Q' be the quorum that replies to the read-next-config occurred in p_2, on configuration $\mathbf{c}_{\sigma'}^{p_2}[\mu].cfg$. By definition $Q \cap Q' \neq \emptyset$, thus there is a server $s \in Q \cap Q'$ that sends $s.nextC = \mathbf{c}_{\sigma_1}^{p_1}[\mu+1]$ to p_2 during π_2. Since $\mathbf{c}_{\sigma_1}^{p_1}[\mu+1] \neq \bot$ then p_2 assigns $\mathbf{c}_*^{p_2}[\mu+1] = \mathbf{c}_{\sigma_1}^{p_1}[\mu+1]$, and repeats the process in the configuration $\mathbf{c}_*^{p_2}[\mu+1].cfg$. Since every configuration $\mathbf{c}_{\sigma_1}^{p_1}[j].cfg$, for $\mu \leq j < v_1$, has a quorum of servers with $s.nextC$, then by a simple induction it can be shown that the process will be repeated for at least $v_1 - \mu$ iterations, and every configuration $\mathbf{c}_{\sigma''}^{p_2}[j] = \mathbf{c}_{\sigma_1}^{p_1}[j]$, at some state σ'' before σ_2. Thus, $\mathbf{c}_{\sigma_2}^{p_2}[j] = \mathbf{c}_{\sigma_1}^{p_1}[j]$, for $0 \leq j \leq v_1$. Hence $v_1 \leq v_2$ and the lemma follows in this case as well. □

Thus far we focused on the configuration member of each element in *cseq*. As operations do get in account the *status* of a configuration, i.e. P or F, in the next lemma we will examine the relationship of the last finalized configuration as detected by two operations. First we present a lemma that shows the monotonicity of the finalized configurations.

LEMMA 16. *Let σ and σ' two states in an execution ξ such that σ appears before σ' in ξ. Then for any process p must hold that $\mu(\mathbf{c}_\sigma^p) \leq \mu(\mathbf{c}_{\sigma'}^p)$.*

PROOF. This lemma follows from the fact that if a configuration k is such that $\mathbf{c}_\sigma^p[k].status = F$ at a state σ, then p will start any future read-config action from a configuration $\mathbf{c}_{\sigma'}^p[j].cfg$ such that $j \geq k$. But $\mathbf{c}_{\sigma'}^p[j].cfg$ is the last finalized configuration at σ' and hence $\mu(\mathbf{c}_{\sigma'}^p) \geq \mu(\mathbf{c}_\sigma^p)$. □

LEMMA 17 (SEQUENCE PROGRESS). *Let π_1 and π_2 two completed* read-config *actions invoked by processes $p_1, p_2 \in \mathcal{I}$ respectively, such that $\pi_1 \to \pi_2$ in an execution ξ. Let σ_1 be the state after the response step of π_1 and σ_2 the state after the response step of π_2. Then $\mu(\mathbf{c}_{\sigma_1}^{p_1}) \leq \mu(\mathbf{c}_{\sigma_2}^{p_2})$.*

PROOF. By Lemma 15 it follows that $\mathbf{c}_{\sigma_1}^{p_1}$ is a prefix of $\mathbf{c}_{\sigma_2}^{p_2}$. Thus, if $v_1 = v(\mathbf{c}_{\sigma_1}^{p_1})$ and $v_2 = v(\mathbf{c}_{\sigma_2}^{p_2})$, $v_1 \leq v_2$. Let $\mu_1 = \mu(\mathbf{c}_{\sigma_1}^{p_1})$, such that $\mu_1 \leq v_1$, be the last element in $\mathbf{c}_{\sigma_1}^{p_1}$, where $\mathbf{c}_{\sigma_1}^{p_1}[\mu_1].status = F$. Let now $\mu_2 = \mu(\mathbf{c}_{\sigma_2}^{p_2})$, be the last element which p_2 obtained in Line A1:2 during π_2 such that $\mathbf{c}_{\sigma'}^{p_2}[\mu_2].status = F$ in some state σ' before σ_2. If $\mu_2 \geq \mu_1$, and since σ_2 is after σ', then by Lemma 16 $\mu_2 \leq \mu(\mathbf{c}_{\sigma_2}^{p_2})$ and hence $\mu_1 \leq \mu(\mathbf{c}_{\sigma_2}^{p_2})$ as well.

It remains to examine the case where $\mu_2 < \mu_1$. Process p_1 sets the status of $\mathbf{c}_{\sigma_1}^{p_1}[\mu_1]$ to F in two cases: (i) either when finalizing a reconfiguration, or (ii) when receiving an $s.nextC = \langle \mathbf{c}_{\sigma_1}^{p_1}[\mu_1].cfg, F \rangle$ from some server s during

a read-config action. In both cases p_1 propagates the $\langle \mathbf{c}_{\sigma_1}^{p_1}[\mu_1].cfg, F\rangle$ to a quorum of servers in $\mathbf{c}_{\sigma_1}^{p_1}[\mu_1 - 1].cfg$ before completing. We know by Lemma 15 that since $\pi_1 \rightarrow \pi_2$ then $\mathbf{c}_{\sigma_1}^{p_1}$ is a prefix in terms of configurations of the $\mathbf{c}_{\sigma_2}^{p_2}$. So it must be the case that $\mu_2 < \mu_1 \leq \nu(\mathbf{c}_{\sigma_2}^{p_2})$. Thus, during π_2, p_2 starts from the configuration at index μ_2 and in some iteration performs get-next-config in configuration $\mathbf{c}_{\sigma_2}^{p_2}[\mu_1 - 1]$. According to Lemma 11, $\mathbf{c}_{\sigma_1}^{p_1}[\mu_1 - 1].cfg = \mathbf{c}_{\sigma_2}^{p_2}[\mu_1 - 1].cfg$. Since π_1 completed before π_2, then it must be the case that σ_1 appears before σ' in ξ. However, p_2 invokes the get-next-config operation in a state σ'' which is either equal to σ' or appears after σ' in ξ. Thus, σ'' must appear after σ_1 in ξ. From that it follows that when the get-next-config is executed by p_2 there is already a quorum of servers in $\mathbf{c}_{\sigma_2}^{p_2}[\mu_1 - 1].cfg$, say Q_1, that received $\langle \mathbf{c}_{\sigma_1}^{p_1}[\mu_1].cfg, F\rangle$ from p_1. Since, p_2 waits from replies from a quorum of servers from the same configuration, say Q_2, and since the $nextC$ variable at each server is monotonic (Lemma 10), then there is a server $s \in Q_1 \cap Q_2$, such that s replies to p_2 with $s.nextC = \langle \mathbf{c}_{\sigma_1}^{p_1}[\mu_1].cfg, F\rangle$. So, $\mathbf{c}_{\sigma_2}^{p_2}[\mu_1]$ gets $\langle \mathbf{c}_{\sigma_1}^{p_1}[\mu_1].cfg, F\rangle$, and hence $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \geq \mu_1$ in this case as well. This completes our proof. □

Using the previous Lemmas we can conclude to the main result of this section.

THEOREM 18. *Let π_1 and π_2 two completed read-config actions invoked by processes $p_1, p_2 \in \mathcal{I}$ respectively, such that $\pi_1 \rightarrow \pi_2$ in an execution ξ. Let σ_1 be the state after the response step of π_1 and σ_2 the state after the response step of π_2. Then the following properties hold:*

(a) **Configuration Consistency**: $\mathbf{c}_{\sigma_2}^{p_2}[i].cfg = \mathbf{c}_{\sigma_1}^{p_1}[i].cfg$, for $1 \leq i \leq \nu(\mathbf{c}_{\sigma_1}^{p_1})$,

(b) **Sequence Prefix**: $\mathbf{c}_{\sigma_1}^{p_1} \preceq_p \mathbf{c}_{\sigma_2}^{p_2}$, and

(c) **Sequence Progress**: $\mu(\mathbf{c}_{\sigma_1}^{p_1}) \leq \mu(\mathbf{c}_{\sigma_2}^{p_2})$

PROOF. Statements (a), (b) and (c) follow from Lemmas 11, 15, and 16. □

6.2 Atomicity Property of ARES

Given the properties satisfied by the reconfiguration algorithm of ARES in any execution, we can now proceed to examine whether our algorithm satisfies the safety (atomicity) conditions. The propagation of the information of the distributed object is achieved using the get-tag, get-data, and put-data actions. We assume that the DAP used satisfy Property 1 as presented in Section 3, and we will show that, given such assumption, ARES satisfies atomicity.

We begin with a lemma stating that if a reconfiguration operation retrieves a configuration sequence of length k during its read-config action, then it installs/finalizes the $k + 1$ configuration in the global configuration sequence.

LEMMA 19. *Let π be a complete reconfiguration operation by a reconfigurer rc in an execution ξ of ARES. if σ_1 is the state in ξ following the termination of the read-config action during π, then π invokes a finalize-config($\mathbf{c}_{\sigma_2}^{rc}$) at a state σ_2 in ξ, with $\nu(\mathbf{c}_{\sigma_2}^{rc}) = \nu(\mathbf{c}_{\sigma_1}^{rc}) + 1$.*

PROOF. This lemma follows directly from the implementation of the reconfig operation. Let π be a reconfiguration operation reconfig(c). At first, π invokes a read-config to retrieve a latest value of the global configuration sequence, $\mathbf{c}_{\sigma_1}^{rc}$, in the state σ_1 in ξ. During the add-config action, π proposes the addition of c, and appends at the end of $\mathbf{c}_{\sigma_1}^{rc}$ the decision d of the consensus protocol. Therefore, if $\mathbf{c}_{\sigma_1}^{rc}$ is extended by $\langle d, P\rangle$ (Line A 2:17), and hence the add-config action returns a configuration sequence $\mathbf{c}_{\sigma_1'}^{rc}$ with length $\nu(\mathbf{c}_{\sigma_1'}^{rc}) = \nu(\mathbf{c}_{\sigma_1}^{rc}) + 1$. As $\nu(\mathbf{c}_{\sigma_1'}^{rc})$ does not change during the update-config action, then $\mathbf{c}_{\sigma_1'}^{rc}$ is passed to the finalize-config action at state σ_2, and hence $\mathbf{c}_{\sigma_2}^{rc} = \mathbf{c}_{\sigma_1'}^{rc}$. Thus, $\nu(\mathbf{c}_{\sigma_2}^{rc}) = \nu(\mathbf{c}_{\sigma_1'}^{rc}) = \nu(\mathbf{c}_{\sigma_1}^{rc}) + 1$ and the lemma follows. □

The next lemma states that only some reconfiguration operation π may finalize a configuration c at index j in a configuration sequence $p.cseq$ at any process p. To finalize c, the lemma shows that π must witness a

configuration sequence such that its last finalized configuration appears at an index $i < j$ in the configuration sequence $p.cseq$. In other words, reconfigurations always finalize configurations that are ahead from their latest observed final configuration, and it seems like "jumping" from one final configuration to the next.

