

SNOW Revisited: Understanding When Ideal READ Transactions Are Possible

Kishori M. Konwar¹, Wyatt Lloyd², Haonan Lu², and Nancy Lynch³

¹RLE, MIT

²Department of Computer Science, Princeton University

³CSAIL, MIT

kishori@mit.edu, wlloyd@cs.princeton.edu, haonanl@cs.princeton.edu, lynch@csail.mit.edu

May 25, 2021

Abstract

READ transactions that read data distributed across servers dominate the workloads of real-world distributed storage systems. The SNOW Theorem [15] stated that ideal READ transactions that have optimal latency and the strongest guarantees—i.e., “SNOW” READ transactions—are impossible in one specific setting that requires three or more clients: at least two readers and one writer. However, it left many open questions.

We close all of these open questions with new impossibility results and new algorithms. First, we prove rigorously the result from [15] saying that it is impossible to have a READ transactions system that satisfies SNOW properties with three or more clients. The insight we gained from this proof led to teasing out the implicit assumptions that are required to state the results and also, resolving the open question regarding the possibility of SNOW with two clients. We show that it is possible to design an algorithm, where SNOW is possible in a multi-writer, single-reader (MWSR) setting when a client can send messages to other clients; on the other hand, we prove it is impossible to implement SNOW in a multi-writer, single-reader (MWSR) setting—which is more general than the two-client setting—when client-to-client communication is disallowed. We also correct the previous claim in [15] that incorrectly identified one existing system, Eiger [14], as supporting the strongest guarantees (SW) and whose read-only transactions had bounded latency. Thus, there were no previous algorithms that provided the strongest guarantees and had bounded latency. Finally, we introduce the first two algorithms to provide the strongest guarantees with bounded latency.

1 Introduction

Today’s web services are built on *distributed storage systems* that provide fault tolerant and scalable access to data. Distributed storage systems scale their capacity and throughput by *sharding* (i.e., partitioning) data across many machines within a datacenter, i.e., each machine stores a subset of the data. They also *geo-replicate* the data across several geographically dispersed datacenters to tolerate failures and to increase their proximity to users.

Distributed storage systems abstract away the complexities of sharding and replication from application code by providing guarantees for accesses to data. These *guarantees* include consistency and transactions. *Consistency* controls the values of data that accesses may observe and *transactions* dictate what accesses may be grouped together. Stronger guarantees provide an abstraction closer to a single-threaded environment, greatly simplifying application code. Ensuring the guarantees hold, however, often comes with worse performance. Therefore, the tradeoff between performance and guarantees lies at the heart of designing such systems.

The performance-guarantee tradeoffs that result from replication have been well-studied with several well-known impossibility results [1, 7, 8, 13]. For instance, the CAP Theorem [8] proves that system designers must choose either availability during network partitions (performance) or strong consistency across replicas (guarantee). However, little prior work exists on what performance-guarantee tradeoffs result from sharding.

Understanding the performance-guarantee tradeoff due to sharding is important because user requests are typically handled across many shards but within a single nearby datacenter (replica). This is particularly true for the reads needed to handle a user request, which are what dominate real-world workloads: Facebook reported 500 reads for every write in their TAO system [4] and Google reported three orders of magnitude more reads than general transactions for their F1 database that runs on their Spanner system [6]. In this work, we focus on clarifying the performance-guarantee tradeoff for reads that results from sharding. Distributed storage systems group read requests (that each individually accesses a separate shard) into *READ transactions* that together return a consistent, cross-shard view of the system. Whether a view is consistent is determined by the consistency model a system provides. The ideal READ transactions would have the strongest guarantees: They would provide strict serializability [17], the strongest consistency model, and they could be used in a system that also includes *WRITE transactions* that group write requests (each to a separate shard) together. The alternative to the latter property are READ transactions that can only be used in systems that have non-transactional, *simple writes*.

The ideal READ transactions would also provide the best performance. In particular, they would provide the lowest possible latency because the prevalence of reads makes them dominate the user response times that are aggressively optimized by web services [5, 12, 18]. The *optimal latency* for a READ transaction is to match the latency of non-transactional, *simple reads*: complete in a single round trip of non-blocking parallel requests to the shards that return only the requested data [15].

1.1 Previous Results and Open Questions

The SNOW Theorem was the first result in the sharding dimension that is relevant to READ transactions [15]. The SNOW Theorem is an impossibility result that proves no READ transaction can provide **S**trict serializability with **N**on-blocking client-server communication that completes with **O**ne response, with only one version of the data, per read in a system with concurrent **W**RITE transactions (§2.1). It shows there is a fundamental tradeoff between the latency and guarantees of READ transactions that system designers must grapple with, they must pick either the strongest guarantees (S and W) or optimal latency (N and O).

SNOW is trivially possible in systems with a single client or a single server because the single entity naturally serializes all transactions. The SNOW Theorem shows SNOW is impossible in systems with at least three clients and at least two servers. It explicitly leaves open the question of the possibility of SNOW in a system with two clients. In addition, the model used in the prior

work implicitly leaves open several questions. It assumed the three clients were a single writer and multiple readers (SWMR). This leaves open the possibility of SNOW with multiple writers and a single reader (MWSR). The SNOW Theorem also implicitly assumed that clients do not directly exchange messages and that write operations in such a system must eventually complete. This also leaves open the question of whether allowing or disallowing client-to-client (C2C) communication has any impact on the feasibility of READ transactions with SNOW properties.

In this work, the new impossibility results are philosophically similar to other impossibility results—such as FLP [7] and CAP [3, 8]—in that they help system designers avoid wasting effort in trying to achieve the impossible. That is, the SNOW Theorem identifies a boundary in the design space of READ transactions, beyond which no algorithms can possibly exist. By revisiting SNOW, our work makes this boundary more precise.

1.2 Our Contributions

Our work builds on the SNOW Theorem to clarify this fundamental tradeoff by providing a thorough proof for the SNOW Theorem [15] and also, answer the open questions mentioned above. First, we formally state the SNOW properties of executions using the I/O automata framework [16] with the additional requirement, for the W property, that any WRITE must eventually complete (§2). Next, we identify and prove some basic results, required in the proofs, for transforming one valid execution to another possible and safe execution in a READ transaction system (§3). Then we present a new, rigorous proof of the impossibility of SNOW with three or more clients even when client-to-client (C2C) communication is allowed (§4). Next, we show that the feasibility of implementing an algorithm with SNOW in MWSR (which also includes the two-client system model) depends on whether C2C communication is allowed: when it is not allowed, SNOW is impossible (§5.1); and when it is allowed, SNOW is possible (§5.2) for any MWSR setting.

Prior to this work, the Eiger [14] algorithm was previously believed to be the only algorithm that provided a bounded number of non-blocking rounds [14] and guaranteed strict serializability. Next, we show this claim is not true by showing that not all execution of Eiger is strictly serializable (§6).

Next, after realizing the limits posed by the SNOW Theorem, we ask ourselves whether it is possible to construct READ transaction algorithms with no C2C communication, as in most practical systems, where one of the SNOW properties is relaxed. One obvious candidate property is the “O” property, where one of the two restrictions (i.e., “one-round” of communication and “one-version” of data) can be relaxed. We provide two algorithms for the *multiple-writers multi-reader* (MWMR) setting: the first algorithm B guarantees SNW and the “one-version” property and completes READ transactions in two rounds (§8); the second algorithm, C , guarantees SNW and the “one-round” property but returns up to as many versions of the data as there are concurrent WRITE transactions (§9). Thereby, making these READ transactions algorithms with a bounded number of non-blocking rounds and guarantees strict serializability. Due to space limitations, we omit most of the proofs and present them in an extended version in arXiv [11].

2 Transactions Processing System

Web services typically have two tiers of machines within a datacenter: a stateless frontend tier and a stateful storage tier. The frontends handle user requests by executing application logic that

Setting	Client-to-Client?	
	Yes	No
2 clients	✓	×
MWSR	✓	×
≥ 3 clients	×	(×)

(a) Is SNOW possible?

Versions	Rounds		
	1	2	∞
1	(×)	✓	(✓)
$ W $	✓		

(b) Bounded SNW algorithms.

Figure 1: A summary of our new results. Previous results are marked in parentheses. × indicates we have proved that such READ transactions are impossible. ✓ indicates we have described such a new READ transaction algorithm. $|W|$ is the number of concurrent WRITE transactions.

generates sub-requests to read/write data in the storage tier that shards (or splits) data across many machines. We refer to the front-ends as the *clients*, the storage machines as the *servers* and the stored data items as *objects*, to match common terminology. While web services are typically geo-replicated, we focus on sharding within a datacenter because the reads that dominate their workloads are handled within a single datacenter.

We consider a transaction processing system that comprises a set of read/write objects $\mathcal{O}$, where each object $o \in \mathcal{O}$ is maintained by a separate server process, and also another set of processes, we refer to as *clients*, that can initiate transactions, after the previous ones, if any, have completed. The system allows two types of transaction: READ transaction, a group of read requests for the values stored in some subset of objects in $\mathcal{O}$; and WRITE transaction, a group of write requests intending to update the values stored in some subset of objects $\mathcal{O}$. A read-client executes only READ transactions, while write-client executes only WRITE transactions; no client executes both types of transaction.

A typical READ transaction, we denote as $R(o_{i_1}, o_{i_2}, \dots, o_{i_q})$ or in short by R , consists a set of individual read requests $read(o_{i_1})$, $read(o_{i_2})$ and $read(o_{i_q})$ to read values in objects $o_{i_1}, o_{i_2}, \dots, o_{i_q}$, respectively. $read(o)$ denotes a read that intends to read the value of object o . A typical WRITE, denoted as $W((o_{i_1}, v_{i_1}), (o_{i_2}, v_{i_2}), \dots, (o_{i_p}, v_{i_p}))$ or in short as W , consists of a set of which requests to update the values of objects $o_{i_1}, o_{i_2}, \dots, o_{i_p}$ with $v_{i_1}, v_{i_2}, \dots, v_{i_p}$, that are values from the domains $V_{i_1}, V_{i_2}, \dots, V_{i_p}$, respectively.

A read (or write) client initiates a READ (or WRITE) transaction with an invocation step $INV(R)$ (or $INV(W)$), then it carries out the read or write operations in the transaction; and eventually completes the transaction with a $RESP(R)$ (or $RESP(W)$). After the completion of the reads or writes in a transaction the client responds, in the case of R , with the values of objects; and, in the case of W , an *ok* status, to the external client.

We assume that the network channels are reliable but asynchronous, i.e., any message sent by a process will eventually arrive at its destination uncorrupted. We assume local computations are asynchronous, i.e., local computations at various processes proceed at arbitrary and unpredictable speeds. When a client receives a transaction request, usually from an external client, such as an user’s device, it executes the transaction, denote by R or W , and finally, responds to the external client with the results.

We model a distributed algorithm using the I/O automata modeling framework (see [16] for a detailed account). In the rest of this paper, for any execution of an automaton $\mathcal{A}$, $\sigma_0, a_1, \dots, a_k, \sigma_k \dots$, where σ ’s and a ’s are states and actions, we use the notation $a_1, \dots, a_k \dots$ which shows only the

actions to simplify notation. We use the notation $prefix(\alpha, a)$ to refer to the finite prefix of any execution α ending with action a such that a occurs within α . In our model, an individual read, such as $read(o)$, in some read transaction R initiated by some read client r consists of the following sequence of actions: after $INV(R)$ at r , r sends a message m (requesting the value stored in o) to a server s via action $send(m)_{r,s}$. When s receives m via action $recv(m)_{r,s}$ it sends the value v (stored in o) to r via action $send(v)_{s,r}$. Then $read(o)$ completes as soon as r receives v_j via action $recv(v)_{s,r}$; R completes with $RESP(R)$ after all the reads in it are complete.

2.1 SNOW Properties

In this subsection, we define the SNOW properties for a transaction processing system. Namely, we require that any fair execution of the system satisfies the following four properties: (i) *Strict serializability* (S), which means there is a total ordering of the transactions such that all transactions in the resulting execution appear to be processed by a single machine one at a time; (ii) *Non-blocking operations* (N), which means that the servers respond immediately to the read requests of a READ transaction without waiting for any input from other processes; (iii) *One response per read* (O), which requires that any read operation consists of one round trip of communication with a server, and also, that the server responds with a message that contains exactly one version of the object value; and (iv) *WRITE transactions that conflict* (W) implies the existence of concurrent WRITE transactions that update the data objects while READ transactions are in progress reading the same objects. Below we describe the individual properties of the SNOW properties in more detail.

Strict serializability (S). By *strict serializability* (for a formal definition please see [10]), we mean each WRITE or READ transaction appears to the clients to be executed atomically, at some point in an execution between the invocation and response events.

Next, we describe the non-blocking and one-response properties. Both are defined as properties of read operations to an individual object. For the purpose of elucidation, we consider an execution α of a transaction processing system T that has a set of objects $\mathcal{O}$, where there is a READ transaction $R(o_{i_1}, o_{i_2}, \dots, o_{i_q})$, in short R , invoked at some reader r , such that R contains a read $read(o_j)$ for some $o_j \in \mathcal{O}$ maintained at server s_j .

Non-blocking reads (N). The *non-blocking* property means that if r , a reader-client, sends any message to any s_i (s_i manages object o_i) during the transaction then s_i can respond to r without waiting for any external input event, such as the arrival of messages, any mutex operations, time, etc. This property ensures that READ transactions are delayed only due to delay in message delivery between r and s_i . We define this property formally as follows.

Definition 2.1 (Non-blocking read (N)). *Suppose in α , following the action $INV(R)$, the actions $recv(m_j^r)_{r,s_j}$ and $send(v_j)_{s_j,r}$ corresponding to $read(o_j)$, occurs at s_j . Then there exists an execution α' of T such that*

- (i) *The execution fragments $prefix(\alpha, recv(m_j^r)_{r,s_j})$ and $prefix(\alpha', recv(m_j^r)_{r,s_j})$ are identical, where $prefix(\alpha, a)$ is the prefix of α ending with a .*
- (ii) *In α' the action $send(v_j)_{s_j,r}$ at s_j occurs after $recv(m_j^r)_{r,s_j}$ without any input action in between.*

One-response per read (O). The *one-response* property requires that each read operation, $read(o_i)$, during any READ transaction, completes successfully in one round of client-to-server communication and the *one-version* states that exactly one version of the value is sent by server

s_i , that manages o_i , to r . *One-round* consists of a read request from the client initiating the read operation to the server and the response containing value sent by the server.

Definition 2.2 (One response per read (O)). *Suppose in α , the action $INV(R)$ occur at r then in α there exists exactly a pair of actions $recv(m_j^r)_{r,s_j}$ and $send(v_j)_{s_j,r}$, corresponding to R , occur at s_j , such that v_j is the object value of o_j .*

If the reads, of some READ transaction, of a transaction processing system respect the *non-blocking* and *one-response* properties then each read includes one-round trip from client to server, where the server returns only the requested value as soon as it receives the request. It is worth noting that the READ transaction can complete only after all the *read*($\cdot$)s in it complete.

Definition 2.3 (Conflicting writes (W)). *Suppose in α , the action $INV(W)$, the actions occur at a write client w then there is an action $RESP(W)$ in α that appears after $INV(W)$.*

WRITE transactions that conflict (W). The *conflicting writes* property states that READ transactions complete even in the presence of concurrent WRITE transactions, where the write operations might update some objects that are also being read by read operations in READ transaction. This shows that READ transactions can be invoked at any point, even in the presence of ongoing WRITE transactions. Note that the liveness of any WRITE transaction is not implied by any of the SNOW properties; however, for useful practical systems the WRITE transactions must eventually complete. Therefore, we assume that every WRITE transaction eventually completes via the *RESP* event, and think of this constraint as a part of the W property.

The SNOW Theorem. Consider a transaction processing system with an asynchronous network where a set $\mathcal{O}$ of objects are maintained by individual server processes, with at least one write client and at least two read clients. Then the SNOW Theorem [15] can be stated as follows.

“For any transaction processing system in an asynchronous setting, with at least one writer and two reader clients, and at least two sharded objects, it is impossible to have an algorithm such that all of its executions guarantee the SNOW properties.”

3 Technical Preliminaries

In this section, we present some useful preliminary results and ideas that we will later use to prove the impossibility results. We assume a simple system with two servers, s_x and s_y , denote the stored values as x and y , respectively, and either two or three clients. One of the clients is a writer w , which initiates only WRITE transactions. One or two of the clients are readers, r_1 and r_2 , which initiate only READ transactions.

Client-to-client (C2C) communication. We consider two types of settings pertinent to communication among the clients: (i) allow C2C communication, where a client can send a message to any other client, and (ii) disallow C2C communications, where a client cannot send any message directly to another client in the system.

Servers s_x and s_y store values for objects o_x and o_y , respectively. the initial values of o_x and o_y are x_0 and y_0 , respectively. Because there is one object on each server the server and object identifiers are often used interchangeably to remove redundancy. For instance, we simply say that s_x returns x_0 to the client that initiated the transaction, which means that s_x returns the value x_0 of object o_x at the end of the READ transaction.

Our proofs often use a special type of execution fragment, named non-blocking fragments, that represent the READ transaction algorithm is non-blocking and returns one version of each object. The one-round property is captured by allowing only one non-blocking fragment on each server for a READ transaction. Our proof strategy plays non-blocking fragments against the requirements of strict serializability and write isolation under the freedom of network asynchrony. We explain non-blocking fragments and helper notations in the context of execution α , of the system described as above, as follows:

1. *Non-blocking fragments.* For a READ transaction R_i by reader r_i , $i \in \{1, 2\}$, suppose there is a execution fragment that starts with $recv(m_j^r)_{r_i, s_j}$ and ends with $send(v_j)_{s_j, r_i}$, both of which occur at s_j . Moreover, suppose there is no other input action at s_j in this fragment. Then we call this execution fragment a *non-blocking fragment* for R_i at s_j and denote by $F_{i,j}(\alpha)^{(v_j)}$, $j \in \{x, y\}$ (Fig. 2). When the context is clear, we omit the first subscript of F . For instance, for a READ transaction R , $F_x(\alpha)^{(x_0)}$ denotes the non-blocking fragment of R on s_x .
2. Suppose READ transaction R_i completes in α . Consider the execution fragment in α between the event $INV(R_i)$ and whichever of the events $send(m_y^r)_{r_i, s_y}$ and $send(m_x^r)_{r_i, s_x}$ that occurs later. If all the actions in this fragment occur at r_i , then we denote this fragment as $I_i(\alpha)$ (Fig. 2).
3. Suppose READ transaction R_i completes in α . Consider the execution fragments in α that occurs between the later of the events $recv(x)_{s_x, r_i}$ or $recv(y)_{s_y, r_i}$, i.e., at the point in α when r_i receives responses from both servers, and the event $RESP(R_i)$. If all the actions in this fragment occur at r_i , then we denote this fragment by $E_i(\alpha)^{(x,y)}$, where R_i returns the values (x, y) (Fig. 2) to the external client.
4. We use $R(\alpha)$ and $W(\alpha)$ to denote the READ and WRITE transactions in the context of α . When the context is clear, we simply use R and W .
5. We use the subscript of a returned value to denote the version identifier, which uniquely identifies a version from a totally ordered set. For instance, x_0 is the 0^{th} version of x (the initial value of object o_x) on server s_x .

In our proofs, we frequently use arguments that rely on the existence of non-blocking fragments and the constraints of strict serializability and write isolation. Below we state a few useful lemmas regarding the executions, of some algorithm $\mathcal{A}$, where all READ transactions are assumed to have all SNOW properties; we will use these lemmas in later sections. Due to space constraints, we explain these lemmas at a high level.

The following lemma states that a READ transaction has to return the same version from both servers in order to satisfy strict serializability and write isolation.

Lemma 1. *Suppose α is any execution of $\mathcal{A}$ such that READ transaction R is in α . Suppose the execution fragment $I(\alpha) \circ F_x(\alpha)^{(x_t)} \circ F_y(\alpha)^{(y_s)} \circ E(\alpha)^{(x_{t'}, y_{s'})}$ in α , corresponds to R , where $x_t, x_{t'} \in V_1$ and $y_s, y_{s'} \in V_2$, then $s = s' = t = t'$.*

Proof. Suppose R is invoked at reader r . Then, via the action $send(x_t)_{s_x, r}$, in execution fragment $F_1(\alpha)^{(x_t)}$, server s_x sends the value x_t to r , which is received at r through the action $recv(x_{t'})_{s_x, r}$ in $E(\alpha)^{(x_{t'}, y_{s'})}$. By the assumptions of the reliable channel automata in our model, we have $x_t = x_{t'}$,

i.e., $t = t'$. Similar argument for $F_2(\alpha)^{(y_s)}$ and $E(\alpha)^{(x_{t'}, y_{s'})}$ leads us to conclude $s = s'$. Next, R responds with $(x_{t'}, y_{s'})$, which implies by the S property for executions of $\mathcal{A}$ that $x_{t'}$ and $y_{s'}$ must correspond to the same version, i.e., $s' = t'$. $\square$

The following lemma states that we can create a new execution α' that is indistinguishable to α by swapping two adjoining fragments, which happen on two distinct automata in α if either (a) both fragments have no input actions or (b) one of the fragments have no external (input or output) actions. Our proofs leverage this lemma to create new executions by swapping such fragments and finally derive an execution that violates strict serializability.

Lemma 2 (Commuting fragments). *Let α be an execution of $\mathcal{A}$. Suppose $G_1(\alpha)$ and $G_2(\alpha)$ are any execution fragments in α such that all actions in each fragment occur only at one automaton and either (a) none of the fragments contain input actions, or (b) at least one of the fragments have no external actions. Suppose $G_1(\alpha)$ and $G_2(\alpha)$ occur at two distinct automata and the execution fragment $G_1(\alpha) \circ G_2(\alpha)$ occurs in α . Then there exists an execution α' of $\mathcal{A}$, where the execution fragment $G_2(\alpha) \circ G_1(\alpha)$ appears in α' , such that (i) $G_1(\alpha) \sim G_1(\alpha')$ and $G_2(\alpha) \sim G_2(\alpha')$ (ii) the prefix in α before $G_1(\alpha) \circ G_2(\alpha)$ is identical to the prefix in α' before $G_1(\alpha') \circ G_2(\alpha')$; and (iii) the suffix in α after $G_1(\alpha) \circ G_2(\alpha)$ is identical to the suffix in α' after the execution fragment $G_2(\alpha') \circ G_1(\alpha')$.*

Proof. This is clear because the adversary can move the actions in G_2 to occur before G_1 at their respective automata, and because either (a) none of the fragments have any input action or (b) at least one of them has no external actions, and hence the actions in one of these fragments cannot affect the actions in the other fragment. $\square$

The following lemma states that if there are two fair executions of $\mathcal{A}$ with READ transaction R in each of them, and suppose at any server the non-blocking fragments of R are identical (in terms of the sequence of states and actions), then R returns the similar values in both executions.

Lemma 3 (Indistinguishability). *Let α and β be executions of $\mathcal{A}$ and let R be any READ transaction. Then (i) if $F_x(\alpha) \stackrel{s_x}{\sim} F_x(\beta)$ then both $R(\alpha)$ and $R(\beta)$ respond with the same value x at s_x ; and (ii) if $F_y(\alpha) \stackrel{s_y}{\sim} F_y(\beta)$ then both $R(\alpha)$ and $R(\beta)$ respond with the same value y at s_y .*

Proof. Suppose R is invoked at some reader r . Let $j \in \{1, 2\}$ and suppose the fragments $F_j(\alpha)$ and $F_j(\beta)$ appears in α and β respectively, where in $F_j(\alpha)$ server s_j sends $v_j \in V_j$ to r . Then $R(\alpha)$ must return v_j for object o_j by the O property of $\mathcal{A}$. Then since $F_j(\alpha) \stackrel{s_j}{\sim} F_j(\beta)$ then in $F_j(\beta)$ the server s_j must also send v_j to r , therefore, both $R(\alpha)$ and $R(\beta)$ must return value v_j for o_j . $\square$

The following lemma shows that for any finite execution of $\mathcal{A}$ that ends with the invocation of READ transaction R_1 , it is always possible to have an extended execution of $\mathcal{A}$ where the fragments I , F_x , F_y and E appear consecutively due to the asynchronous network.