LEMMA 20. *Suppose ξ is an execution of* ARES. *For any state σ in ξ, if $\mathbf{c}_\sigma^p[j].status = F$ for some process $p \in \mathcal{I}$, then there exists a* reconfig *operation π by a reconfigurer $rc \in \mathcal{G}$, such that (i) rc invokes* finalize-config($\mathbf{c}_{\sigma'}^{rc}$) *during π at some state σ' in ξ, (ii) $v(\mathbf{c}_{\sigma'}^{rc}) = j$, and (iii) $\mu(\mathbf{c}_{\sigma'}^{rc}) < j$.*

PROOF. A process sets the status of a configuration c to F in two cases: (i) either during a finalize-config(seq) action such that $v(seq) = \langle c, P \rangle$ (Line A2:33), or (ii) when it receives $\langle c, F \rangle$ from a server s during a read-next-config action. Server s sets the status of a configuration c to F only if it receives a message that contains $\langle c, F \rangle$ (Line A3:10). So, (ii) is possible only if c is finalized during a reconfig operation.

Let, w.l.o.g., π be the first reconfiguration operation that finalizes $\mathbf{c}_\sigma^p[j].cfg$. To do so, process rc invokes finalize-config($\mathbf{c}_{\sigma'_1}^{rc}$) during π, at some state σ' that appears before σ in ξ. By Lemma 11, $\mathbf{c}_\sigma^p[j].cfg = \mathbf{c}_{\sigma'}^{rc}[j].cfg$. Since, rc finalizes $\mathbf{c}_{\sigma'}^{rc}[j]$, then this is the last entry of $\mathbf{c}_{\sigma'}^{rc}$ and hence $v(\mathbf{c}_{\sigma'}^{rc}) = j$. Also, by Lemma 20 it follows that the read-config action of π returned a configuration $\mathbf{c}_{\sigma''}^{rc}$ in some state σ'' that appeared before σ' in ξ, such that $v(\mathbf{c}_{\sigma''}^{rc}) < v(\mathbf{c}_{\sigma'}^{rc})$. Since by definition, $\mu(\mathbf{c}_{\sigma''}^{rc}) \leq v(\mathbf{c}_{\sigma''}^{rc})$, then $\mu(\mathbf{c}_{\sigma''}^{rc}) < j$. However, since only $\langle c, P \rangle$ is added to $\mathbf{c}_{\sigma''}^{rc}$ to result in $\mathbf{c}_{\sigma'}^{rc}$, then $\mu(\mathbf{c}_{\sigma''}^{rc}) = \mu(\mathbf{c}_{\sigma'}^{rc})$. Therefore, $\mu(\mathbf{c}_{\sigma'}^{rc}) < j$ as well and the lemma follows. □

We now reach an important lemma of this section. By ARES, before a read/write/reconfig operation completes it propagates the maximum tag it discovered by executing the put-data action in the last configuration of its local configuration sequence (Lines A2:18, A4:16, A4:38). When a subsequent operation is invoked, it reads the latest configuration sequence by beginning from the last finalized configuration in its local sequence and invoking read-data to all the configurations until the end of that sequence. The lemma shows that the latter operation will retrieve a tag which is higher than the tag used in the put-data action of any preceding operation.

LEMMA 21. *Let π_1 and π_2 be two completed read/write/reconfig operations invoked by processes p_1 and p_2 in $\mathcal{I}$, in an execution ξ of* ARES, *such that, $\pi_1 \rightarrow \pi_2$. If c_1.put-data($\langle \tau_{\pi_1}, v_{\pi_1} \rangle$) is the last put-data action of π_1 and σ_2 is the state in ξ after the completion of the first read-config action of π_2, then there exists a c_2.put-data($\langle \tau, v \rangle$) action in some configuration $c_2 = \mathbf{c}_{\sigma_2}^{p_2}[k].cfg$, for $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \leq k \leq v(\mathbf{c}_{\sigma_2}^{p_2})$, such that (i) it completes in a state σ' before σ_2 in ξ, and (ii) $\tau \geq \tau_{\pi_1}$.*

PROOF. Note that from the definitions of $v(\cdot)$ and $\mu(\cdot)$, we have $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \leq v(\mathbf{c}_{\sigma_2}^{p_2})$. Let σ_1 be the state in ξ after the completion of c_1.put-data($\langle \tau_{\pi_1}, v_{\pi_1} \rangle$) and σ'_1 be the state in ξ following the response step of π_1. Since any operation executes put-data on the last discovered configuration then c_1 is the last configuration found in $\mathbf{c}_{\sigma_1}^{p_1}$, and hence $c_1 = \mathbf{c}_{\sigma_1}^{p_1}[v(\mathbf{c}_{\sigma_1}^{p_1})].cfg$. By Lemma 16 we have $\mu(\mathbf{c}_{\sigma_1}^{p_1}) \leq \mu(\mathbf{c}_{\sigma'_1}^{p_1})$ and by Lemma 17 we have $\mu(\mathbf{c}_{\sigma'_1}^{p_1}) \leq \mu(\mathbf{c}_{\sigma_2}^{p_2})$, since π_2 (and thus its first read-config action) is invoked after σ'_1 (and thus after the last read-config action during π_1). Hence, combining the two implies that $\mu(\mathbf{c}_{\sigma_1}^{p_1}) \leq \mu(\mathbf{c}_{\sigma_2}^{p_2})$. Now from the last implication and the first statement we have $\mu(\mathbf{c}_{\sigma_1}^{p_1}) \leq v(\mathbf{c}_{\sigma_2}^{p_2})$. Therefore, it remains to examine whether the last finalized configuration witnessed by p_2 appears before or after c_1, i.e.: (a) $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \leq v(\mathbf{c}_{\sigma_1}^{p_1})$ and (b) $\mu(\mathbf{c}_{\sigma_2}^{p_2}) > v(\mathbf{c}_{\sigma_1}^{p_1})$.

Case (*a*): Since $\pi_1 \rightarrow \pi_2$ then, by Theorem 18, $\mathbf{c}_{\sigma_2}^{p_2}$ value returned by read-config at p_2 during the execution of π_2 satisfies $\mathbf{c}_{\sigma_1}^{p_1} \preceq_p \mathbf{c}_{\sigma_2}^{p_2}$. Therefore, $v(\mathbf{c}_{\sigma_1}^{p_1}) \leq v(\mathbf{c}_{\sigma_2}^{p_2})$, and hence in this case $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \leq v(\mathbf{c}_{\sigma_1}^{p_1}) \leq v(\mathbf{c}_{\sigma_2}^{p_2})$. Since c_1 is the last configuration in $\mathbf{c}_{\sigma_1}^{p_1}$, then it has index $v(\mathbf{c}_{\sigma_1}^{p_1})$. So if we take $c_2 = c_1$ then the c_1.put-data($\langle \tau_{\pi_1}, v_{\pi_1} \rangle$) action trivially satisfies both conditions of the lemma as: (i) it completes in state σ_1 which appears before σ_2, and (ii) it puts a pair $\langle \tau, v \rangle$ such that $\tau = \tau_{\pi_1}$.

Case (b): This case is possible if there exists a reconfiguration client rc that invokes reconfig operation ρ, during which it executes the finalize-config($\mathbf{c}_*^{rc}$) that finalized configuration with index $\nu(\mathbf{c}_*^{rc}) = \mu(\mathbf{c}_{\sigma_2}^{p_2})$. Let σ be the state immediately after the read-config of ρ. Now, we consider two sub-cases: (i) σ appears before σ_1 in ξ, or (ii) σ appears after σ_1 in ξ.

Subcase (b)(i): Since read-config at σ completes before the invocation of last read-config of operation π_1 then, either $\mathbf{c}_\sigma^{rc} \prec_p \mathbf{c}_{\sigma_1}^{p_1}$, or $\mathbf{c}_\sigma^{rc} = \mathbf{c}_{\sigma_1}^{p_1}$ due to Lemma 15. Suppose $\mathbf{c}_\sigma^{rc} \prec_p \mathbf{c}_{\sigma_1}^{p_1}$, then according to Lemma 19 rc executes finalize-config on configuration sequence $\mathbf{c}_*^{rc}$ with $\nu(\mathbf{c}_*^{rc}) = \nu(\mathbf{c}_\sigma^{rc}) + 1$. Since $\nu(\mathbf{c}_*^{rc}) = \mu(\mathbf{c}_{\sigma_2}^{p_2})$, then $\mu(\mathbf{c}_{\sigma_2}^{p_2}) = \nu(\mathbf{c}_\sigma^{rc}) + 1$. If however, $\mathbf{c}_\sigma^{rc} \prec_p \mathbf{c}_{\sigma_1}^{p_1}$, then $\nu(\mathbf{c}_\sigma^{rc}) < \nu(\mathbf{c}_{\sigma_1}^{p_1})$ and thus $\nu(\mathbf{c}_\sigma^{rc}) + 1 \leq \nu(\mathbf{c}_{\sigma_1}^{p_1})$. This implies that $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \leq \nu(\mathbf{c}_{\sigma_1}^{p_1})$ which contradicts our initial assumption for this case that $\mu(\mathbf{c}_{\sigma_2}^{p_2}) > \nu(\mathbf{c}_{\sigma_1}^{p_1})$. So this sub-case is impossible.