Lemma 4. *If any finite execution of $\mathcal{A}$ ends with $INV(R)$, for a READ transaction R_1 then there exists an extension α which is a fair execution of $\mathcal{A}$ and is of the form $P(\alpha) \circ I(\alpha) \circ F_{1,x}(\alpha)^{(x)} \circ F_{1,y}(\alpha)^{(y)} \circ E(\alpha)^{(x,y)} \circ S(\alpha)$, where $P(\alpha)$ is the prefix and $S(\alpha)$ denotes the rest of the execution.*

Proof. Consider a finite execution of $\mathcal{A}$ that end with $INV(R)$, which occurs at some reader r , then the adversary induces the execution fragment $I(\alpha)$ by delaying all actions, except the internal




Figure 2: The relevant actions in the execution fragments $I_i(\alpha)$, $F_{i,x}(\alpha)^x$, $F_{i,y}(\alpha)^y$ and $E_i(\alpha)^{(x,y)}$ for any READ transaction R_i , $i \in \{1, 2\}$ of a fair execution α of $\mathcal{A}$.

and output actions at r , between the actions $INV(R)$ and the later of the actions $send(m_1^r)_{r, s_x}$ and $send(m_2^r)_{r, s_y}$. Next, the adversary delivers m_1^r at s_x (via the action $recv(m_1^r)_{r, s_x}$) and delays all actions, other than internal and output actions at s_x , until s_x responds with x , via $send(x)_{s_x, r}$; we identify this execution fragment as $F_1(\alpha)^{(x)}$. Subsequently, in a similar manner, the adversary delivers the message m_2^r and delays appropriate actions to induce the execution fragment $F_2(\alpha)^{(y)}$. Finally, the adversary delivers the values x and y to r (via the events $recv(x)_{s_x, r}$ and $recv(y)_{s_y, r}$), and delays all actions at other automata until R completes with action $RESP(R)$ by returning (x, y) . As a result, we arrive at a fair execution of $\mathcal{A}$ of the form $I(\alpha) \circ F_1(\alpha)^{(x)} \circ F_2(\alpha)^{(y)} \circ E(\alpha)^{(x,y)} \circ S(\alpha)$. $\square$

4 No SNOW with Three Clients and C2C

This section provides the sketch of a formal proof of the SNOW Theorem with 3 clients, i.e., SNOW is impossible in a system with 3 or more clients even when client-to-client communication is allowed. The main result of this section is captured by the following theorem.

Theorem 1. *The SNOW properties cannot be implemented in a system with two readers and one writer, for two servers even in the presence of client-to-client communication.*

Our proof strategy is to assume the existence of an algorithm $\mathcal{A}$ that satisfies all SNOW properties and create an execution α of $\mathcal{A}$ that contradicts the S property. We begin with an execution of $\mathcal{A}$ that contains READ transactions R_1 and R_2 , which both read s_x and s_y , and WRITE transaction W that writes (x_1, y_1) to s_x and s_y respectively (both servers have initial values x_0, y_0). R_1 begins after W completes, and R_2 begins after R_1 completes. By the S property both R_1 and R_2 should return (x_1, y_1) . Then we create a sequence of executions of $\mathcal{A}$ (Fig. 3), where we interchange




Figure 3: Executions of $\mathcal{A}$ with three clients and operations W , R_1 and R_2 leading to the contradiction of S in α_{10} . Arrows show the transposition of execution fragments from the previous execution.

the fragments until we finally reach an execution in which R_2 completes before R_1 begins, but R_2 returns (x_1, y_1) and R_1 returns (x_0, y_0) which contradicts the S property.

The following lemma shows that in an execution of $\mathcal{A}$ with a $WRITE$ transaction W and a $READ$ transaction R_1 , there exists a point in the execution such that if R_1 is invoked before that point then R_1 returns (x_0, y_0) and if R_1 is invoked after that point then R_1 returns (x_1, y_1) .

Lemma 5 (Existence of α_0 and α_1). *There exist executions α_0 and α_1 of $\mathcal{A}$ that contain transactions W and R_1 that satisfy the following properties where k is some positive integer and $a_1, \dots, a_k$ is a prefix of $a_1, \dots, a_{k+1}$: (i) α_0 can be written as $a_1, \dots, a_k \circ R_1(\alpha_0)^{(x_0, y_0)} \circ S(\alpha_0)$; (ii) α_1 can be written as $a_1, \dots, a_{k+1} \circ R_1(\alpha_1)^{(x_1, y_1)} \circ S(\alpha_1)$; and (iii) a_{k+1} in α_1 occurs at r_1 .*

Proof. Now we describe the construction of a sequence of finite executions of $\mathcal{A}$, $\{\gamma_k\}_{k=0}^\infty$ such that each γ_k contains W and R_1 . Consider an execution α of $\mathcal{A}$ that contains W . Suppose R_1 is invoked at r_1 after the execution fragment $a_1, \dots, a_{k+1}$, a prefix of α . Allowed by network asynchrony, let $INV(R)$ be followed by only internal and external actions at r_1 until both $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ occur, thereby creating an execution fragment of the form $a_1, \dots, a_{k+1} \circ I_1(\alpha)$. We denote $a_1, \dots, a_{k+1}$ by P_{k+1} .

Next, consider the network delivers the message $m_x^{r_1}$ at s_x , and delays all actions at other automata and also any input action at s_x until s_x sends x to r_1 . Therefore, we achieve the execution fragment $P_{k+1} \circ I_{1,x}(\alpha) \circ F_{1,x}(\alpha)$ of $\mathcal{A}$. Next, the network delivers $m_y^{r_1}$ at s_y and delays all actions at other automata and input actions at s_y until s_y sends y to r_1 . Then the network

delivers x and y at r_1 but it delays actions at other automata and any other input action at r_1 until $RESP(R_1)$ occurs. Now we have an execution fragment of $\mathcal{A}$, which can be written as $P_{k+1} \circ I_1(\alpha) \circ F_{1,x}(\alpha)^{(x)} \circ F_{1,y}(\alpha)^{(y)} \circ E_1(\alpha)^{(x,y)}$, where R_1 responds with (x, y) such that $(x, y) \in \{(x_0, y_0), (x_1, y_1)\}$. We denote this finite execution prefix as γ_k . Therefore, there exists a sequence of such finite executions $\{\gamma_k\}_{k=0}^\infty$.

Because R_1 precedes W , by the S property R_1 must respond with (x_0, y_0) in γ_0 . If k is large enough such that a_k occurs in α after the completion of W then by the S property, R_1 must return (x_1, y_1) in γ_{k+1} . Therefore, there exists a minimum k where in γ_k READ transaction R_1 returns (x_0, y_0) and in γ_{k+1} , R_1 returns (x_1, y_1) . We denote this minimum by k^* . Note that γ_{k^*} corresponds to α_0 and γ_{k^*+1} corresponds to α_1 in (i) and (ii) respectively.

Now, we prove case (iii) by eliminating the possibility of a_{k^*+1} occurring at s_x , s_y , w or r_2 . The S property requires that R_1 must retrieve the same version from both s_x and s_y , which implies that s_x and s_y must send values of the same version. Observe that R_1 returns the 0th version in α_0 and the 1st version in α_1 , while the prefixes P_{k^*} and P_{k^*+1} differ by a single action a_{k^*+1} . Importantly, just one action at any of s_x , s_y , r_2 or w is not enough for s_x and s_y to coordinate the same version to send. Therefore, a_{k^*+1} must occur at r_1 , which can possibly help coordinate by sending some information via m_x and m_y sent to s_x and s_y respectively.

Case a_{k^*+1} occurs at s_x : Consider the prefix of execution α_0 up to a_{k^*} . Suppose the network invokes R_1 immediately after action a_{k^*} via $INV(R_1)$. By Lemma 4 there exists an execution α' that contains an execution fragment of the form $P_{k^*} \circ I_1(\alpha') \circ F_{1,x}(\alpha')^{(x)} \circ F_{1,y}(\alpha')^{(y)} \circ E(\alpha')^{(x,y)}$. Then, $I_1(\alpha_1) \stackrel{r_1}{\sim} I_1(\alpha')$ and $F_{1,y}(\alpha_1) \stackrel{s_y}{\sim} F_{1,y}(\alpha')$ because in both executions the actions of I_1 occur entirely at r_1 and those of $F_{1,y}$ occur entirely at s_y , and thus they are unaffected by the addition of the single action a_{k^*+1} at s_x . As a result, $F_{1,y}(\alpha')$ must send the same value y_1 to r_1 as in $F_{1,y}(\alpha_1)$. Then in α' , $R_1(\alpha')$ returns y_1 by Lemma 3, and thus $R_1(\alpha')$ returns (x_1, y_1) by the S property. However, this contradicts the fact that in γ_{k^*} R_1 responds with (x_0, y_0) .

Case a_{k^*+1} occurs at s_y : A contradiction can be shown by following a line of reasoning similar to the preceding case.

Case a_{k^*+1} occurs at w : This can be argued in a similar manner as the previous case with the trivial fact that $F_{1,x}(\alpha_1) \stackrel{s_x}{\sim} F_{1,x}(\alpha')$ and $F_{1,y}(\alpha_1) \stackrel{s_y}{\sim} F_{1,y}(\alpha')$.

Case a_{k^*+1} occurs at r_2 : A contradiction can be derived using a line of reasoning as in the previous case.

So we conclude that a_{k^*+1} must occur at r_1 in α_1 . □

In the remainder of the section, we suppress the explicit reference to the execution. For instance, we use I_i , $F_{i,x}^{(x)}$, $F_{i,y}^{(y)}$, $E_i^{(x,y)}$ and S , where we drop α , instead of $I_i(\alpha)$, $F_{i,x}(\alpha)^{(x)}$, $F_{i,y}(\alpha)^{(y)}$, $E_i(\alpha)^{(x,y)}$ and $S(\alpha)$. If a READ transaction R_i has an execution fragment of the form $I_i \circ F_{i,x}^{(x)} \circ F_{i,y}^{(y)} \circ E_i^{(x,y)}$ we denote it as $R_i^{(x,y)}$. In the rest of the section, α_0 , α_1 , and the value of k are the same as in the discussion above. We denote the execution fragments $a_1, \dots, a_k$ and $a_1, \dots, a_{k+1}$ as P_k and P_{k+1} respectively. Our proof proceeds by stating a sequence set of lemmas (Fig. 3). The first lemma states there exists an execution in which two consecutive READ transactions follow a WRITE transaction, and both READ transactions return the new values by the WRITE transaction.

Lemma 6 (Existence of α_2). *There exists an execution α_2 of $\mathcal{A}$ that contains W , R_1 , and R_2 , and can be written in the form $P_{k+1} \circ R_1^{(x_1, y_1)} \circ R_2^{(x_1, y_1)} \circ S$, where both R_1 and R_2 return (x_1, y_1) .*

Proof. We can construct an execution α_2 of $\mathcal{A}$ as follows. Consider the prefix $a_1, \dots, a_{k+1} \circ R_1(\alpha_1)^{(x_1, y_1)}$ of the execution α_1 , from Lemma 5. At the end of this prefix, the network invokes R_2 . Now, by Lemma 4, due to $INV(R_2)$ there is an extension of the prefix of the form $a_1, \dots, a_{k+1} \circ R_1(\alpha_1)^{(x_1, y_1)} \circ I(\alpha) \circ F_1(\alpha)^{(x)} \circ F_2(\alpha)^{(y)} \circ E(\alpha)^{(x, y)}$. By the S property, we have $x = x_1$ and $y = y_1$. Therefore, α_2 (Fig. 3) can be written in the form $P_{k+1} \circ R_1^{(x_1, y_1)} \circ R_2^{(x_1, y_1)} \circ S$, where S is the rest of the execution. $\square$

Based on the previous execution, the following lemma proves that there is an execution of $\mathcal{A}$ where I_2 occurs earlier than the action a_{k+1} and the invocation of R_1 .

Lemma 7 (Existence of α_3). *There exists execution α_3 of $\mathcal{A}$ that contains transactions W , R_1 and R_2 , and can be written in the form $P_k \circ I_2 \circ a_{k+1} \circ R_1^{(x_1, y_1)} \circ F_{2,x} \circ F_{2,y} \circ E_2 \circ S$, where both R_1 and R_2 return (x_1, y_1) .*

Proof. Consider the execution α_2 as in Lemma 6. In the execution fragment $I_1 \circ F_{1,x}^{(x_1)} \circ F_{1,y}^{(y_1)} \circ E_1^{(x_1, y_1)}$ in α_2 , none of the actions occur at r_2 and by Lemma 5, a_{k+1} occurs at r_1 , also the actions in I_2 occur only at r_2 . Starting with α_2 , and by repeatedly using Lemma 2, we create a sequence of four executions of $\mathcal{A}$ by repeatedly swapping I_2 with the execution fragments $E_1^{(x_1, y_1)}$, $F_{1,y}^{(y_1)}$, $F_{1,x}^{(x_1)}$ and I_1 , which appears in $I_1 \circ F_{1,x}^{(x_1)} \circ F_{1,y}^{(y_1)} \circ E_1^{(x_1, y_1)} \circ I_2$, where the following sequence of execution fragments $I_1 \circ F_{1,x}^{(x_1)} \circ F_{1,y}^{(y_1)} \circ I_2 \circ E_1^{(x_1, y_1)}$ (by commuting I_2 and $E_1^{(x_1, y_1)}$); $I_1 \circ F_{1,x}^{(x_1)} \circ I_2 \circ F_{1,y}^{(y_1)} \circ E_1^{(x_1, y_1)}$ (by commuting I_2 and $F_{1,y}^{(y_1)}$); $I_1 \circ I_2 \circ F_{1,x}^{(x_1)} \circ F_{1,y}^{(y_1)} \circ E_1^{(x_1, y_1)}$ (by commuting I_2 and $F_{1,x}^{(x_1)}$) appear. Finally, we have an execution α' of the form $P_{k+1} \circ I_2 \circ R_1^{(x_1, y_1)} \circ F_{2,x}^{(x_1)} \circ F_{2,y}^{(y_1)} \circ E_2^{(x_1, y_1)} \circ S$ (by commuting I_2 and I_1) Next, from α' , by using Lemma 2 and swapping a_{k+1} with I_2 we have shown the existence of an execution α_3 . $\square$

In the following lemma, we show that we can create an execution α_4 of $\mathcal{A}$, where $F_{2,y}$ occurs immediately before $E_1^{(x_1, y_1)}$, while R_1 and R_2 both return (x_1, y_1) .

Lemma 8 (Existence of α_4). *There exists execution α_4 of $\mathcal{A}$ that contains transactions W , R_1 and R_2 and can be written in the form $P_k \circ I_2 \circ a_{k+1} \circ I_1 \circ F_{1,x} \circ F_{1,y} \circ F_{2,y} \circ E_1 \circ F_{2,x} \circ E_2 \circ S$, where both R_1 and R_2 return (x_1, y_1) .*

Proof. We start with an execution α_3 , as in Lemma 7, and apply Lemma 2 twice.

First, by Lemma 2, we know there exists an execution α' of $\mathcal{A}$ where $F_{2,x}$ (identify as G_1) and $F_{2,y}$ (identify as G_2) are interchanged since actions of $F_{2,x}$ occurs solely at s_x and those of $F_{2,y}$ at s_y , and $F_{2,x}$ and $F_{2,y}$ return x_1 and y_1 , respectively, to r_2 .

Next, by Lemma 2 there is execution of $\mathcal{A}$, say α_4 where the fragments E_1 (identify as G_1) and $F_{2,y}$ (identify as G_2) are interchanged, with respect to α' , because the actions in E_1 occur at r_1 and those of $F_{2,y}$ at s_y . Furthermore, α_4 can be written in the form $P_k \circ I_2 \circ a_{k+1} \circ I_1 \circ F_{1,x}^{(x_1)} \circ F_{1,y}^{(y_1)} \circ F_{2,y}^{(y_1)} \circ E_1^{(x_1, y_1)} \circ F_{2,x}^{(x_1)} \circ E_2^{(x_1, y_1)} \circ S$. $\square$

Next, we create an execution α_5 where $F_{2,y}$ occurs before $F_{1,y}$.

Lemma 9 (Existence of α_5). *There exists execution α_5 of $\mathcal{A}$ that contains transactions W , R_1 and R_2 and can be written in the form $P_k \circ I_2 \circ a_{k+1} \circ I_1 \circ F_{1,x} \circ F_{2,y} \circ F_{1,y} \circ E_1 \circ F_{2,x} \circ E_2 \circ S$, where both R_1 and R_2 return (x_1, y_1) .*

Proof. Given all actions in $F_{1,y}$ and $F_{2,y}$ occur at s_y in α_4 , consider the prefix of α_4 that ends with $F_{1,x}$. We extend this prefix as follows. In this prefix, the actions $send(m_y^{r_2})_{r_2,s_y}$ and $send(m_y^{r_1})_{r_1,s_y}$ do not have their corresponding $recv$ actions. Suppose the network delivers $m_y^{r_2}$ at s_y (via the action $recv(m_y^{r_2})_{r_2,s_y}$) and delays all actions, other than internal and output actions at s_y , until s_y responds with y , via action $send(y)_{s_y,r_2}$. This extended execution fragment is of the form $F_{2,y}$. Similarly, the network further extends the execution by placing the action $recv(m_y^{r_1})_{r_1,s_y}$ at s_y and creates the execution fragment of the form $F_{1,y}$. Note that, so far, the actions due to the above extensions are entirely at s_y . Suppose the network makes the execution fragments E_1 happen next by delivering values sent during $F_{1,x}$ and $F_{1,y}$ via the actions $recv(x)_{s_x,r_1}$ and $recv(y)_{s_y,r_1}$ respectively at r_1 . Then $F_{2,x}$ occurs next, such that this fragment contains exactly the same sequence of actions as in the corresponding execution fragment in α_4 . This is possible because they are not influenced by any output action in $F_{2,y}$ or $F_{1,y}$. Suppose the network places the execution fragment E_2 next. Let us denote the execution that is an extension of this finite execution so far as α_5 , which is of the form $P_k \circ I_2 \circ a_{k+1} \circ I_1 \circ F_{1,x} \circ F_{2,y} \circ F_{1,y} \circ E_1 \circ F_{2,x} \circ E_2 \circ S$. Now we need to argue about the values returned by the reads.

Note that both α_4 and α_5 have the same execution fragment $F_{1,x}(\alpha_4)$. Therefore, $F_{1,x}(\alpha_4) \stackrel{s_x}{\sim} F_{1,x}(\alpha_5)$, and thus s_x also returns x_1 in $F_{1,x}$ in α_5 . Next by Lemma 3 for R_1 , s_y returns y_1 in $F_{1,y}$ and hence by the S property, $R_1(\alpha_5)$ returns (x_1, y_1) , i.e., that r_1 returns the new version of object values. Therefore, $F_{1,x}(\alpha_4)$, $F_{1,y}$ and E_1 are of the form $F_{1,x}^{(x_1)}$, $F_{1,y}^{(y_1)}$ and $E_1^{(x_1,y_1)}$, respectively.

Note that by construction of α_5 above, the execution fragment $F_{2,x}$ in both α_4 and α_5 is the same, therefore, $F_{2,x}(\alpha_4) \stackrel{s_x}{\sim} F_{2,x}(\alpha_5)$. Hence as in α_4 , s_x returns x_1 in the execution fragment $F_{2,x}(\alpha_5)$ in α_5 , i.e., of the form $F_{2,x}(\alpha_5)^{(x_1)}$. Since s_x returns x_1 in $F_{2,x}$ in α_5 , by Lemma 3 and the S property, R_2 returns (x_1, y_1) and hence E_2 is of the form $E_2^{(x_1,y_1)}$.

From the above argument we know that α_5 is of the form $P_k \circ I_2 \circ a_{k+1} \circ I_1 \circ F_{1,x}^{(x_1)} \circ F_{2,y}^{(y_1)} \circ F_{1,y}^{(y_1)} \circ E_1^{(x_1,y_1)} \circ F_{2,x}^{(x_1)} \circ E_2^{(x_1,y_1)} \circ S$. $\square$

In the next lemma, we show the existence of an execution of $\mathcal{A}$ where R_1 returns (x_0, y_0) and I_2 occurs immediately after a_k and R_2 responds with (x_1, y_1) .

Lemma 10 (Existence of α_6). *There exists execution α_6 of $\mathcal{A}$ that contains transactions W , R_1 and R_2 and can be written in the form $P_k \circ I_2 \circ I_1 \circ F_{1,x} \circ F_{2,y} \circ F_{1,y} \circ E_1 \circ F_{2,x} \circ E_2 \circ S$, where R_1 returns (x_0, y_0) and R_2 returns (x_1, y_1) .*

Proof. The crucial part of this proof is to carefully use the result of Lemma 5 so that R_1 returns (x_0, y_0) , instead of (x_1, y_1) . Note that the same prefix P_k appears in α_5 of Lemma 9 as well as in α_0 and α_1 of Lemma 5, where k is defined as in Lemma 5.

By Lemma 5, action a_{k+1} occurs at r_1 . In the execution fragment $a_{k+1} \circ I_1 \circ F_{1,x}^{(x_1)} \circ F_{2,y}^{(y_1)}$ of α_5 , the actions in $a_{k+1} \circ I_1$ occur at r_1 ; actions in $F_{1,x}^{(x_1)}$ occur at s_x ; and actions in $F_{2,y}^{(y_1)}$ occur at s_y . Now consider the prefix of execution α_4 ending with I_2 and the network invokes R_1 immediately after I_2 (instead of after a_{k+1}) and extends it by the execution fragment $I_1 \circ F_{1,x} \circ F_{2,y}$ to create a new finite execution ϵ , which is of the form $P_k \circ I_2 \circ I_1 \circ F_{1,x} \circ F_{2,y}$. As a result, a_{k+1} may not be in ϵ because we introduce changes before a_{k+1} occurs.

Note that if in the prefix $P_k \circ I_2(\epsilon) \circ I_1(\epsilon) \circ F_{1,x}(\epsilon) \circ F_{2,y}(\epsilon)$ of ϵ we ignore the actions in $I_2(\epsilon)$ then the remaining execution is the same as the prefix $P_k \circ I_1(\alpha_0) \circ F_{1,x}(\alpha_0) \circ F_{2,y}(\alpha_0)$ of α_0 in Lemma 5. Here we explicitly use the notations ϵ and α_0 to avoid confusion. Since the actions in

$I_2(\epsilon)$ have no impact on the actions in $I_1(\epsilon) \circ F_{1,x}(\epsilon) \circ F_{2,y}(\epsilon)$, we have $F_{1,x}(\epsilon) \stackrel{s_x}{\sim} F_{1,x}(\alpha_0)$. Therefore, by Lemma 1 $F_{1,x}(\epsilon)$ returns x_0 as in $F_{1,x}(\alpha_0)$, i.e., s_x returns x_0 in $F_{1,x}$. Now by Lemma 3 we conclude that for any extension of ϵ , say γ , READ transaction $R_1(\gamma)$ returns x_0 at s_x and by the S property $R_1(\gamma)$ returns (x_0, y_0) . Also, since $F_{2,y}(\alpha_5) \stackrel{s_y}{\sim} F_{2,y}(\epsilon) \stackrel{s_y}{\sim} F_{2,y}(\gamma)$ by Lemma 3 and the S property, $R_2(\gamma)$ must return (x_1, y_1) . Therefore, γ has an extension of α_6 (Fig. 3) which is of the form $P_k \circ I_2 \circ I_1 \circ F_{1,x}^{(x_0)} \circ F_{2,y}^{(y_1)} \circ F_{1,y}^{(y_0)} \circ E_1^{(x_0,y_0)} \circ F_{2,x}^{(x_1)} \circ E_2^{(x_1,y_1)} \circ S$ as in the statement of the lemma. $\square$

The following lemma shows that there exists an execution α_7 for $\mathcal{A}$ where $F_{2,x}$ appears before $F_{1,y} \circ E_1$, where R_1 returns (x_0, y_0) and R_2 returns (x_1, y_1) . The lemma can be proven by starting from α_6 in Lemma 10 and moving the execution fragments of R_2 earlier, a little at a time, until finally we have R_2 finishing before R_1 starts. This simply uses commutativity since the actions in the swapped execution fragments occur at different automata.

Lemma 11 (Existence of α_7). *There exists execution α_7 of $\mathcal{A}$ that contains transactions W , R_1 and R_2 , and can be written in the form $P_k \circ I_2 \circ I_1 \circ F_{1,x} \circ F_{2,y} \circ F_{2,x} \circ F_{1,y} \circ E_1 \circ E_2 \circ S$ where R_1 returns (x_0, y_0) and R_2 returns (x_1, y_1) .*

Proof. This result is proved by applying the result of Lemma 2 to the execution created in Lemma 10. Suppose, α_6 (Fig. 3) is a execution as in Lemma 10, where in the execution fragment $E_1^{(x_0,y_0)} \circ F_{2,x}$ we identify $E_1^{(x_0,y_0)}$ as G_1 and $F_{2,x}^{(y_0)}$ as G_2 . The actions of G_1 and G_2 occur at two distinct automata, therefore, we can use the result of Lemma 2, to argue that there exists an execution α' of $\mathcal{A}$ that contains the execution fragment $F_{2,x} \circ E_1^{(x_0,y_0)}$, and α_6 and α' are identical in the prefixes and suffixes corresponding to G_1 and G_2 .

Now, α' contains $F_{1,y} \circ F_{2,x}$, where the actions in $F_{1,y}$ (identified as G_1) and $F_{2,x}$ (identify as G_2) occur at distinct automata. Hence, by Lemma 2 there exists an execution α_7 of the form $P_k \circ I_2 \circ I_1 \circ F_{1,x}^{(x_0)} \circ F_{2,y}^{(y_1)} \circ F_{2,x}^{(x_1)} \circ F_{1,y}^{(y_0)} \circ E_1^{(x_0,y_0)} \circ E_2^{(x_1,y_1)} \circ S$. $\square$

The following lemma leverages Lemma 2 to show the existence of an execution α_8 of $\mathcal{A}$ where $F_{2,y}$ appears before $I_1 \circ F_{1,x}$, and R_1 returns (x_0, y_0) while R_2 returns (x_1, y_1) .