Now suppose, that $\mathbf{c}_\sigma^{rc} = \mathbf{c}_{\sigma_1}^{p_1}$. Then it follows that $\nu(\mathbf{c}_\sigma^{rc}) = \nu(\mathbf{c}_{\sigma_1}^{p_1})$, and that $\mu(\mathbf{c}_{\sigma_2}^{p_2}) = \nu(\mathbf{c}_{\sigma_1}^{p_1}) + 1$ in this case. Since σ_1 is the state after the last put-data during π_1, then if σ_1' is the state after the completion of the last read-config of π_1 (which follows the put-data), it must be the case that $\mathbf{c}_{\sigma_1}^{p_1} = \mathbf{c}_{\sigma_1'}^{p_1}$. So, during its last read-config process p_1 does not read the configuration indexed at $\nu(\mathbf{c}_{\sigma_1}^{p_1}) + 1$. This means that the put-config completes in ρ at state σ_ρ after σ_1' and the update-config operation is invoked at state σ_ρ' after σ_ρ with a configuration sequence $\mathbf{c}_{\sigma_\rho'}^{rc}$. During the update operation ρ invokes get-data operation in every configuration $\mathbf{c}_{\sigma_\rho'}^{rc}[i].cfg$, for $\mu(\mathbf{c}_{\sigma_\rho'}^{rc}) \leq i \leq \nu(\mathbf{c}_{\sigma_\rho'}^{rc})$. Notice that $\nu(\mathbf{c}_{\sigma_\rho'}^{rc}) = \mu(\mathbf{c}_{\sigma_2}^{p_2}) = \nu(\mathbf{c}_{\sigma_1}^{p_1}) + 1$ and moreover the last configuration of $\mathbf{c}_{\sigma_\rho'}^{rc}$ was just added by ρ and it is not finalized. From this it follows that $\mu(\mathbf{c}_{\sigma_\rho'}^{rc}) < \nu(\mathbf{c}_{\sigma_\rho'}^{rc})$, and hence $\mu(\mathbf{c}_{\sigma_\rho'}^{rc}) \leq \nu(\mathbf{c}_{\sigma_1}^{p_1})$. Therefore, ρ executes get-data in configuration $\mathbf{c}_{\sigma_\rho'}^{rc}[j].cfg$ for $j = \nu(\mathbf{c}_{\sigma_1}^{p_1})$. Since p_1 invoked put-data($\langle \tau_{\pi_1}, \upsilon_{\pi_1} \rangle$) at the same configuration c_1, and completed in a state σ_1 before σ_ρ', then by **C1** of Property 1, it follows that the get-data action will return a tag $\tau \geq \tau_{\pi_1}$. Therefore, the maximum tag that ρ discovers is $\tau_{max} \geq \tau \geq \tau_{\pi_1}$. Before invoking the finalize-config action, ρ invokes c_1.put-data($\langle \tau_{max}, \upsilon_{max} \rangle$). Since $\nu(\mathbf{c}_{\sigma_\rho'}^{rc}) = \mu(\mathbf{c}_{\sigma_2}^{p_2})$, and since by Lemma 11, then the action put-data is invoked in a configuration $c_2 = \mathbf{c}_{\sigma_2}^{p_2}[j].cfg$ such that $j = \mu(\mathbf{c}_{\sigma_2}^{p_2})$. Since the read-config action of π_2 observed configuration $\mu(\mathbf{c}_{\sigma_2}^{p_2})$, then it must be the case that σ_2 appears after the state where the finalize-config was invoked and therefore after the state of completion of the put-data action during ρ. Thus, in this case both properties are satisfied and the lemma follows.

Subcase (b)(ii): Suppose in this case that σ occurs in ξ after σ_1. In this case the last put-data in π_1 completes before the invocation of the read-config in ρ in execution ξ. Now we can argue recursively, ρ taking the place of operation π_2, that $\mu(\mathbf{c}_\sigma^{rc}) \leq \nu(\mathbf{c}_\sigma^{rc})$ and therefore, we consider two cases: (a) $\mu(\mathbf{c}_\sigma^{rc}) \leq \nu(\mathbf{c}_{\sigma_1}^{p_1})$ and (b) $\mu(\mathbf{c}_\sigma^{rc}) > \nu(\mathbf{c}_{\sigma_1}^{p_1})$. Note that there are finite number of operations invoked in ξ before π_2 is invoked, and hence the statement of the lemma can be shown to hold by a sequence of inequalities. □

The following lemma shows the consistency of operations as long as the DAP used satisfy Property 1.

LEMMA 22. *Let π_1 and π_2 denote completed read/write operations in an execution ξ, from processes $p_1, p_2 \in \mathcal{I}$ respectively, such that $\pi_1 \rightarrow \pi_2$. If τ_{π_1} and τ_{π_2} are the local tags at p_1 and p_2 after the completion of π_1 and π_2 respectively, then $\tau_{\pi_1} \leq \tau_{\pi_2}$; if π_1 is a write operation then $\tau_{\pi_1} < \tau_{\pi_2}$.*

PROOF. Let $\langle \tau_{\pi_1}, \upsilon_{\pi_1} \rangle$ be the pair passed to the last put-data action of π_1. Also, let σ_2 be the state in ξ that follows the completion of the first read-config action during π_2. Notice that π_2 executes a loop after the first read-config operation and performs c.get-data (if π_2 is a read) or c.get-tag (if π_2 is a write) from all $c = \mathbf{c}_{\sigma_2}^{p_2}[i].cfg$, for $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \leq i \leq \nu(\mathbf{c}_{\sigma_2}^{p_2})$. By Lemma 21, there exists a c'.put-data($\langle \tau, \upsilon \rangle$) action by some operation π' on some configuration $c' = \mathbf{c}_{\sigma_2}^{p_2}[j].cfg$, for $\mu(\mathbf{c}_{\sigma_2}^{p_2}) \leq j \leq \nu(\mathbf{c}_{\sigma_2}^{p_2})$, that completes in some state σ' that appears before σ_2 in ξ. Thus, the get-data or get-tag invoked by p_2 on $\mathbf{c}_{\sigma_2}^{p_2}[j].cfg$, occurs after state σ_2 and thus after σ'. Since the DAP

primitives used satisfy **C1** and **C2** of Property 1, then the get-tag action will return a tag τ'_{π_2} or a get-data action will return a pair $\langle \tau'_{\pi_2}, v'_{\pi_2} \rangle$, with $\tau'_{\pi_2} \geq \tau$. As p_2 gets the maximum of all the tags returned, then by the end of the loop p_2 will retrieve a tag $\tau_{max} \geq \tau'_{\pi_2} \geq \tau \geq \tau_{\pi_1}$.

If now π_2 is a read, it returns $\langle \tau_{max}, v_{max} \rangle$ after propagating that value to the last discovered configuration. Thus, $\tau_{\pi_2} \geq \tau_{\pi_1}$. If however π_2 is a write, then before propagating the new value the writer increments the maximum timestamp discovered (Line A4:13) generating a tag $\tau_{\pi_2} > \tau_{max}$. Therefore the operation π_2 propagates a tag $\tau_{\pi_2} > \tau_{\pi_1}$ in this case. □

And the main result of this section follows:

THEOREM 23 (ATOMICITY). *In any execution ξ of* ARES, *if in every configuration $c \in \mathcal{G}_L$, c.get-data(), c.put-data(), and c.get-tag() satisfy Property 1, then* ARES *satisfy atomicity.*

As algorithm ARES handles each configuration separately, then we can observe that the algorithm may utilize a different mechanism for the put and get primitives in each configuration. So the following remark:

REMARK 24. *Algorithm* ARES *satisfies atomicity even when the implementaton of the DAPs in two different configurations c_1 and c_2 are not the same, given that the c_i.get-tag, c_i.get-data, and the c_i.put-data primitives in each c_i satisfy Property 1.*

7 PERFORMANCE ANALYSIS OF ARES

A major challenge in reconfigurable atomic services is to examine the latency of terminating read and write operations, especially when those are invoked concurrently with reconfiguration operations. In this section we provide an in depth analysis of the latency of operations in ARES. Additionally, a storage and communication analysis is shown when ARES utilizes the erasure-coding algorithm presented in Section 5, in each configuration.

7.1 Latency Analysis

The idea behind our latency analysis is quite straight forward: we construct the worst case execution that would allow all concurrent reconfigurations to add their proposed configuration. This leads to the longest configuration sequence that a read/write operation needs to traverse before completing. Thus, given a bounded delay, we compute the delay for each operation and we finally compute how long it is going to take for a read/write operation to catch up in the worst case and complete.

Liveness (termination) properties cannot be specified for ARES, without restricting asynchrony or the rate of arrival of reconfig operations, or if the consensus protocol never terminates. Here, we provide some conditional performance analysis of the operation, based on latency bounds on the message delivery. We assume that local computations take negligible time and the latency of an operation is due to the delays in the messages exchanged during the execution. We measure delays in *time units* of some global clock, which is visible only to an external viewer. No process has access to the clock. Let d and D be the minimum and maximum durations taken by messages, sent during an execution of ARES, to reach their destinations. Also, let $T(\pi)$ denote the duration of an operation (or action) π. In the statements that follow, we consider any execution ξ of ARES, which contains k reconfig operations. For any configuration c in an execution of ARES, we assume that any $c.Con.$propose operation, takes at least $T_{min}(CN)$ time units.

Let us first examine what is the action delays based on the boundaries we assume. It is easy to see that actions put-config, get-next-config perform two message exchanges thus take time $2d \leq T(\phi) \leq 2D$. From this we can derive the delay of a read-config action.

LEMMA 25. *Let ϕ be a* read-config *operation invoked by a non-faulty reconfiguration client rc, with the input argument and returned values of ϕ as $\mathbf{c}_\sigma^{rc}$ and $\mathbf{c}_{\sigma'}^{rc}$ respectively. Then the delay of ϕ is: $4d(v(\mathbf{c}_{\sigma'}^{rc}) - \mu(\mathbf{c}_{\sigma'}^{rc}) + 1) \leq T(\phi) \leq 4D(v(\mathbf{c}_{\sigma'}^{rc}) - \mu(\mathbf{c}_\sigma^{rc}) + 1)$.*

From Lemma 25 it is clear that the latency of a read-config action depends on the number of configurations installed since the last finalized configuration known to the recon client.

Given the latency of a read-config, we can compute the minimum amount of time it takes for k configurations to be installed.

The following lemma shows the maximum latency of a read or a write operation, invoked by any non-faulty client. From Ares algorithm, the latency of a read/write operation depends on the delays of the DAPs operations. For our analysis we assume that all get-data, get-tag and put-data primitives use two phases of communication. Each phase consists of a communication between the client and the servers.

Lemma 26. *Suppose π, ϕ and ψ are operations of the type* put-data, get-tag *and* get-data, *respectively, invoked by some non-faulty reconfiguration clients, then the latency of these operations are bounded as follows:* (*i*) $2d \leq T(\pi) \leq 2D$; (*ii*) $2d \leq T(\phi) \leq 2D$; *and* (*iii*) $2d \leq T(\psi) \leq 2D$.

In the following lemma, we estimate the time taken for a read or a write operation to complete, when it discovers k configurations between its invocation and response steps.