Lemma 12 (Existence of α_8). *There exists execution α_8 of $\mathcal{A}$ that contains transactions W , R_1 and R_2 and can be written in the form $P_k \circ I_2 \circ F_{2,y} \circ I_1 \circ F_{1,x} \circ F_{2,x} \circ F_{1,y} \circ E_1 \circ E_2 \circ S$, where R_1 returns (x_0, y_0) and R_2 returns (x_1, y_1) .*

Proof. Consider the execution α_7 of $\mathcal{A}$ as in Lemma 11. In the context of of Lemma 2, in α_7 (Fig. 3) the actions in $F_{1,x}$ (identify as G_1) occur at s_x and those in $F_{2,y}$ (identify as G_2) at s_y . Then by Lemma 2 there exists an execution α' of $\mathcal{A}$, of the form $P_k \circ I_2 \circ I_1 \circ F_{2,y} \circ F_{1,x} \circ F_{2,x} \circ F_{1,y} \circ E_1 \circ E_2 \circ S$, where $F_{2,y}$ and $F_{1,x}$ are interchanged.

Since actions in $F_{2,y}$ (identify as G_1) occur at s_y and those in I_1 (identify as G_1) occur at r_1 then by Lemma 2 there is a execution of $\mathcal{A}$, α_8 where $F_{2,y}$ appear before I_1 , i.e., of the form $P_k \circ I_2 \circ F_{2,y} \circ I_1 \circ F_{1,x} \circ F_{2,x} \circ F_{1,y} \circ E_1 \circ E_2 \circ S$, where $F_{2,y}$ and I_1 are interchanged.

By (ii) of Lemma 3 we have $F_{2,x}(\alpha') \stackrel{s_x}{\sim} F_{2,x}(\alpha_8)$ hence $F_{2,x}$ sends x_1 and $F_{1,x}$ and $F_{1,y}$ sends x_0 and y_0 , respectively. So considering these returned values we have α_8 (Fig. 3) in the form as stated in the lemma. $\square$

The following lemma shows the existence of an execution α_9 , of $\mathcal{A}$, where $F_{2,x}$ appears before $F_{1,x}$.

Lemma 13 (Existence of α_9). *There exists execution α_9 of $\mathcal{A}$ that contains transactions W , R_1 and R_2 and can be written in the form $P_k \circ I_2 \circ F_{2,y}^{(y_1)} \circ I_1 \circ F_{2,x}^{(x_1)} \circ F_{1,x}^{(x_0)} \circ F_{1,y}^{(y_0)} \circ E_1^{(x_0,y_0)} \circ E_2^{(x_1,y_1)} \circ S$ where R_1 returns (x_0, y_0) and R_2 returns (x_1, y_1) .*

Proof. In α_8 from Lemma 12, all the actions in I_1 occur at r_1 ; those in $F_{1,x}$ occur at s_x ; and the actions in $F_{2,x}$ occur only at s_x . Note that actions of both execution fragments $F_{2,x}$ and $F_{1,x}$ occur at r_1 . Consider the prefix of α_8 that ends with I_1 then suppose the network extends this prefix by adding an execution fragment of the form $F_{2,x} \circ F_{1,x}$ as follows. First note that the actions $send(m_x^{r_2})_{r_2,s_x}$ and $send(m_x^{r_1})_{r_1,s_x}$ appears in the prefix but do not have corresponding *recv* actions. The network places action $recv(m_x^{r_2})_{r_2,s_x}$, and allows an execution fragment of the form $F_{2,x}$ to appear. Now, immediately after this the network further extends it with an execution fragment of the form $F_{1,x}$ by placing action $recv(m_x^{r_1})_{r_1,s_x}$. Next the fragment $F_{1,y}$ is added and is the same as $F_{1,y}(\alpha_8)$. This last step can be argued by the fact that none of the actions in $F_{1,y}$ can be affected by any of the output actions at $F_{2,x}$ and $F_{1,x}$. Following this the network allows the rest of the execution by adding an execution fragment of the form $E_1 \circ E_2 \circ S$. The resulting execution is of the form $P_k \circ I_2 \circ F_{2,y}^{(y_1)} \circ I_1 \circ F_{2,x} \circ F_{1,x} \circ F_{1,y}^{(y_0)} \circ E_1 \circ E_2 \circ S$, where we retained the values wherever it is known, and we denote this execution by α_9 .

Now, we argue about the return values in α_9 . Applying Lemma 3 to R_2 and $F_{2,y}$ implies that R_2 returns (x_1, y_1) . Similarly, applying Lemma 3 to R_1 and $F_{1,y}$ implies that R_1 must return (x_0, y_0) in α_9 . $\square$

Now by constructing a sequence of executions, α_3 through α_{10} (Fig. 3, lemmas and proofs omitted due to lack of space) realize the existence of an execution α_{10} of $\mathcal{A}$ where the execution fragments corresponding to R_2 appears before R_1 , where R_1 returns (x_0, y_0) and R_2 completes by returning (x_1, y_1) but R_1 is in real time after R_2 , therefore, violates the S property.

Lemma 14 (Existence of α_{10}). *There exists an execution α_{10} of $\mathcal{A}$ that contains transactions W , R_1 and R_2 and can be written in the form $P_k \circ R_2^{(x_1,y_1)} \circ R_1^{(x_0,y_0)} \circ S$. where R_1 returns (x_0, y_0) and R_2 returns (x_1, y_1) .*

Proof. Now, by applying Lemma 2 to α_9 , we can swap $F_{2,x}$ and I_1 to create an execution α_{10} (Fig. 3) of $\mathcal{A}$, which is of the form $P_k \circ I_2 \circ F_{2,y}^{(y_1)} \circ F_{2,x}^{(x_1)} \circ R_1^{(x_0,y_0)} \circ E_2^{(x_1,y_1)} \circ S$, where the returned values are determined by Lemma 1.

Note that none of the actions in $I_1 \circ F_{1,x}^{(x_0)} \circ F_{1,y}^{(y_0)} \circ E_1^{(x_0,y_0)}$ occur at r_2 and all actions in $E_2^{(x_1,y_1)}$ occur at r_2 . Therefore, by applying Lemma 2, we can consecutively swap E_2 with E_1 , $F_{1,y}$, I_1 , and $F_{1,x}$. Therefore, we create a sequence of four executions of $\mathcal{A}$ to arrive at execution α_{10} (Fig. 3) of the form $P_k \circ R_2^{(x_1,y_1)} \circ R_1^{(x_0,y_0)} \circ S$. $\square$

Using the above results, we prove Theorem 1 for 3-clients by showing the existence of α_{10} , where R_2 completes before R_1 is invoked and R_2 returns (x_1, y_1) whereas R_1 returns (x_0, y_0) , which violates the S property.

5 Two Client Open Question

This section closes the open question of whether SNOW properties can be implemented with two clients. We first prove that SNOW remains impossible in a 2-client 2-server system if the clients

cannot directly send messages to each other. However, in the presence of client to client communication, it is possible to have all SNOW properties with two clients and at least two servers.

5.1 No SNOW Without C2C Messages

In this section, we prove the following results that states it is impossible guarantee the SNOW properties in a transaction processing system with two clients, without client-to-client communication.

Theorem 2. *The SNOW properties cannot be implemented in a system with two clients and two servers, where the clients do not communicate with each other.*

We use system model with two servers s_x and s_y with two clients, a reader r_1 that issues only READ transactions and a writer w that issues only WRITE transactions. A WRITE transaction W writes (x_1, y_1) to s_x and s_y , and a READ transaction R reads both servers. We assume that there is a bi-directional communication channel between any pair of client and server and any pair of servers. There is no communication channel between clients. We assume that each transaction can be identified by a unique number, e.g., transaction identifier.

Our strategy is still proof by contradiction: We assume there exists some algorithm $\mathcal{A}$ that satisfies all SNOW properties, and then we show the existence of a sequence of executions of $\mathcal{A}$, eventually leading to an execution that contradicts the S property. First, we show the existence of an execution α of $\mathcal{A}$ where R_1 is invoked after W completes, where the send actions $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ at the r_1 occur consecutively in $P(\alpha)$, which is a prefix of α . Then we show that α can be written in the form $P(\alpha) \circ F_{1,x}(\alpha)$ (Fig. 4 (a), Lemma 15). We then prove the existence of another execution β , which can be written in the form $P(\beta) \circ F_{1,x}(\beta) \circ F_{1,y}(\beta)$ by extending α with an execution fragment $F_{1,y}(\beta)$, such that $F_{1,x}(\beta) \stackrel{s_x}{\sim} F_{1,x}(\alpha)$ (Fig. 4 (b); Lemma 16). Note that in any arbitrary extension of β , R_1 eventually returns (x_1, y_1) . Next, we show the existence of an execution γ of the form $P(\gamma) \circ F_{1,x}(\gamma) \circ F_{1,y}(\gamma)$, where the send actions $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ at r_1 occur before W is invoked (Fig. 4 (c); γ), but $F_{1,x}(\gamma)$ and $F_{1,y}(\gamma)$ occur after $RESP(W)$ as in β . Based on γ , we show the existence of an execution δ of the form $P(\eta) \circ F_{1,x}(\eta) \circ F_{1,y}(\eta) \circ S(\eta)$, where R_1 responds with (x_1, y_1) . Finally, starting with η , we create a sequence of executions $\delta(\equiv \eta)$, $\delta^{(1)}$, $\dots$ $\delta^{(f)}$, of $\mathcal{A}$, where in each of them R_1 responds with (x_1, y_1) (Fig. 4 (e) and (g); Lemma 2). Additionally, for any $\delta^{(i)}$, the fragments $F_{1,x}(\delta^{(i)})$ and $F_{1,y}(\delta^{(i)})$ appear before $\delta^{(i-1)}$. Based on $\delta^{(f)}$, we prove the existence of an execution ϕ , where R_1 returns (x_1, y_1) even before W begins, which violates the S property.

Now, we explain the relevant lemmas used in the main proof. The first lemma states that there exists a finite execution of $\mathcal{A}$, where R_1 begins after W completes, and the actions $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ occur before either of the servers s_x and s_y receives $m_x^{r_1}$ or $m_y^{r_1}$ from r_1 ; also, server s_x responds to r_1 in a non-blocking manner (execution α in Fig. 4).

Lemma 15. *There exists a finite execution α of $\mathcal{A}$ that contains transactions $R_1(\alpha)$ and $W(\alpha)$ where $INV(R)$ appears after $RESP(W)$, and the following conditions hold:*

- (i) *The actions $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ appear consecutively in $trace(\alpha)|_{r_1}$; and*
- (ii) *α contains the execution fragment $F_{1,x}(\alpha)$.*

Proof. Consider a finite execution fragment of $\mathcal{A}$ with a completed transaction W , where after W completes the adversary invokes R , i.e., $INV(R_1)$ occurs. Note that each of the read operations $op_1^{r_1}$ and $op_2^{r_1}$, in R_1 , can be invoked by the adversary at any point in the execution. Following $INV(R)$, the adversary introduces the invocation action $inv(op_1^r)$; by the O property of the read operations of $\mathcal{A}$ the action $send(m_x^{r_1})_{r_1, s_x}$ eventually occurs. Next, the adversary introduces $inv(op_2^{r_1})$ and also, delays the arrival of $m_x^{r_1}$ until action $send(m_y^{r_1})_{r_1, s_y}$ eventually occurs, which must occur in accordance with the property O of read operations. Let us call this finite execution α^0 .

Next, suppose at the end of α^0 the adversary delivers the message $m_x^{r_1}$, which has been delayed so far, via the action $recv(m_x^{r_1})_{r, s_x}$, at s_x , but it delays any other input actions at s_x . Note that by the N property of read operations s_x eventually responds with $send(x)_{s_x, r_1}$, with one value x by o property, where $x = x_1$ by the S property, since R begins after W completes. Let us call this execution α . Note that α satisfies conditions (i) and (ii) by the design of the execution. $\square$

The following lemma states that there is an execution β where R begins after W completes, and the two send events at r_1 occurs before $F_{1,x}(\beta)$, which thus occurs before $F_{1,y}(\beta)$ (β in Fig. 4).

Lemma 16. *There exists an execution β of $\mathcal{A}$ that contains transactions R_1 and W where $INV(R_1)$ appears after $RESP(W)$ and the following conditions hold:*

- (i) *The actions $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ appear consecutively in $trace(\beta)|_{r_1}$; and*
- (ii) *β contains the execution fragment $F_{1,x}(\beta) \circ F_{1,y}(\beta)$.*

Proof. Consider the execution α of $\mathcal{A}$ as constructed in Lemma 15. At the end of the execution fragment α , the adversary delivers the previously delayed message m_y^r , which is sent via the action $send(m_y^r)_{r, s_y}$, by introducing the action $recv(m_y^r)_{r, s_y}$. The adversary then delays any other input action in $\mathcal{A}$. By the N property, server s_y must respond to r , with some value y , and hence the output action $send(y)_{s_y, r}$ must eventually occur at s_y . Let us call this finite execution as β . Note that β satisfies the properties (i) and (ii) in the statement of the lemma. $\square$

The following result shows that starting with β there is an execution γ , where R_1 is invoked before W is invoked while $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ occur before $INV(W)$ (Fig. 4 (c)) and $m_x^{r_1}$ and $m_y^{r_1}$ from r_1 reach s_x and s_y after the action $RESP(W)$.

Lemma 17. *There exists an execution γ that contains R_1 and W , where the action $INV(R_1)$ appears before $INV(W)$ and $RESP(R_1)$ appears after $RESP(W)$, and the following conditions hold for γ :*

- (i) *The actions $send(m_x^{r_1})_{r_1, s_x}$ and $send(m_y^{r_1})_{r_1, s_y}$ appear before $INV(W)$ and they appear consecutively in $trace(\gamma)|_{r_1}$;*
- (ii) *γ contains the execution fragment $F_{1,x}(\gamma) \circ F_{1,y}(\gamma)$; and*
- (iii) *action $RESP(W)$ occurs before $F_{1,x}(\gamma)$.*

Proof. Consider the execution β of $\mathcal{A}$ as in Lemma 16. Note that β is an execution of the composed automaton $\mathcal{A} (\equiv S_1 \times r)$. In β , the actions $send(m_x^r)_{r, s_x}$ and $send(m_y^r)_{r, s_y}$ occur at r ; and following that, the actions $recv(m_x^r)_{r, s_x}$, $recv(m_y^r)_{r, s_y}$, $send(x)_{s_x, r}$ and $send(y)_{s_y, r}$ occur at S_1 . Consider the executions $\alpha_r \equiv \beta|_r$ and $\alpha_{S_1} \equiv \beta|_{S_1}$. Let s_β denote $trace(\beta)$.

In s_β , $send(m_x^r)_{r,s_x}$, $send(m_y^r)_{r,s_y}$ appear after $RESP(W)$, as in $trace(\beta)$. Let s'_β be the sequence of external actions of S_2 which we construct from s_β by moving $send(m_x^r)_{r,s_x}$, $send(m_y^r)_{r,s_y}$ before $INV(W)$, which is also an external action of $\mathcal{A}$, and leaving the rest of the actions in s_β as it is.

In β , $INV(R)$, $recv(x)_{s_x,r}$ and $recv(y)_{s_y,r}$ are the only input actions at r , therefore, $s'_\beta|_r = trace(\alpha_r)$. On the other hand, $recv(m_x^r)_{r,s_x}$, $recv(m_y^r)_{r,s_y}$ are the only input actions at s_x , therefore, $s'_\beta|_{S_1} = trace(\alpha_{S_1})$. Now, by Theorem 6, there exists an execution γ of $\mathcal{A}$ such that, $s'_\beta = trace(\gamma)$ and $\alpha_r = \gamma|_r$ and $\alpha_{S_1} = \gamma|_{S_1}$. Therefore, in γ , $send(m_x^r)_{r,s_x}$, $send(m_y^r)_{r,s_y}$ appear before $INV(W)$ (condition (i)) and since $s'_\beta = trace(\gamma)$ condition (ii) holds. Condition (iii) holds trivially. $\square$

In the following lemma we show that in any execution, of $\mathcal{A}$, that is an extension of either execution β or execution γ , as in the preceding lemmas, R_1 eventually returns (x_1, y_1) .

Lemma 18. *Let ξ be an execution of $\mathcal{A}$ that is an extension of the either execution β from Lemma 16 or execution γ from Lemma 17, then $R(\xi)$ responds with (x_1, y_1) .*

Proof. Note that in executions β and γ , the traces $trace(\beta)|_{s_x}$ is a prefix of $trace(\gamma)|_{s_x}$, since β ends with $F_{1,x}(\beta)$ and γ ends with $F_{1,y}(\gamma)$. Therefore, in both β and γ , the respective $send(x)_{s_x,r}$ actions have the same value for their x 's. Now, in any extended execution η of $\mathcal{A}$, which starts with β or γ , by the properties N and O the transaction R completes; and by the property S , R returns (x_1, y_1) . Therefore, $R_1(\xi)$ returns (x_1, y_1) . $\square$

In the following lemma, we show there exists an execution η of $\mathcal{A}$ of the form $P(\eta) \circ F_{1,x}(\eta) \circ F_{1,y}(\eta) \circ S(\eta)$ where $RESP(R)$ appears in $S(\eta)$ (Fig. 4 (d)) and $R(\eta)$ returns (x_1, y_1) .

Lemma 19. *There exists an execution η of $\mathcal{A}$ that contains transactions R and W where $INV(R)$ appears before $INV(W)$; $RESP(R_1)$ appears after $RESP(W)$ and the following conditions hold for η :*

- (i) η can be written in the form $P(\eta) \circ F_{1,x}(\eta) \circ F_{1,y}(\eta) \circ S(\eta)$, for some $P(\eta)$ and $S(\eta)$;
- (ii) The actions $send(m_x^r)_{r_1,s_x}$ and $send(m_y^r)_{r_1,s_y}$ appear before $INV(W)$ and they appear consecutively in $trace(\eta)|_{r_1}$;
- (iii) action $RESP(W)$ occurs before $F_{1,x}(\eta)$; and
- (iv) $R_1(\eta)$ returns (x_1, y_1) .

Proof. Let γ be an execution of $\mathcal{A}$, as described in Lemma 17. Let γ^0 be the execution fragment of γ up to the action $send(y)_{s_y,r}$. Now, by Theorem 7 (1), there exists an execution $\gamma^0 \circ \mu$, of $\mathcal{A}$, where μ denotes the extended portion of the execution.

Clearly, by the N and O properties, the actions $resp(op_1^r)$ and $resp(op_2^r)$ must eventually occur in $\gamma^0 \circ \mu$. Now, identify η as $\gamma^0 \circ \mu$, where $P(\eta) \circ F_{1,x}(\eta) \circ F_{1,y}(\eta)$ is γ^0 , and μ is $S(\eta)$, thereby, proving condition (i).

Note the condition (ii) is satisfied by η because $RESP(W)$ appears in $P(\eta)$, therefore, the execution γ is equivalent to the execution fragment of $P(\eta)$ up to the event $INV(W)$, and also, γ satisfies condition (ii) as stated in Lemma 17.

Condition (iii) is true because $F_{1,x}(\eta)$ begins with action $recv(m_x^r)_{r,s_x}$, which occurs after $RESP(W)$. Condition (iv) is satisfied by η because η is an extension of γ and due to the result of Lemma 18. $\square$




Figure 4: Schematic representation of executions α , β , γ and η , of $\mathcal{A}$ with transactions R and W . The executions evolve from left to right. The dots denote external events at clients. The up-arrow marks denote external actions at s_x . The down-arrow marks denote external actions at s_y .

The following proof proves Theorem 2. In the proof we start with an execution η and create a sequence of executions of $\mathcal{A}$, where each one is of the form $P(\cdot) \circ F_{1,x}(\cdot) \circ F_{1,y}(\cdot) \circ S(\cdot)$, with progressively shorter $P(\cdot)$ until we have a final execution that contradicts the S property.

Proof. Consider an execution $\delta^{(\ell)}$ of $\mathcal{A}$ as in Lemma 19, and let $P(\delta^{(\ell)})$ be the execution fragment $\sigma_0, a_1, \dots, a_\ell, \sigma_\ell$. By Lemma 19, $RESP(R_1(\delta^{(\ell)}))$ returns (x_1, y_1) and $\delta^{(\ell)}$ is also of the form $\sigma_0, a_1, \dots, a_\ell, \sigma_\ell \circ F_{1,x}(\delta^{(\ell)}) \circ F_{1,y}(\delta^{(\ell)}) \circ S(\delta^{(\ell)})$.

Next, we inductively prove the existence of a finite sequence of executions of $\mathcal{A}$ —i.e., by proving the existence of a new execution based on the existence of a previous one—as $\delta^{(\ell)}, \delta^{(\ell-1)}, \dots, \delta^{(i)}, \delta^{(i-1)}, \dots, \delta^{(f)}$, for some positive integer f , with the following properties: (a) Each of the execution in the sequence can be written in the form $\sigma_0, a_1, \dots, a_i, \sigma_i \circ F_{1,x}(\delta^{(i)}) \circ F_{1,y}(\delta^{(i)}) \circ S(\delta^{(i)})$ (or $P(\cdot) \circ F_{1,x}(\cdot) \circ F_{1,y}(\cdot) \circ S(\cdot)$); (b) for each i , $f \leq i < \ell$, we have $P(\delta^{(i)})$ to be a prefix of $P(\delta^{(i+1)})$; and (c) $R_1(\delta^{(f)})$ returns (x_0, y_0) , and for any i , $f < i \leq \ell$, we have $R_1(\delta^{(i)})$ that returns (x_1, y_1) . Note that there is a final execution of the form $\delta^{(f)}$ because of the initial values of x_0 and y_0 and the WRITE transaction W .

Clearly, there exists an integer k , $f \leq k < \ell$, such that $R(\delta^{(k)})$ returns (x_0, y_0) and $R_1(\delta^{(k+1)})$ returns (x_1, y_1) . Now we start with execution $\delta^{(k+1)}$ and construct an execution $\delta^{(k)}$ as described in the rest of the proof. The following argument will show that $R(\delta^{(k)})$ must also return (x_1, y_1) , which contradicts the assumption of having the S property.

Consider the execution $\delta^{(k+1)}$ of the form $\sigma_0, a_1, \dots, a_{k+1}, \sigma_{k+1} \circ F_{1,x}(\delta^{(k+1)}) \circ F_{1,y}(\delta^{(k+1)}) \circ S(\delta^{(k+1)})$. The action a_{k+1} can occur at any of the automata r , w , s_x , and s_y . Therefore, we consider the following four possible cases.

Case (i) a_{k+1} occurs at w : The execution fragments $F_{1,x}(\delta^{(k+1)})$ and $F_{1,y}(\delta^{(k+1)})$ do not contain any input action at s_x or s_y . a_{k+1} does not occur at s_x or s_y . Therefore, the asynchronous network can delay the occurrence of a_{k+1} at w to create the finite execution $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\delta^{(k+1)}) \circ F_{1,y}(\delta^{(k+1)})$, of. Also, there exists an execution $\delta^{(k)}$ that is an

extension of the above finite execution, where R_1 completes in $\delta^{(k)}$. Clearly, $\delta^{(k)}$ can be written as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\delta^{(k)}) \circ F_{1,y}(\delta^{(k)}) \circ S(\delta^{(k)})$, where $S(\delta^{(k)})$ is the tail of the execution resulting from the extension. Moreover, $F_{1,x}(\delta^{(k+1)})$ is indistinguishable from $F_{1,x}(\delta^{(k)})$ at s_x , i.e., $F_{1,x}(\delta^{(k+1)}) \stackrel{s_x}{\approx} F_{1,x}(\delta^{(k)})$. Therefore, $send(x)_{s_x, r_1}$ has the same object value x in both fragments, which means R_1 returns x_1 , and thus $R_1(\delta^{(k)})$ must return (x_1, y_1) by the property S .

Case (ii) a_{k+1} occurs at r : Similar to Case (i).

Case (iii) a_{k+1} occurs at s_x : Observe that the two execution fragments $a_{k+1}\sigma_{k+1} \circ F_{1,x}(\delta^{(k+1)})$ and $F_{1,y}(\delta^{(k+1)})$ occur at separate automata, i.e., at s_x and s_y respectively. Also, the execution fragments $a_{k+1}\sigma_{k+1} \circ F_{1,x}(\delta^{(k+1)})$ and $F_{1,y}(\delta^{(k+1)})$ do not contain any input actions at s_x or s_y . Therefore, we can create an execution ϵ , which can be expressed as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,y}(\epsilon) \circ a_{k+1}, \sigma_{k+1} \circ F_{1,x}(\epsilon) \circ S(\epsilon)$. Clearly, $F_{1,x}(\epsilon) \stackrel{s_x}{\approx} F_{1,x}(\delta^{(k+1)})$ and $F_{1,y}(\epsilon) \stackrel{s_y}{\approx} F_{1,y}(\delta^{(k+1)})$. Because $send(x)_{s_x, r_1}$ occurs in both $F_{1,x}(\epsilon)$ and $F_{1,x}(\delta^{(k+1)})$, it sends the same value x_1 to r_1 . Therefore, R returns (x_1, y_1) .