Lemma 27. *Consider any execution of Ares where at most k reconfiguration operations are invoked. Let σ_s and σ_e be the states before the invocation and after the completion step of a read/write operation π, in some fair execution ξ of Ares. Then we have $T(\pi) \leq 6D(k+2)$ to complete.*

Proof. Let σ_s and σ_e be the states before the invocation and after the completion step of a read/write operation π by p respectively, in some execution ξ of Ares. By algorithm examination we can see that any read/write operation performs the following actions in this order: (*i*) read-config, (*ii*) get-data (or get-tag), (*iii*) put-data, and (*iv*) read-config. Let σ_1 be the state when the first read-config, denoted by read-config$_1$, action terminates. By Lemma 25 the action will take time:

$$T(\text{read-config}_1) \leq 4D(v(\mathbf{c}^p_{\sigma_1}) - \mu(\mathbf{c}^p_{\sigma_s}) + 1)$$

The get-data action that follows the read-config (Lines Alg. 4:34-35) also took at most $(v(\mathbf{c}^p_{\sigma_1}) - \mu(\mathbf{c}^p_{\sigma_s}) + 1)$ time units, given that no new finalized configuration was discovered by the read-config action. Finally, the put-data and the second read-config actions of π may be invoked at most $(v(\mathbf{c}^p_{\sigma_e}) - v(\mathbf{c}^p_{\sigma_1}) + 1)$ times, given that the read-config action discovers one new configuration every time it runs. Merging all the outcomes, the total time of π can be at most:

$$
\begin{aligned}
T(\pi) &\leq 4D(v(\mathbf{c}^p_{\sigma_1}) - \mu(\mathbf{c}^p_{\sigma_s}) + 1) + 2D(v(\mathbf{c}^p_{\sigma_1}) - \mu(\mathbf{c}^p_{\sigma_s}) + 1) + (4D + 2D)(v(\mathbf{c}^p_{\sigma_e}) - v(\mathbf{c}^p_{\sigma_1}) + 1) \\
&\leq 6D\left[v(\mathbf{c}^p_{\sigma_e}) - \mu(\mathbf{c}^p_{\sigma_s}) + 2\right] \leq 6D(k+1)
\end{aligned}
$$

where $v(\mathbf{c}^p_{\sigma_e}) - \mu(\mathbf{c}^p_{\sigma_s}) \leq k+1$ since there can be at most k new configurations installed. and the result of the lemma follows. □

It remains now to examine the conditions under which a read/write operation may "catch up" with an infinite number of reconfiguration operations. For the sake of a worst case analysis we will assume that reconfiguration operations suffer the minimum delay d, whereas read and write operations suffer the maximum delay D in each message exchange. We first show how long it takes for k configurations to be installed.

Lemma 28. *Let σ be the last state of a fair execution of Ares, ξ. Then k configurations can be installed to $\mathbf{c}_\sigma$, in time $T(k) \geq 4d \sum_{i=1}^{k} i + k(T_{min}(CN) + 2d)$ time units.*

Proof. In Ares a reconfig operation has four phases: (*i*) read-config(*cseq*), reads the latest configuration sequence, (*ii*) add-config(*cseq*, *c*), attempts to add the new configuration at the end of the global sequence $\mathcal{G}_L$, (*iii*) update-config(*cseq*), transfers the knowledge to the added configuration, and (*iv*) finalize-config(*cseq*)
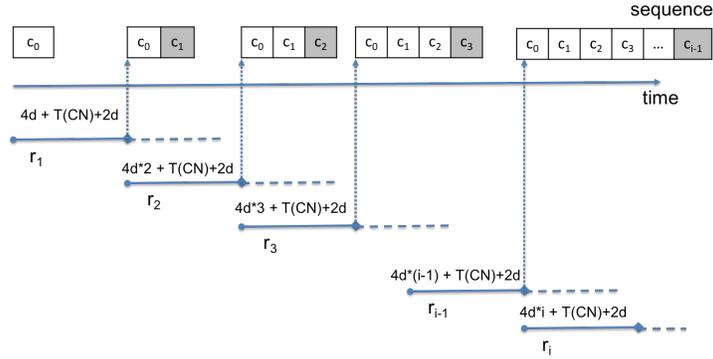
Fig. 2. Successful reconfig operations.

finalizes the added configuration. So, a new configuration is appended to the end of the configuration sequence (and it becomes visible to any operation) during the add-config action. In turn, the add-config action, runs a consensus algorithm to decide on the added configuration and then invokes a put-config action to add the decided configuration. Any operation that is invoked after the put-config action observes the newly added configuration.

Notice that when multiple reconfigurations are invoked concurrently, then it might be the case that all participate to the same consensus instance and the configuration sequence is appended by a single configuration. The worst case scenario happens when all concurrent reconfigurations manage to append the configuration sequence by their configuration. In brief, this is possible when the read-config action of each reconfig operation appears after the put-config action of another reconfig operation.

More formally we can build an execution where all reconfig operations append their configuration in the configuration sequence. Consider the partial execution ξ that ends in a state σ. Suppose that every process $p \in \mathcal{I}$ knows the same configuration sequence, $\mathbf{c}_\sigma^p = \mathbf{c}_\sigma$. Also let the last finalized operation in $\mathbf{c}_\sigma$ be the last configuration of the sequence, e.g. $\mu(\mathbf{c}_\sigma) = \nu(\mathbf{c}_\sigma)$. Notice that $\mathbf{c}_\sigma$ can also be the initial configuration sequence $\mathbf{c}_{\sigma_0}^p$. We extend ξ_0 by a series of reconfig operations, such that each reconfiguration rc_i is invoked by a reconfigurer r_i and attempts to add a configuration c_i. Let rc_1 be the first reconfiguration that performs the following actions without being concurrent with any other reconfig operation:

- read-config starting from $\mu(\mathbf{c}_\sigma)$
- add-config completing both the consensus proposing c_1 and the put-config action writing the decided configuration

Since rc_1 its not concurrent with any other reconfig operation, then is the only process to propose a configuration in $\mu(\mathbf{c}_\sigma)$, and hence by the consensus algorithm properties, c_1 is decided. Thus, $\mathbf{c}_\sigma$ is appended by a tuple $\langle c_1, P \rangle$.

Let now reconfiguration rc_2 be invoked immediately after the completion of the add-config action from rc_1. Since the local sequence at the beginning of rc_2 is equal to $\mathbf{c}_\sigma$, then the read-config action of rc_2 will also start from $\mu(\mathbf{c}_\sigma)$. Since, rc_1 already propagated c_1 to $\mu(\mathbf{c}_\sigma)$ during is put-config action, then rc_2 will discover c_1 during the first iteration of its read-config action, and thus it will repeat the iteration on c_1. Configuration c_1 is the last in the sequence and thus the read-config action of rc_2 will terminate after the second iteration. Following the read-config action, rc_2 attempts to add c_2 in the sequence. Since rc_1 is the only reconfiguration that might be concurrent with rc_2, and since rc_1 already completed consensus in $\mu(\mathbf{c}_\sigma)$, then rc_2 is the only operation to run consensus in c_1. Therefore, c_2 is accepted and rc_2 propagates c_2 in c_1 using a put-config action.

So in general we let configuration rc_i to be invoked after the completion of the add-config action from rc_{i-1}. As a result, the read-config action of rc_i performs i iterations, and the configuration c_i is added immediately after configuration c_{i-1} in the sequence. Figure 2 illustrates our execution construction for the reconfiguration operations.

It is easy to notice that such execution results in the worst case latency for all the reconfiguration operations $rc_1, rc_2, \ldots, rc_i$. As by Lemma 25 a read-config action takes at least $4d$ time to complete, then as also seen in Figure 2, k reconfigs may take time $T(k) \geq \sum_{i=1}^{k} [4d * i + (T_{min}(CN) + 2d)]$. Therefore, it will take time $T(k) \geq 4d \sum_{i=1}^{k} i + k(T_{min}(CN) + 2d)$ and the lemma follows. □

The following theorem is the main result of this section, in which we define the relation between $T_{min}(CN)$, d and D so to guarantee that any read or write issued by a non-faulty client always terminates.

THEOREM 29. *Suppose $T_{min}(CN) \geq 3(6D - d)$, then any read or write operation π completes in any execution ξ of* ARES *for any number of reconfiguration operations in ξ.*

PROOF. By Lemma 28, k configurations may be installed in: $T(k) \geq 4d \sum_{i=1}^{k} i + k(T_{min}(CN) + 2d)$. Also by Lemma 27, we know that operation π takes at most $T(\pi) \leq 6D \left(\nu(\mathbf{c}_{\sigma_e}^p) - \mu(\mathbf{c}_{\sigma_s}^p) + 2 \right)$. Assuming that $k = \nu(\mathbf{c}_{\sigma_e}^p) - \mu(\mathbf{c}_{\sigma_s}^p)$, the total number of configurations observed during π, then π may terminate before a $k + 1$ configuration is added in the configuration sequence if $6D(k + 2) \leq 4d \sum_{i=1}^{k} i + k(T_{min}(CN) + 2d)$ then we have $d \geq \frac{3D}{k} - \frac{T_{min}(CN)}{2(k+2)}$. And that completes the lemma. □

7.2 Storage and Communication Costs for ARES.

Storage and Communication costs for ARES highly depends on the DAP that we use in each configuration. For our analysis we assume that each configuration utilizes the algorithms and the DAPs presented in Section 5.

Recall that by our assumption, the storage cost counts the size (in bits) of the coded elements stored in variable *List* at each server. We ignore the storage cost due to meta-data. For communication cost we measure the bits sent on the wire between the nodes.

LEMMA 30. *The worst-case total storage cost of Algorithm 5 is $(\delta + 1)\frac{n}{k}$.*

PROOF. The maximum number of (tag, coded-element) pair in the *List* is $\delta+1$, and the size of each coded element is $\frac{1}{k}$ while the tag variable is a metadata and therefore, not counted. So, the total storage cost is $(\delta + 1)\frac{n}{k}$. □

We next state the communication cost for the write and read operations in Aglorithm 5. Once again, note that we ignore the communication cost arising from exchange of meta-data.

LEMMA 31. *The communication cost associated with a successful write operation in Algorithm 5 is at most $\frac{n}{k}$.*

PROOF. During read operation, in the get-tag phase the servers respond with their highest tags variables, which are metadata. However, in the put-data phase, the reader sends each server the coded elements of size $\frac{1}{k}$ each, and hence the total cost of communication for this is $\frac{n}{k}$. Therefore, we have the worst case communication cost of a write operation is $\frac{n}{k}$. □

LEMMA 32. *The communication cost associated with a successful read operation in Algorithm 5 is at most $(\delta + 2)\frac{n}{k}$.*

PROOF. During read operation, in the get-data phase the servers respond with their *List* variables and hence each such list is of size at most $(\delta + 1)\frac{1}{k}$, and then counting all such responses give us $(\delta + 1)\frac{n}{k}$. In the put-data phase, the reader sends each server the coded elements of size $\frac{1}{k}$ each, and hence the total cost of communication for this is $\frac{n}{k}$. Therefore, we have the worst case communication cost of a read operation is $(\delta + 2)\frac{n}{k}$. □

From the above Lemmas we get.

THEOREM 33. *The* ARES *algorithm has: (i) storage cost* $(\delta + 1)\frac{n}{k}$*, (ii) communication cost for each write at most to* $\frac{n}{k}$*, and (iii) communication cost for each read at most* $(\delta + 2)\frac{n}{k}$*.*

8 FLEXIBILITY OF DAPS

In this section, we argue that various implementations of DAPs can be used in ARES. In fact, via reconfig operations, one can implement a highly adaptive atomic DSS: replication-based can be transformed into erasure-code based DSS; increase or decrease the number of storage servers; study the performance of the DSS under various code parameters, etc. The insight to implementing various DAPs comes from the observation that the simple algorithmic template A (see Alg. 7) for reads and writes protocol combined with any implementation of DAPs, satisfying Property 1 gives rise to a MWMR atomic memory service. Moreover, the read and writes operations terminate as long as the implemented DAPs complete.

Algorithm 7 Template A for the client-side read/write steps.

```
       operation read()                          6:  operation write(v)
2:       ⟨t, v⟩ ← c.get-data()                        t ← c.get-tag()
         c.put-data(⟨t, v⟩)                      8:     t_w ← inc(t)
4:       return ⟨t, v⟩                                 c.put-data(⟨t_w, v⟩)
       end operation
                                                10:  end operation
```

A read operation in A performs c.get-data() to retrieve a tag-value pair, $\langle \tau, v \rangle$ from a configuration c, and then it performs a c.put-data($\langle \tau, v \rangle$) to propagate that pair to the configuration c. A write operation is similar to the read but before performing the put-data action it generates a new tag which associates with the value to be written. The following result shows that A is atomic and live, if the DAPs satisfy Property 1 and live.

THEOREM 34 (ATOMICITY OF TEMPLATE A). *Suppose the DAP implementation satisfies the consistency properties* $C1$ *and* $C2$ *of Property 1 for a configuration* $c \in C$*. Then any execution* ξ *of algorithm* A *in configuration* c *is atomic and live if each DAP invocation terminates in* ξ *under the failure model* $c.\mathcal{F}$*.*

PROOF. We prove the atomicity by proving properties $A1$, $A2$ and $A3$ presented in Section 2 for any execution of the algorithm.

Property $A1$: Consider two operations ϕ and π such that ϕ completes before π is invoked. We need to show that it cannot be the case that $\pi \prec \phi$. We break our analysis into the following four cases:

Case (a): *Both ϕ and π are writes.* The c.put-data($*$) of ϕ completes before π is invoked. By property **C1** the tag τ_π returned by the c.get-data() at π is at least as large as τ_ϕ. Now, since τ_π is incremented by the write operation then π puts a tag τ_π' such that $\tau_\phi < \tau_\pi'$ and hence we cannot have $\pi \prec \phi$.

Case (b): *ϕ is a write and π is a read.* In execution ξ since c.put-data($\langle t_\phi, * \rangle$) of ϕ completes before the c.get-data() of π is invoked, by property **C1** the tag τ_π obtained from the above c.get-data() is at least as large as τ_ϕ. Now $\tau_\phi \leq \tau_\pi$ implies that we cannot have $\pi \prec \phi$.

Case (c): *ϕ is a read and π is a write.* Let the id of the writer that invokes π we w_π. The c.put-data($\langle \tau_\phi, * \rangle$) call of ϕ completes before c.get-tag() of π is initiated. Therefore, by property **C1** get-tag(c) returns τ such that, $\tau_\phi \leq \tau$. Since τ_π is equal to $inc(\tau)$ by design of the algorithm, hence $\tau_\pi > \tau_\phi$ and we cannot have $\pi \prec \phi$.

Case (d): *Both ϕ and π are reads.* In execution ξ the c.put-data($\langle t_\phi, * \rangle$) is executed as a part of ϕ and completes before c.get-data() is called in π. By property **C1** of the data-primitives, we have $\tau_\phi \leq \tau_\pi$ and hence we cannot have $\pi \prec \phi$.

Property A2: Note that because the tag set $\mathcal{T}$ is well-ordered we can show that A2 holds by first showing that every write has a unique tag. This means that any two pair of writes can be ordered. Note that a read can be ordered w.r.t. any write operation trivially if the respective tags are different, and by definition, if the tags are equal the write is ordered before the read.

Observe that two tags generated from different writers are necessarily distinct because of the id component of the tag. Now if the operations, say ϕ and π are writes from the same writer then, by well-formedness property, the second operation will witness a higher integer part in the tag by property **C1**, and since the c.get-tag() is followed by c.put-data($*$). Hence π is ordered after ϕ.

Property A3: By **C2** the c.get-data() may return a tag τ, only when there exists an operation π that invoked a c.put-data($\langle \tau, * \rangle$). Otherwise it returns the initial value. Since a write is the only operation to put a new tag τ in the system then Property $A3$ follows from **C2**. $\qquad\square$

8.1 Representing Known Algorithms in terms of data-access primitives

A number of known tag-based algorithms that implement atomic read/write objects (e.g., ABD [11], FAST[21]), can be expressed in terms of DAP. In this subsection we demonstrate how we can transform the very celebrated ABD algorithm [11].

MWABD *Algorithm.* The multi-writer version of the ABD can be transformed to the generic algorithm Template A. Algorithm 8 illustrates the three DAP for the ABD algorithm. The get-data primitive encapsulates the query phase of MWABD, while the put-data primitive encapsulates the propagation phase of the algorithm.

Algorithm 8 Implementation of DAP for ABD at each process p using configuration c

---

Data-Access Primitives at process p:

2:    **procedure** c.put-data($\langle \tau, v \rangle$))
      send (WRITE, $\langle \tau, v \rangle$) to each $s \in c.Servers$
4:      **until** $\exists Q, Q_{\in} c.Quorums$ s.t. p receives ACK from $\forall s \in Q$
   end procedure

6:    **procedure** c.get-tag()
      send (QUERY-TAG) to each $s \in c.Servers$
8:      **until** $\exists Q, Q \in c.Quorums$ s.t.
        p receives $\langle \tau_s, v_s \rangle$ from $\forall s \in Q$
10:     $\tau_{max} \leftarrow \max(\{\tau_s : p \text{ received } \langle \tau_s, v_s \rangle \text{ from } s\})$

        return τ_{max}
12:    **end procedure**

   procedure c.get-data()
14:      **send** (QUERY) to each $s \in c.Servers$
      until $\exists Q, Q \in c.Quorums$ s.t.
16:        p receives $\langle \tau_s, v_s \rangle$ from $\forall s \in Q$
     $\tau_{max} \leftarrow \max(\{\tau_s : r_i \text{ received } \langle \tau_s, v_s \rangle \text{ from } s\})$
18:     **return** $\{\langle \tau_s, v_s \rangle : \tau_s = \tau_{max} \wedge p \text{ received } \langle \tau_s, v_s \rangle \text{ from } s\}$
   end procedure

20: _____

Primitive Handlers at server s_i in configuration c:

22:    **Upon receive** (QUERY-TAG) **from** q
     send τ to q
24:    **end receive**

   Upon receive (QUERY) **from** q
26:     send $\langle \tau, v \rangle$ to q
   end receive

28:    **Upon receive** (WRITE, $\langle \tau_{in}, v_{in} \rangle$) **from** q
     if $\tau_{in} > \tau$ **then**
30:        $\langle \tau, v \rangle \leftarrow \langle \tau_{in}, v_{in} \rangle$
     send ACK to q
32:    **end receive**

Let us now examine if the primitives satisfy properties **C1** and **C2** of Property 1. We begin with a lemma that shows the monotonicity of the tags at each server.

LEMMA 35. *Let σ and σ' two states in an execution ξ such that σ appears before σ' in ξ. Then for any server $s \in \mathcal{S}$ it must hold that $s.tag|_\sigma \leq s.tag|_{\sigma'}$.*

Proof. According to the algorithm, a server s updates its local tag-value pairs when it receives a message with a higher tag. So if $s.tag|_\sigma = \tau$ then in a state σ' that appears after σ in ξ, $s.tag|_{\sigma'} \geq \tau$. □

In the following two lemmas we show that property **C1** is satisfied, that is if a put-data action completes, then any subsequent get-data and get-tag actions will discover a higher tag than the one propagated by that put-data action.

LEMMA 36. *Let ϕ be a c.put-data($\langle \tau, v \rangle$) action invoked by p_1 and γ be a c.get-tag() action invoked by p_2 in a configuration c, such that $\phi \rightarrow \gamma$ in an execution ξ of the algorithm. Then γ returns a tag $\tau_\gamma \geq \tau$.*

Proof. The lemma follows from the intersection property of quorums. In particular, during the c.put-data($\langle \tau, v \rangle$) action, p_1 sends the pair $\langle \tau, v \rangle$ to all the servers in $c.Servers$ and waits until all the servers in a quorum $Q_i \in c.Quorums$ reply. When those replies are received then the action completes.

During a c.get-data() action on the other hand, p_2 sends query messages to all the servers in $c.Servers$ and waits until all servers in a quorum $Q_j \in c.Quorums$ (not necessarily different than Q_i) reply. By definition $Q_i \cap Q_j \neq \emptyset$, thus any server $s \in Q_i \cap Q_j$ reply to both ϕ and γ actions. By Lemma 35 and since s received a tag τ, then s replies to p_2 with a tag $\tau_s \geq \tau$. Since γ returns the maximum tag it discovers then $\tau_\gamma \geq \tau_s$. Therefore $\tau_\gamma \geq \tau$ and this completes the proof. □

With similar arguments and given that each value is associated with a unique tag then we can show the following lemma.

LEMMA 37. *Let π be a c.put-data($\langle \tau, v \rangle$) action invoked by p_1 and ϕ be a c.get-data() action invoked by p_2 in a configuration c, such that $\pi \rightarrow \phi$ in an execution ξ of the algorithm. Then ϕ returns a tag-value $\langle \tau_\phi, v_\phi \rangle$ such that $\tau_\phi \geq \tau$.*

Finally we can now show that property **C2** also holds.

LEMMA 38. *If ϕ is a c.get-data() that returns $\langle \tau_\pi, v_\pi \rangle \in \mathcal{T} \times \mathcal{V}$, then there exists π such that π is a c.put-data($\langle \tau_\pi, v_\pi \rangle$) and $\phi \nrightarrow \pi$.*

Proof. This follows from the facts that (i) servers set their tag-value pair to a pair received by a put-data action, and (ii) a get-data action returns a tag-value pair that it received from a server. So if a c.get-data() operation ϕ returns a tag-value pair $\langle \tau_\pi, v_\pi \rangle$, there should be a server s that replied to that operation with $\langle \tau_\pi, v_\pi \rangle$, and s received $\langle \tau_\pi, v_\pi \rangle$ from some c.put-data($\langle \tau_\pi, v_\pi \rangle$) action, π. Thus, π can proceed or be concurrent with ϕ, and hence $\phi \nrightarrow \pi$. □

9 EXPERIMENTAL EVALUATION

The theoretical findings suggest that Ares is an algorithm to provide robustness and flexibility on shared memory implementations, without sacrificing strong consistency. In this section we present a *proof-of-concept* implementation of Ares and we run preliminary experiments to get better insight on the efficiency and adaptiveness of Ares. In particular, our experiments measure the latency of each read, write, and reconfig operations, and examine the persistence of consistency even when the service is reconfigured between configurations that add/remove servers and utilize different shared memory algorithms.