Now let us denote the execution fragment $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,y}(\epsilon)$ by ϵ' , which is simply a finite prefix of ϵ . Allowed by the asynchronous network, we append $recv(m_x^{r_1})_{r_1, s_x}$ to ϵ' , and create a finite execution ϵ'' as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,y}(\epsilon''), recv(m_x^{r_1})_{r_1, s_x}$, and delay any input action at s_y .

Let ϵ''' be an extension of ϵ'' . Clearly, $F_{1,y}(\epsilon''') \stackrel{s_y}{\approx} F_{1,y}(\epsilon'')$, and thus $send(y)_{s_y, r_1}$ sends the same value y_1 in ϵ'' and ϵ''' . By the N property, $send(x)_{s_x, r_1}$ eventually occurs, and by the O property x is sent to r_1 . Therefore, R_1 completes in ϵ''' , which implies that $R(\epsilon''')$ must return (x_1, y_1) .

Note that the execution fragment of ϵ''' has no input actions of s_x between $recv(m_x^{r_1})_{r_1, s_x}$ and $send(x)_{s_x, r_1}$, which can be identified as $F_{1,x}(\epsilon''')$. Therefore, ϵ''' can be written as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,y}(\epsilon''') \circ F_{1,x}(\epsilon''') \circ S(\epsilon''')$.

Next, since $F_{1,x}(\epsilon''')$ and $F_{1,y}(\epsilon''')$ contain actions of different automata, we can create an execution prefix $\epsilon^{(iv)}$ as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\epsilon^{(iv)}) \circ F_{1,y}(\epsilon^{(iv)})$, where $F_{1,x}(\epsilon^{(iv)})$ appears before $F_{1,y}(\epsilon^{(iv)})$. Next, we create an execution $\delta^{(k)}$ as an extension of $\epsilon^{(iv)}$, which can be written as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\delta^{(k)}) \circ F_{1,y}(\delta^{(k)}) \circ S(\delta^{(k)})$. Then by the O and N properties, R_1 completes in $\delta^{(k)}$. Since $F_{1,x}(\epsilon^{(iv)}) \stackrel{s_x}{\approx} F_{1,x}(\epsilon''')$, $send(x)_{s_x, r_1}$ returns x_1 in $\epsilon^{(iv)}$. Similarly, because $F_{1,y}(\delta^{(k)}) \stackrel{s_y}{\approx} F_{1,y}(\epsilon^{(iv)})$, $send(y)_{s_y, r_1}$ returns y_1 in $\delta^{(k)}$. So, $R_1(\delta^{(k)})$ returns (x_1, y_1) .

Case (iv) a_{k+1} occurs at s_y : Because a_{k+1} occurs at server s_y and $F_{1,x}(\delta^{(k+1)})$ occurs at server s_x (different automata), we can create a new execution ϵ of $\mathcal{A}$ (Fig. 4 (f)) as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\epsilon) \circ a_{k+1}, \sigma_{k+1} \circ F_{1,y}(\epsilon) \circ S(\epsilon)$, such that $F_{1,x}(\epsilon) \stackrel{s_x}{\approx} F_{1,x}(\delta^{(k+1)})$, where a_{k+1}, σ_{k+1} occurs after $F_{1,x}(\delta^{(k+1)})$ and $R(\epsilon)$ returns (x_1, y_1) .

Now, consider the finite execution $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\epsilon)$ at the end of which we append $recv(m_y^{r_1})_{r_1, s_y}$ to create a finite execution of $\mathcal{A}$ as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\epsilon), recv(m_y^{r_1})_{r_1, s_y}$. Then, there exists an execution ϵ' of $\mathcal{A}$, where the network delays the input actions at s_y . By the N and O properties, $send(y)_{s_y, r_1}$ occurs in ϵ' . Clearly, since $F_{1,x}(\epsilon') \stackrel{s_x}{\approx} F_{1,x}(\epsilon)$, $send(x)_{s_x, r_1}$ sends the same value in ϵ and ϵ' . Therefore, R_1 returns (x_1, y) .

In ϵ' , we denote the fragment that begins with $recv(m_y^{r_1})_{r_1, s_y}$ and ends with $send(y)_{s_y, r_1}$ by $F_{1,y}(\epsilon')$ to have ϵ' as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\epsilon') \circ F_{1,y}(\epsilon')$. Then, there exists an execution $\delta^{(k)}$ of $\mathcal{A}$, which is an extension of ϵ' . Clearly, $\delta^{(k)}$ can be written as $\sigma_0, a_1, \dots, a_k, \sigma_k \circ F_{1,x}(\delta^{(k)}) \circ F_{1,y}(\delta^{(k)}) \circ S(\delta^{(k)})$, where $S(\delta^{(k)})$ is the tail of the extended execution. Clearly, $F_{1,x}(\delta^{(k)}) \stackrel{s_x}{\approx} F_{1,x}(\epsilon')$. Therefore, $R_1(\delta^{(k)})$ returns x_1 in $\delta^{(k)}$, which implies $R_1(\delta^{(k)})$ must return (x_1, y_1) . $\square$

5.2 SNOW with C2C Communication

In this section, we show that SNOW is possible in the *multiple-writers single-reader* (MWSR) setting when client-to-client communication is allowed. In particular, we present an algorithm A , which has all SNOW properties in such setting. We consider a system that has $\ell \geq 1$ writers with ids $w_1, w_2 \dots w_\ell \in \mathcal{W}$, one reader r , and $k \geq 1$ servers with ids $s_1, s_2 \dots s_k \in \mathcal{S}$. Client-to-client communication is allowed. The pseudocode for algorithm A is presented in Pseudocode 4. We use keys to uniquely identify a WRITE transaction. A key $\kappa \in \mathcal{K}$ is defined as a pair (z, w) , where $z \in \mathbb{N}$, and $w \in \mathcal{W}$ is the id of a writer. $\mathcal{K}$ denotes the set of all possible keys. Also, with each transaction we associate a tag $t \in \mathbb{N}$.

State variables: (i) Each *writer* w stores a counter z corresponding to the number of WRITE transactions it has invoked so far, initially 0. (ii) The *reader* r has an ordered list of elements, $List$, as $(\kappa, (b_1, \dots, b_k))$, where $\kappa \in \mathcal{K}$ and $(b_1, \dots, b_k) \in \{0, 1\}^k$. Initially, $List = [(\kappa^0, (1, \dots, 1))]$, where $\kappa^0 \equiv (0, w_0)$, and w_0 is any place holder identifier for writer id. (iii) Each *server* $s_i \in \mathcal{S}$ stores a set variable $Vals$ with elements of key-value pairs $(\kappa, v_i) \in \mathcal{K} \times \mathcal{V}_i$. Initially, $Vals = \{(\kappa^0, v_i^0)\}$.

Writer steps: Any writer client, $w \in \mathcal{W}$, may invoke a WRITE transaction $W((o_{i_1}, v_{i_1}), (o_{i_2}, v_{i_2}), \dots, (o_{i_p}, v_{i_p}))$, comprising a set of write operations, where $I = \{i_1, i_2, \dots, i_p\}$ is some subset of p indices of $[k]$. We define the set $S_I \triangleq \{s_{i_1}, s_{i_2}, \dots, s_{i_p}\}$. This procedure consists of two consecutive phases: *write-value* and *info-reader*. In the *write-value* phase, w creates a key κ as $\kappa \equiv (z+1, w)$; and also increments the local counter z by one. Then it sends (WRITE-VAL, (κ, v_i)) to each server s_i in S_I , and awaits ACKs from each server in S_I . After receiving all ACKs, w initiates the *info-reader* phase during which it sends (INFO-READER, $(\kappa, (b_1, \dots, b_k))$) to r , where for any $i \in [k]$, b_i is a boolean variable, such that $b_i = 1$ if $s_i \in S_I$, otherwise $b_i = 0$. Essentially, such a $(k+1)$ -tuple identifies the set of objects that are updated during that WRITE transaction, i.e., if $b_i = 1$ then object o_i was updated during the execution of the WRITE transaction, otherwise $b_i = 0$. After w receives ACK from r it completes the WRITE transaction.

Reader steps: We use the same notations for I and S_I as above for the set of indices and corresponding servers, possibly different across transactions. The procedure $R(o_{i_1}, o_{i_2}, \dots, o_{i_p})$, for any READ transaction, is initiated at reader r , where $o_{i_1}, o_{i_2}, \dots, o_{i_p}$ denotes the subset of objects r intends to read. This procedure consists of only one phase, *read-value*, of communication between the reader and the servers in S_I . Here r sends the message (READ-VAL, κ_i) to each server $s_i \in S_I$, where the κ_i is the key in the tuple $(\kappa_i, (b_1, \dots, b_k))$ in $List$ located at index j^* such that $b_i = 1$ such that $i \in I$. After receiving the values $v_{i_1}, v_{i_2}, \dots, v_{i_p}$ from all servers in S_I , where $S_I \triangleq \{s_{i_1}, s_{i_2}, \dots, s_{i_p}\}$, the transaction completes by returning $(v_{i_1}, \dots, v_{i_p})$.

On receiving a message (INFO-READER, $(\kappa, (b_1, \dots, b_k))$) from any writer w , r appends $(\kappa, (b_1, \dots, b_k))$ to its $List$, and responds to w with ACK and $t_w = |List|$, i.e., number of elements in $List$. The order of the elements in $List$ corresponds to the order the WRITE transactions, the order of the incoming INFO-READER updates, as seen by the reader.

Server steps: The server responds to messages containing the tags WRITE-VAL and READ-VAL. The first procedure is used if a server s_i receives a message (WRITE-VAL, (κ, v_i)) from a writer w , it adds (κ, v_i) to its set variable $Vals$ and sends ACK back to w . The second procedure is used if s_i receives a message, i.e., (READ-VAL, κ_i), from r , then it responds with v_i such that (κ_i, v_i) is in its $Vals$.

A respects the SNOW properties as stated below.

Theorem 3. *Any well-formed and fair execution of A guarantees all of the SNOW properties.*

Pseudocode. 4 Steps at writer w , reader r and server s_i in A .

At writer w <i>State Variables at w:</i> $z \in \mathbb{N}$, initially 0 $W((o_{i_1}, v_{i_1}), \dots, (o_{i_p}, v_{i_p}))$ 2: <u><i>write-value:</i></u> $\kappa \leftarrow (z + 1, w)$ 4: $z \leftarrow z + 1$ $I \triangleq \{i_1, i_2, \dots, i_p\}$ 6: for $i \in I$ do Send (WRITE-VAL, (κ, v_{s_i})) to s_i	8: Await ACK from $s_i \forall i \in I$. <u><i>info-reader:</i></u> 10: for $i \in [k]$ do if $i \in I$ then 12: $b_i \leftarrow 1$ else 14: $b_i \leftarrow 0$ Send (INFO-READER, 16: $(\kappa, (b_1, \dots, b_k))$) to r Receive (ACK, t_w) from r
18: At reader r <i>State Variables at r:</i> $List$, a list of elements in $\mathcal{K} \times \{0, 1\}^k$, initially $[(\kappa^0, 1, \dots, 1)]$ 20: $R(o_{i_1}, o_{i_2}, \dots, o_{i_p})$ <u><i>read-value:</i></u> 22: $I \triangleq \{i_1, i_2, \dots, i_p\}$ for $i \in I$ do 24: $j^* \leftarrow \max_{1 \leq j \leq List } \{j : List[j].b_i = 1\}$ 26: $\kappa_i \leftarrow List[j^*].\kappa$ Send (READ-VAL, κ_i) to s_i	28: Await responses v_i from $s_i \forall i \in I$ Return $(v_{i_1}, v_{i_2}, \dots, v_{i_p})$ 30: <u>Response routines</u> <u>On recv (INFO-READER,</u> <u>32: $(\kappa, (b_1, \dots, b_k))$ from w:</u> $List \leftarrow List \oplus (\kappa, (b_1, \dots, b_k))$ 34: $/* \oplus$ for append $*/$ $tag \leftarrow List /* \cdot$ list size $*/$ 36: Send (ACK, tag) to w
At server s_i for any $i \in [k]$ 38: <i>State Variables:</i> $Vals \subset \mathcal{K} \times \mathcal{V}_i$, initially $\{(t_{key}^0, v_i^0)\}$ <u>On recv (WRITE-VAL, (κ, v)) from w:</u>	40: $Vals \leftarrow Vals \cup \{(\kappa, v)\}$ Send ACK to w. 42: <u>On recv (READ-VAL, κ) from r:</u> Send v s.t. $(\kappa, v) \in Vals$ to r

Proof. Below we show that A satisfies the SNOW properties.