9.1 Experimental Testbed

We ran experiments on two different setups: (i) simulated locally on a single machine, and (ii) on a LAN. Both type of experiments run on *Emulab* [2], an emulated WAN environment testbed used for developing, debugging, and evaluating the systems. We used nodes with two 2.4 GHz 64-bit 8-Core E5-2630 "Haswell" processors, 64 GB RAM, with 1 GB and 10 GB NICs. In both setups we used an external implementation of Raft[45] consensus

algorithms, which was used for the service reconfiguration (line 16 of Alg. 2) and was deployed on top of small RPi devices. Small devices introduced further delays in the system, reducing the speed of reconfigurations and creating harsh conditions for longer periods in the service. The Python implementation of Raft used for consensus is PySyncObj [5]. Some modifications were done to allow the execution of Raft in the Ares environment. We built an HTTP API for the management of the Raft subsystem. A reconfigurer can propose a configuration at a particular index in the configuration sequence by sending a POST request to the url of each Raft node, and receives a response from the RAFT on which configuration is decided for that index.

Local Experimental Setup: The local setup was used to have access to a global synchronized clock (the clock of the local machine) in order to examine whether our algorithm preserves global ordering and hence atomicity even when using different algorithms between configurations. Therefore, all the instances are hosted on the same physical machine avoiding the skew between computer clocks in a distributed system. Furthermore, the use of one clock guarantees that when an event occurs after another, it will assign a later time.

Distributed Experimental Setup: The distributed experiments in Emulab enabled the examination of the performance of the algorithm in a close to real environment. For the deployment and remote execution of the experimental tasks on the Emulab, we used *Ansible Playbooks* [1]. All physical nodes were placed on a single LAN using a DropTail queue without delay or packet loss. Each physical machine runs one server or client process. This guarantees a fair communication delay between a client and a server node.

Node Types: In all experiments, we use four distinct types of nodes, writers, readers , reconfigurers and servers. Their main role is listed below:

- **writer** $w \in W \subseteq C$ **:** a client that sends write requests to all servers and waits for a quorum of the servers to reply
- **reader** $r \in R \subseteq C$**:** a client that sends read requests to servers and waits for a quorum of the servers to reply
- **reconfigurer** $g \in G \subseteq C$**:** a client that sends reconfiguration requests to servers and waits for a quorum of the servers to reply
- **server** $s \in S$**:** a server listens for read and write requests, it updates its object replica according to the atomic shared memory and replies to the process that originated the request.

Performance Metric: The metric for evaluating the algorithms is the operational latency. This includes both communication and computational delays. The operation latency is computed as the average of all clients' average operation latencies. For better estimations, each experiment in every scenario was repeated 6 times. In the graphs, we use error bars to illustrate the standard error of the mean (SEM) from the 6 repeated experiments.

9.2 Experimental Scenarios

In this section, we describe the scenarios we constructed and the settings for each of them. In our scenarios we constructed the DAP_s and used two different atomic storage algorithms in Ares: (i) the erasure coding based algorithm presented in Section 5, and (ii) the ABD algorithm (see Section 8.1).

Implementation of DAP_s: Clients initialize the appropriate configuration objects to handle any request. Notice that the client creates a configuration object when it is the first time that the client requests an operation or when doing a reconfiguration operation. Once the configuration object is initialized, it is stored on the client *cseq* and it is retrieved directly on any subsequent request from the client. Therefore, the DAP_s procedures are called from a configuration object. The asynchronous communication between components is achieved by using DEALER and ROUTER sockets, from the ZeroMQ Python library [6].

Erasure Coding: The type of erasure coding we use is (n,k)-Reed-Solomon code, which guarantees that any k of n coded fragments is enough to reassemble the original object. The parameter k is the number of encoded data fragments, n is the total number of servers and m is the number of parity fragments, i.e. $n - k$. A high number of k and consequently a small number of m means less redundancy with the system tolerating fewer failures. When

$k = 1$ we essentially converge to replication. In practice, the get-data and put-data functions from algorithm 5 integrate the standard Reed-Solomon implementation provided by liberasurecode from the PyEClib Python library [4].

Fixed Parameters: In all scenarios, the number of servers is fixed to 10. The number of writers and the value of delta are set to 5; delta being the maximum number of concurrent put-data operations. The parity value of the EC is set to 2 in order to minimize the redundancy, leading to 8 data servers and 2 parity servers. It is worth mentioning that the quorum size of the EC algorithm is $\lceil \frac{10+8}{2} \rceil = 9$, while the quorum size of ABD algorihtm is $\lfloor \frac{10}{2} \rfloor + 1 = 6$. In relation to the EC algorithm, we can conclude that the parameter k is directly proportional to the quorum size. But as the value of k and quorum size increase, the size of coded elements decreases.

Distributed Experiments: For the distributed experiments we use a *stochastic* invocation scheme in which readers and writers pick a random time uniformly distributed (discrete) between intervals to invoke their next operations. Respectively the intervals are $[1...rInt]$ and $[1..wInt]$, where $rInt, wInt = 2sec$. In total, each writer performs 60 writes and each reader 60 reads. The reconfigurer invokes its next operation every $15sec$ and performs a total of 6 reconfigurations. The intervals are set within these values in order to generate a continuous flow of operations and stress the concurrency in the system. Note that these values are not based on any real world scenario.

In particular, we present six types of scenarios:

- **File Size Scalability (Emulab):** The first scenario is made to evaluate how the read and write latencies are affected by the size of the shared object. There are two separated, one for each examined storage algorithm. The file size is doubled from 1 MB to 128 MB. The number of readers is fixed to 5, without any reconfigurers.
- **Reader Scalability (Emulab):** This scenario is constructed to compare the read and write latency of the system with two different storage algorithms, while the readers increase. In particular, we execute two separate runs, one for each storage algorithm. We use only one reconfigurer which requests recon operations that lead to the same shared memory emulation and server nodes. The size of the file used is 4 MB.
- **Changing Reconfigurations (Emulab):** In this scenario, we evaluate how the read and write latencies are affected when increasing the number of readers, while also changing the storage algorithm. We run two different runs which differ in the way the reconfigurer chooses the next storage algorithm: (i) the reconfigurer chooses randomly between the two storage algorithms, and (ii) the reconfigurer switches between the two storage algorithms. The size of the file used, in both scenarios, is 4 MB.
- **k Scalability (Emulab, EC only):** In this scenario, we examine the read and write latencies with different numbers of k (a parameter of Reed-Solomon). We increase the k of the EC algorithm from 1 to 9. The number of readers is fixed to 5, without any reconfigurers. The size of the file used is 4 MB.
- **Changing the number of Reads/Writes (Emulab):** In these scenarios, we examine the read and write latencies with different numbers of read and write operations respectively. We change the number of reads/writes that each reader/writer performs, from 10 to 60, increasing by 10. We calculate all possible pairs of writes and reads. The number of readers is fixed to 5. The reconfigurer switches between the two storage algorithms. The size of the file used is 4 MB
- **Consistency Persistence (Local):** In this scenario, we run multiple client operations in order to check if the data is consistent across servers. The number of readers is set to 5. The readers and writers invoke their next operations without any time delay, while the reconfigurer waits $15sec$ for the next invocation. We run two different scenarios which differ in the reconfigurations. In both scenarios, the reconfigurer switches between the two storage algorithms. In the second scenario, the reconfigurer changes concurrently the
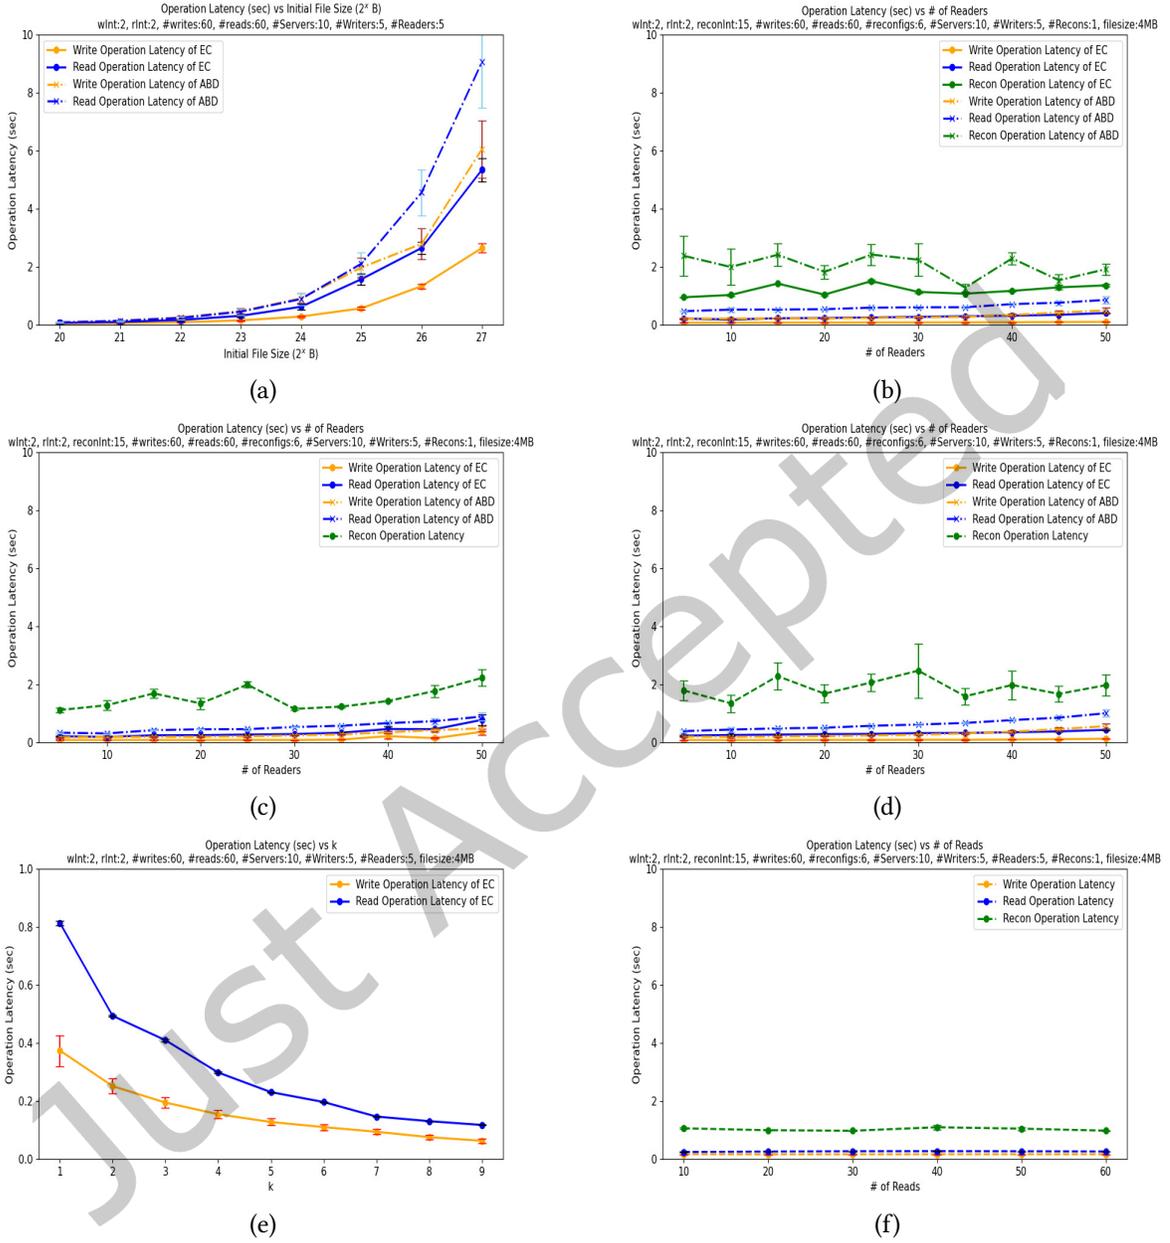
Fig. 3. Simulation results.

quorum of servers. In total, each writer performs 500 writes, each reader 500 reads and the reconfigurer 50 reconfigurations. The size of the file used is 4 MB.

9.3 Experimental Results

In this section, we present and explain the evaluation results of each scenario. As a general observation, the Ares algorithm with the EC storage provides data redundancy with a lower communicational and storage cost compared to the ABD storage that uses a strict replication technique.

File Size Scalability Results: Fig. 3(a) shows the results of the file size scalability experiments. The read and write latencies of both storage algorithms remain in low levels until 16 MB. In bigger sizes we observe the latancies of all operations to grow significantly. It is worth noting that the fragmentation applied by the EC algorithm, benefits its write operations which follow a slower increasing curve than the rest of the operations. From the rest the reads seem to suffer the worst delay hit, as they are engaged in more communication phases. Nevertheless, the larger messages sent by ABD result in slower read operations. We had noticed that EC has lower SEM values than ABD, which indicates that the calculated mean latencies of EC align very closely throughout the experiments. As EC breaks each file into smaller fragments, in combination with the fact that the variation is smaller when using smaller files in ABD, may lead to the conclusion that the file size has a significant impact on the error variation. To this end it appears that larger file sizes introduce higher variation on the delivery times of the file and hence higher statistical errors.

Reader Scalability Results: The results of reader scalability experiments can be found in Fig. 3(b). The read and write latencies of both algorithms remain almost unchanged, while the number of readers increases. This indicates that the system does not reach a state where it can not handle the concurrent read operations. Still, the reduced message size of read and write operations in EC keep their latencies lower than the corresponding latencies of ABD. On the other hand, the reconfiguration latency in both algorithms witnesses wild fluctuations between about 1 sec and 4 sec. This is probably due to the unstable connection in the external service which handles the reconfigurations. Notice that the number of readers does not have a great impact on the results we obtain in each experiment as the SEM error bars are small. The same goes for the next scenario where the number of readers changes while switching algorithms between reconfigurations.

Changing Reconfigurations Results: Fig. 3(c) illustrates the results of experiments with the random storage change. While, in Fig. 3(d), we can find the results of the experiments when the reconfigurer switches between storage algorithms. During both experiments, there are cases where a single read/write operation may access configurations that implement both ABD and EC algorithms, when concurrent with a recon operation. Thus, the latencies of such operations are accounted in both ABD and EC latencies. As we mentioned earlier, our choice of k minimizes the coded fragment size but introduces bigger quorums and thus larger communication overhead. As a result, in smaller file sizes, Ares may not benefit from the coding, bringing the delays of the two algorithms closer to each other. It is again obvious that the reconfiguration delays are higher than the delays of all other operations.

k Scalability Results: From Figs. 3(e) we can infer that when smaller k are used, the write and read latencies reach their highest values. In both cases, small k results in the generation of smaller number of data fragments and thus bigger sizes of the fragments and higher redundancy. For example we can see that for RS(10,8) and RS(10,7) we have the same size of quorum, equal to 9, whereas the latter has more redundant information. As a result, with a higher number of m (i.e. smaller k) we achieve higher levels of fault-tolerance, but that it would waste storage efficiency. The write latency seems to be less affected by the number of k since the encoding is considerably faster as it requires less computation. In conclusion, there appears to be a trade-off between operation latency and fault-tolerance in the system: the further increase of the k (and thus lower fault-tolerance) the smaller the latency of read/write operations. This experiment proves that the object size plays a significant role on the error variation. Notice that while k is small, and thus the object we send out is bigger, the error is higher. As k goes bigger and the fragments get smaller the SEM minimizes. This is an indication that communication of larger data over the wire may fluctuate the delivery times (as also seen in the file size scenario).

Changing the number of Reads/Writes Results: Fig. 3(f) shows a subset of the results of the experiments where the number of read operations changes and the number of write operations is fixed to 60. The experiments show that the total read/write latency (both EC and ABD) has very similar values for all the combinations of writes and reads, which indicates that the system performance is not affected by the number of reads and writes. This is expected since the number of participants is the same in all cases and by well-formedness (i.e., each participant invokes a single operation at a time) at most 10 operations will be concurrent in the execution at any given state. Higher concurrency can be captured by the scalability scenario. Note again that the read latency is higher than the write one, since the read operation actually transfers data twice: once to fetch the data from the servers, and once during the propagation phase.

Consistency Persistence Results: Though Ares protocol is probably strongly consistent, it is important to ensure that our implementation is correct. Validating strong consistency of an execution requires precise clock synchronization across all processes, so that one can track operations with respect to a global time. This is impossible to achieve in a distributed system where clock drift is inevitable. To circumvent this, we deploy all the processes in a single beefy machine so that every process observes the same clock running in the same physical machine.

Our checker gathers data regarding an execution, and this data includes start and end times of all the operations, as well as other parameters like logical timestamps used by the protocol. The checker logic is based on the conditions appearing in Lemma 13.16 [38], which provide a set of sufficient conditions for guaranteeing strong consistency. The checker validates strong consistency property for every atomic object individually for the execution under consideration. Note that consistency holds despite the existence of concurrent read/write and reconfiguration operations that may add/remove servers and switch the storage algorithm in the system.

10 CONCLUSIONS

We presented an algorithmic framework suitable for reconfigurable, erasure code-based atomic memory service in asynchronous, message-passing environments. In particular, we provide a new modular framework, called Ares, which abstracts the implementation of the underlying shared memory within a set of DAPs with specific correctness properties. Using these structures, Ares may implement a large class of atomic shared algorithms (those that can be expressed using the proposed DAPs) allowing any such algorithm to work on a reconfigurable environment. A set of erasure-coded-based atomic memory algorithms are included in this class. To demonstrate the use of our framework, we provided a new two-round erasure code-based algorithm that has near optimal storage cost, implemented in terms of the proposed DAPs. Such implementation gave rise to the first (to our knowledge) reconfigurable erasure-coded atomic shared memory object. We provided a proof-of-concept implementation of our framework and obtained initial experimental results proving the feasibility of the presented approach, demonstrating its correctness and comparing its performance with traditional approaches.

Ares is designed to address the real-world problem of system migration from a replicated system to a system that uses erasure codes and vice-versa. Ares can also enable replacing of failed nodes with new non-failed nodes. It's key difference with existing state of the art systems is that it can perform such reconfigurations with relatively minimal interruption of service unlike current implementations that would block ongoing operations for reconfiguration. We anticipate that Ares will be very useful for workloads that are prone to fast changes in properties, and have stringent constraints on the latencies. For such workloads, Ares enables the system to adapt itself to the changes in the workload in an agile manner - utilizing the full flexibility that EC brings to the system - without causing latency constraint or consistency violations due to interruptions.

Our main goal was to establish that non-blocking reconfiguration is feasible and compatible with EC based atomic data storage.Our experimental study is designed as a *proof-of-concept* prototype to verify the correctness properties we have developed, and show some benefits. It must be emphasized that a full-fledged system study of

our algorithms - albeit an interesting area of future work that is motivated by our paper - is outside our current scope. In particular, although our study provides some initial hints, we anticipate such a future study would examine the following questions in more detail:

- Real-world applications that would indeed benefit from our design
- Workloads generated from these real-world applications to test our algorithms, and competing ones
- The synergies between our reconfiguration algorithm and existing failure detection and recovery mechanisms.
- Adding efficient repair and reconfiguration using regenerating codes

REFERENCES

[1] Ansible. https://www.ansible.com/overview/how-ansible-works. (????).
[2] Emulab network testbed. https://www.emulab.net/. (????).
[3] Intel Storage Acceleration Library (Open Source Version). (????). https://goo.gl/zkVl4N.
[4] PyEClib. (????). https://github.com/openstack/pyeclib
[5] PySyncObj. (????). https://github.com/bakwc/PySyncObj
[6] ZeroMQ. (????). https://zeromq.org
[7] M. Abebe, K. Daudjee, B. Glasbergen, and Y. Tian. 2018. EC-Store: Bridging the Gap between Storage and Latency in Distributed Erasure Coded Systems. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*. 255–266. DOI:https://doi.org/10.1109/ICDCS.2018.00034
[8] Marcos Kawazoe Aguilera, Idit Keidar, Dahlia Malkhi, and Alexander Shraer. 2009. Dynamic atomic storage without consensus. In *Proceedings of the 28th ACM symposium on Principles of distributed computing (PODC '09)*. ACM, New York, NY, USA, 17–25.
[9] Marcos K. Aguilera, Idit Keidary, Dahlia Malkhi, Jean-Philippe Martin, and Alexander Shraery. 2010. Reconfiguring Replicated Atomic Storage: A Tutorial. *Bulletin of the EATCS* 102 (2010), 84–081.
[10] Antonio Fernández Anta, Nicolas Nicolaou, and Alexandru Popa. 2015. Making "Fast" Atomic Operations Computationally Tractable. In *International Conference on Principles Of Distributed Systems (OPODIS'15)*.
[11] H. Attiya, A. Bar-Noy, and D. Dolev. 1996. Sharing Memory Robustly in Message Passing Systems. *J. ACM* 42(1) (1996), 124–142.
[12] Dorian Burihabwa, Pascal Felber, Hugues Mercier, and Valerio Schiavoni. 2016. A performance evaluation of erasure coding libraries for cloud-based data stores. In *Distributed Applications and Interoperable Systems*. Springer, 160–173.
[13] Christian Cachin and Stefano Tessaro. 2006. Optimal Resilience for Erasure-Coded Byzantine Distributed Storage. In *Dependable Systems and Networks, International Conference on*. IEEE Computer Society, Los Alamitos, CA, USA, 115–124.
[14] V. R. Cadambe, N. Lynch, M. Médard, and P. Musial. 2014. A Coded Shared Atomic Memory Algorithm for Message Passing Architectures. In *Network Computing and Applications (NCA), 2014 IEEE 13th International Symposium on*. 253–260. DOI:https://doi.org/10.1109/NCA.2014.44
[15] Viveck R. Cadambe, Nancy A. Lynch, Muriel Médard, and Peter M. Musial. 2017. A coded shared atomic memory algorithm for message passing architectures. *Distributed Computing* 30, 1 (2017), 49–73.
[16] Yu Lin Chen Chen, Shuai Mu, and Jinyang Li. 2017. Giza: Erasure Coding Objects across Global Data Centers. In *Proceedings of the 2017 USENIX Annual Technical Conference (USENIX ATC '17)*. 539–551.
[17] Gregory Chockler, Seth Gilbert, Vincent Gramoli, Peter M. Musial, and Alexander A. Shvartsman. 2009. Reconfigurable distributed storage for dynamic networks. *J. Parallel and Distrib. Comput.* 69, 1 (2009), 100–116.
[18] Gregory Chockler and Dahlia Malkhi. 2005. Active Disk Paxos with infinitely many processes. *Distributed Computing* 18, 1 (2005), 73–84.
[19] Dan Dobre, Ghassan O. Karame, Wenting Li, Matthias Majuntke, Neeraj Suri, and Marko Vukolić. 2019. Proofs of Writing for Robust Storage. *IEEE Transactions on Parallel and Distributed Systems* 30, 11 (2019), 2547–2566. DOI:https://doi.org/10.1109/TPDS.2019.2919285
[20] Partha Dutta, Rachid Guerraoui, and Ron R. Levy. 2008. Optimistic Erasure-Coded Distributed Storage. In *DISC '08: Proceedings of the 22nd international symposium on Distributed Computing*. Springer-Verlag, Berlin, Heidelberg, 182–196.
[21] Partha Dutta, Rachid Guerraoui, Ron R. Levy, and Arindam Chakraborty. 2004. How fast can a distributed atomic read be?. In *Proceedings of the 23rd ACM symposium on Principles of Distributed Computing (PODC)*. 236–245.
[22] Rui Fan and Nancy Lynch. 2003. Efficient Replication of Large Data Objects. In *Distributed algorithms (Lecture Notes in Computer Science)*, Faith Ellen Fich (Ed.), Vol. 2848. 75–91.
[23] Antonio Fernández Anta, Theophanis Hadjistasi, and Nicolas Nicolaou. 2016. Computationally Light "Multi-Speed" Atomic Memory. In *International Conference on Principles Of Distributed Systems (OPODIS'16)*.

[24] Eli Gafni and Dahlia Malkhi. 2015. Elastic Configuration Maintenance via a Parsimonious Speculating Snapshot Solution. In *International Symposium on Distributed Computing*. Springer, 140–153.

[25] Chryssis Georgiou, Nicolas C. Nicolaou, and Alexander A. Shvartsman. 2008. On the Robustness of (Semi) Fast Quorum-Based Implementations of Atomic Shared Memory. In *DISC '08: Proceedings of the 22nd international symposium on Distributed Computing*. Springer-Verlag, Berlin, Heidelberg, 289–304.

[26] Chryssis Georgiou, Nicolas C. Nicolaou, and Alexander A. Shvartsman. 2009. Fault-tolerant semifast implementations of atomic read/write registers. *J. Parallel and Distrib. Comput.* 69, 1 (2009), 62–79.

[27] Seth Gilbert. 2003. *RAMBO II: Rapidly Reconfigurable Atomic Memory for Dynamic Networks.* Master's thesis. MIT.

[28] S. Gilbert, N. Lynch, and A.A. Shvartsman. 2003. RAMBO II: Rapidly reconfigurable atomic memory for dynamic networks. In *Proceedings of International Conference on Dependable Systems and Networks (DSN)*. 259–268.

[29] Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems* 12, 3 (1990), 463–492.

[30] Jianzhong Huang, Xianhai Liang, Xiao Qin, Ping Xie, and Changsheng Xie. 2015. Scale-RS: An Efficient Scaling Scheme for RS-Coded Storage Clusters. *IEEE Transactions on Parallel and Distributed Systems* 26, 6 (2015), 1704–1717. DOI:https://doi.org/10.1109/TPDS.2014.2326156

[31] W. C. Huffman and V. Pless. 2003. *Fundamentals of error-correcting codes.* Cambridge university press.

[32] Leander Jehl, Roman Vitenberg, and Hein Meling. 2015. Smartmerge: A new approach to reconfiguration for atomic storage. In *International Symposium on Distributed Computing*. Springer, 154–169.

[33] Gauri Joshi, Emina Soljanin, and Gregory Wornell. 2017. Efficient redundancy techniques for latency reduction in cloud systems. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS)* 2, 2 (2017), 12.

[34] K. M. Konwar, N. Prakash, E. Kantor, N. Lynch, M. Médard, and A. A. Schwarzmann. 2016. Storage-Optimized Data-Atomic Algorithms for Handling Erasures and Errors in Distributed Storage Systems. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 720–729.

[35] Kishori M Konwar, N Prakash, Nancy Lynch, and Muriel Médard. 2016. RADON: Repairable atomic data object in networks. In *The International Conference on Distributed Systems (OPODIS)*.

[36] Avinash Lakshman and Prashant Malik. 2010. Cassandra: a decentralized structured storage system. *SIGOPS Oper. Syst. Rev.* 44, 2 (April 2010), 35–40. DOI:https://doi.org/10.1145/1773912.1773922

[37] Leslie Lamport. 1998. The part-time parliament. *ACM Transactions on Computer Systems* 16, 2 (1998), 133–169. DOI:https://doi.org/10.1145/279227.279229

[38] N.A. Lynch. 1996. *Distributed Algorithms.* Morgan Kaufmann Publishers.

[39] N. Lynch and A.A. Shvartsman. 2002. RAMBO: A reconfigurable atomic memory service for dynamic networks. In *Proceedings of 16th International Symposium on Distributed Computing (DISC)*. 173–190.

[40] Nancy A. Lynch and Alexander A. Shvartsman. 1997. Robust Emulation of Shared Memory Using Dynamic Quorum-Acknowledged Broadcasts. In *Proceedings of Symposium on Fault-Tolerant Computing*. 272–281.

[41] Ellis Michael, Dan R. K. Ports, Naveen Kr. Sharma, and Adriana Szekeres. 2017. Recovering Shared Objects Without Stable Storage. In *31st International Symposium on Distributed Computing (DISC 2017) (Leibniz International Proceedings in Informatics (LIPIcs))*, Andréa W. Richa (Ed.), Vol. 91. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 36:1–36:16. DOI:https://doi.org/10.4230/LIPIcs.DISC.2017.36

[42] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).

[43] Nicolas Nicolaou, Viveck Cadambe, Kishori Konwar, N Prakash, Nancy Lynch, and Muriel Médard. 2018. ARES: Adaptive, Reconfigurable, Erasure coded, atomic Storage. *CoRR* abs/1805.03727 (2018).

[44] Nicolas Nicolaou, Viveck Cadambe, N. Prakash, Kishori Konwar, Muriel Medard, and Nancy Lynch. 2019. ARES: Adaptive, Reconfigurable, Erasure Coded, Atomic Storage. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. 2195–2205. DOI:https://doi.org/10.1109/ICDCS.2019.00216

[45] Diego Ongaro and John Ousterhout. 2014. In Search of an Understandable Consensus Algorithm. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference (USENIX ATC'14)*. USENIX Association, Berkeley, CA, USA, 305–320.

[46] KV Rashmi, Mosharaf Chowdhury, Jack Kosaian, Ion Stoica, and Kannan Ramchandran. 2016. EC-Cache: Load-Balanced, Low-Latency Cluster Caching with Online Erasure Coding.. In *OSDI*. 401–417.

[47] Alexander Shraer, Jean-Philippe Martin, Dahlia Malkhi, and Idit Keidar. 2010. Data-centric reconfiguration with network-attached disks. In *Proceedings of the 4th Int'l Workshop on Large Scale Distributed Sys. and Middleware (LADIS '10)*. 22–26.

[48] Alexander Spiegelman, Idit Keidar, and Dahlia Malkhi. 2017. Dynamic Reconfiguration: Abstraction and Optimal Asynchronous Solution. In *31st International Symposium on Distributed Computing (DISC 2017)*, Vol. 91. 40:1–40:15.

[49] Shuang Wang, Jianzhong Huang, Xiao Qin, Qiang Cao, and Changsheng Xie. 2017. WPS: A Workload-Aware Placement Scheme for Erasure-Coded In-Memory Stores. In *Networking, Architecture, and Storage (NAS), 2017 International Conference on*. IEEE, 1–10.

[50] Chentao Wu and Xubin He. 2012. GSR: A Global Stripe-Based Redistribution Approach to Accelerate RAID-5 Scaling. In *2012 41st International Conference on Parallel Processing*. 460–469. DOI:https://doi.org/10.1109/ICPP.2012.32

[51] Chentao Wu and Xubin He. 2013. A Flexible Framework to Enhance RAID-6 Scalability via Exploiting the Similarities among MDS Codes. In *2013 42nd International Conference on Parallel Processing*. 542–551. DOI:https://doi.org/10.1109/ICPP.2013.68

[52] Yu Xiang, Tian Lan, Vaneet Aggarwal, and Yih-Farn R Chen. 2015. Multi-tenant latency optimization in erasure-coded storage with differentiated services. In *2015 IEEE 35th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 790–791.

[53] Yu Xiang, Tian Lan, Vaneet Aggarwal, Yih-Farn R Chen, Yu Xiang, Tian Lan, Vaneet Aggarwal, and Yih-Farn R Chen. 2016. Joint latency and cost optimization for erasure-coded data center storage. *IEEE/ACM Transactions on Networking (TON)* 24, 4 (2016), 2443–2457.

[54] Yinghao Yu, Renfei Huang, Wei Wang, Jun Zhang, and Khaled Ben Letaief. 2018. SP-cache: load-balanced, redundancy-free cluster caching with selective partition. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. IEEE Press, 1–13.

[55] Guangyan Zhang, Keqin Li, Jingzhe Wang, and Weimin Zheng. 2015. Accelerate RDP RAID-6 Scaling by Reducing Disk I/Os and XOR Operations. *IEEE Trans. Comput.* 64, 1 (2015), 32–44. DOI:https://doi.org/10.1109/TC.2013.210

[56] Heng Zhang, Mingkai Dong, and Haibo Chen. 2016. Efficient and Available In-memory KV-Store with Hybrid Erasure Coding and Replication. In *14th USENIX Conference on File and Storage Technologies (FAST 16)*. USENIX Association, Santa Clara, CA, 167–180.

[57] Xiaoyang Zhang, Yuchong Hu, Patrick P. C. Lee, and Pan Zhou. 2018. Toward Optimal Storage Scaling via Network Coding: From Theory to Practice. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 1808–1816. DOI:https://doi.org/10.1109/INFOCOM.2018.8485961

[58] P. Zhou, J. Huang, X. Qin, and C. Xie. 2018. PaRS: A Popularity-aware Redundancy Scheme for In-Memory Stores. *IEEE Trans. Comput.* (2018), 1–1. DOI:https://doi.org/10.1109/TC.2018.2876827