S property: Let β be any fair execution of A and suppose all clients in β behave in a well-formed manner. Suppose β contains no incomplete transactions and let Π be the set of transactions in β . We define an irreflexive partial ordering ($\prec$) among the transactions in Π as follows: if ϕ and π are any two distinct transactions in Π then we say $\phi \prec \pi$ if either (i) $tag(\phi) < tag(\pi)$ or (ii) $tag(\phi) = tag(\pi)$ and ϕ is a WRITE and π is a READ. We will prove the S (strict-serializability) property of A by proving that the properties $P1$, $P2$, $P3$ and $P4$ of Lemma 20 hold for β .

P1: If π is a READ then since all READS are invoked by a single reader r and in a well-formed manner, therefore, there cannot be an infinite number of READS such that they all precede π (w.r.t. $\prec$). Now, suppose π is a WRITE. Clearly, from an inspection of the algorithm, $tag(\pi) \in \mathbb{N}$. From inspection of the algorithm, each WRITE increases the size of $List$, and the value of the tags are defined by the size of $List$. Therefore, there can be at most a finite number of WRITES such that can precede π (w.r.t. $\prec$) in β .

P2: Suppose ϕ and π are any two transactions in Π , such that, π begins after ϕ completes. Then we show that we cannot have $\pi \prec \phi$. Now, we consider four cases, depending on whether ϕ and π are READS or WRITES.

- (a) ϕ and π are WRITES invoked by writers w_ϕ and w_π , respectively. Since the size of $List$, in r , grows monotonically with each WRITE hence w_π receives the tag at least as high as $tag(\phi)$, so $\pi \not\prec \phi$.
- (b) ϕ is a WRITE, π is a READ transactions invoked by writer w_ϕ and r , respectively. Since the size of $List$, in r , grows monotonically, and because w_π invokes π after ϕ completes hence $tag(\pi)$ is at least as high as $tag(\phi)$, so $\pi \not\prec \phi$.
- (c) ϕ and π are READS invoked by reader r . Since the size of $List$, in r , grows monotonically, hence w_π invoked π after ϕ completes hence $tag(\pi)$ is at least as high as $tag(\phi)$, so $\pi \not\prec \phi$.
- (d) ϕ is a READ, π is a WRITE invoked by reader r and w_π , respectively. This case is simple because new values are added to $List$ only by writers, and $tag(\pi)$ is larger than the tag of ϕ and hence $\pi \not\prec \phi$.

P3: This is clear by the fact that any WRITE transaction always creates a unique tag and all tags are totally ordered since they all belong to $\mathbb{N}$

P4: Consider a READ ρ as $READ(o_{i_1}, o_{i_2}, \dots, o_{i_q})$, in β . Let the returned value from ρ be $\mathbf{v} \equiv (v_{i_1}, v_{i_2}, \dots, v_{i_q})$ such that $1 \leq i_1 < i_2 < \dots < i_q \leq k$, where value v_{i_j} corresponds to o_{i_j} . Suppose $tag(\rho) \in \mathbb{N}$ was created during some WRITE transaction, say ϕ , i.e., ϕ is the WRITE that added the elements in index $(tag(\rho) - 1)$ of $List$. Note that element in index 0 contains the initial value. Now we consider two cases:

Case $tag(\rho) = 1$. We know that it corresponds the initial default value v_i^0 at each sub-object o_i , and this equates to ρ returning the default initial value for each sub-object.

Case $tag(\rho) > 1$. Then we argue that there exists no WRITE transaction, say π , that updated object o_{i_j} , in β , such that, $\pi \neq \phi$ and ρ returns values written by π and $\phi \prec \pi \prec \rho$. Suppose we assume the contrary, which means $tag(\phi) < tag(\pi) < tag(\rho)$. The latter implies $tag(\phi) = tag(\pi)$ which is not possible because this contradicts the fact that for any two distinct WRITES $tag(\phi) \neq tag(\pi)$ in any execution of A .

N property: By inspection of algorithm A for the response steps of the servers to the reader.

O property: By inspection of the *read-value* phase: it consists of one round of communication between the reader and the servers, where the servers send only one version of the value of the object it maintains.

W property: By inspection of the WRITE transaction steps, and and that writers always get to complete the transactions they invoke. $\square$

6 No Prior Bounded Latency for SW

The SNOW work [15] claimed, after examining existing, work there existed only one system, Eiger [14], whose READ transactions had bounded latency—i.e., non-blocking and finish in three rounds—while providing the strongest guarantees—i.e., having properties W and S—because Eiger claimed that its READ transactions provide strict serializability within a datacenter. In this section, we correct this claim and show there were no existing algorithms that had bounded latency while providing the strongest guarantees by proving Eiger’s READ transactions are not strictly serializable.

The insight on why Eiger’s READ transactions are not strictly serializable is that Eiger uses logical timestamps, i.e., Lamport clocks, to track the ordering of operations, and logical clocks are




Figure 5: An example execution that shows Eiger’s READ transaction is not strictly serializable. w_1 , w_2 , and w_3 are three writes, where w_3 is issued after w_2 finishes. r_A and r_B are read operations from the same READ transaction R , which is concurrent with all three writes. Each number is a value of the logical clock on the machine (process) as a result of message exchange.

not able to identify the real-time ordering between operations that do not have causal relationship, i.e., operations from different processes. Strict serializability, however, requires the real-time ordering to be respected.

Figure 5 is an example execution, which is allowed by Eiger but violates strict serializability. Real time goes downwards in the diagram. Each number is a value of the Lamport clock on the machine (process) as a result of message exchange. Initially all processes have Lamport clock value 0, and no messages happened before the execution in Figure 5. A READ transaction $R = \{r_A, r_B\}$ reads values on S_A and S_B respectively. Due to the asynchronous nature of the network, r_B arrives on S_B before w_2 and r_A arrives after w_3 . Following the READ transaction algorithm of Eiger, r_A returns the value of w_3 and its valid logical duration, i.e., $[2, 3]$. Similarly, r_B returns the value of w_1 and its valid logical duration, i.e., $[2, 3]$. Because the two logical durations overlap, Eiger claims the combined values of r_A and r_B are consistent and accept them. However, because w_3 starts after w_2 finishes, w_3 is in real time after w_2 . By strict serializability, if a READ transaction sees the value of w_3 , then it must observe the effect of w_2 . Hence, R , which returns the values of w_1 and w_3 , violates strict serializability.

7 Condition for proving strict serializability

In this section, we describe the data types used in our algorithms. Also, we derive useful properties in proving *strict serializability* property (S property) of executions of the algorithms that implement objects of the data type.

7.1 Data type

In this section, we formally describe the data type, which we denote as $\mathcal{O}_T$, for the transaction processing systems considered in this paper. We assume there is a set of k objects, where k is some positive integer, and o_k denotes the k^{th} object. Object o_k stores a value from some non-empty domain V_k and supports two types of operations: *read* and *write* operations. A read operation on object o_k , denoted by $read(o_k)$, on completion returns the value stored in o_k . A write operation on o_k with some value v_k from V_k , denoted as $write(o_k, v_k)$, on completion, updates the value of o_k .

A WRITE transaction consists of a subset of p distinct write operations for a subset of p distinct objects, where $1 \leq p \leq k$. For example, a WRITE transaction with the set of operations $\{write(o_{i_1}, v_{i_1}), write(o_{i_2}, v_{i_2}) \cdots write(o_{i_p}, v_{i_p})\}$, means value v_{i_1} is to be written to object o_{i_1} , value v_{i_2} to object o_{i_2} , and so on. We denote such a WRITE transaction as $WRITE((o_{i_1}, v_{i_1}), (o_{i_2}, v_{i_2}), \dots, (o_{i_p}, v_{i_p}))$.

A READ transaction consisting of a set of read operations is denoted as $READ(o_{i_1}, o_{i_2}, \dots, o_{i_q})$, where $o_{i_1}, o_{i_2}, \dots, o_{i_q}$ is a set of distinct objects, q is any positive integer, $1 \leq q \leq k$, which upon completion returns the values $(v_{i_1}, v_{i_2}, \dots, v_{i_q})$, where $v_{i_j} \in V_{i_j}$ is the value returned by $read(o_{i_j})$, for any $i_j \in \{i_1, i_2, \dots, i_q\}$ and $1 \leq i_1 < i_2 < \dots < i_q \leq k$.

Formally, we define the value type used in this paper for a k object data-type as follows:

- (i) a tuple $(v_1, v_2, \dots, v_k) \in \prod_{i=1}^k V_i$, where V_i is a domain of values, for each $i \in \{1, \dots, k\}$;
- (ii) an initial value $v_i^0, v_i^0 \in V_i$ for each object o_i ;
- (iii) *invocations*: $READ(o_{i_1}, o_{i_2}, \dots, o_{i_p})$ and $WRITE((o_{i_1}, v_{i_1}), (o_{i_2}, v_{i_2}), \dots, (o_{i_p}, v_{i_p}))$, such that $v_{i_j} \in V_{i_j}$ for any $i_j \in \{i_1, i_2, \dots, i_p\}$ where $1 \leq i_1 < i_2 < \dots < i_p \leq k$ and p is some integer with $1 \leq p \leq k$;
- (iv) *responses*: a tuple $(v_1, v_2, \dots, v_k) \in \prod_{i=1}^k V_i$ and ok ; and
- (v) for any subset of p objects we define $f : \text{invocations} \times V \rightarrow \text{responses} \times V$, such that:
 - (a) $f(READ(o_{i_1}, o_{i_2}, \dots, o_{i_p}), (v_1, v_2, \dots, v_k)) = ((v_{i_1}, v_{i_2}, \dots, v_{i_p}), (v_1, v_2, \dots, v_k))$; and
 - (b) $f(WRITE((o_{i_1}, u_{i_1}), (o_{i_2}, u_{i_2}), \dots, (o_{i_p}, u_{i_p})), (v_1, v_2, \dots, v_k)) = (ack, (w_1, w_2, \dots, w_k))$, where for any i , $w_i = u_i$ if $i \in \{i_1, i_2, \dots, i_p\}$, and for all other values of i , $w_i = v_i$.

7.2 Strict-serializability in $\mathcal{O}_T$

Although the strict serializability property in transaction-processing systems is a well-studied topic, the specific setting considered in this paper is much simpler. Therefore, this allows us to derive simpler conditions to prove the safety of these algorithms. A wide range of transaction types and transaction processing systems are considered in the literature. For example, in [17], Papadimitriou defined the strict serializability conditions as a part of developing a theory for analyzing transaction processing systems. In this work, each transaction T consists of a set of write operations W , at individual objects, and a set of read operations R from individual objects, where the operations in W must complete before the operations in R execute. Other types of transaction processing systems allow nested transactions [9], where the transactions may contain sub-transactions [2] which may further contain a mix of read or write operations, or even child-transactions. In most transaction

processing systems considered in the literature, transactions can be *aborted* so as to handle failed transactions. As a result, the serializability theories are developed while considering the presence of aborts. However, in our system, we do not consider any abort, nor any client or server failures. A transaction in our system is either a set of independent writes or a set of reads with all the reads or writes in a transaction operating on different objects. Such simplifications allow us to formulate an equivalent condition for the execution of an algorithm to prove the S property of such algorithms while implementing an object of data type $\mathcal{O}_T$.

We note that an execution of a variable of type $\mathcal{O}_T$ is a finite sequence $\mathbf{v}_0, INV_1, RESP_1, \mathbf{v}_1, INV_2, RESP_2, \mathbf{v}_2, \dots, \mathbf{v}_r$ or an infinite sequence $\mathbf{v}_0, INV_1, RESP_1, \mathbf{v}_1, INV_2, RESP_2, \mathbf{v}_2, \dots$, where INV 's and $RESP$'s are invocations and responses, respectively. The $\mathbf{v}_i$'s are tuples of the form $(v_1, v_2, \dots, v_k) \in \prod_{i=1}^q V_i$, that corresponds to the latest values stored across the objects $o_1, o_2 \dots o_k$, and the values in $\mathbf{v}_0$ are the initial values of the objects. Any adjacent quadruple such as $\mathbf{v}_i, INV_{i+1}, RESP_{i+1}, \mathbf{v}_{i+1}$ is consistent with the f function for an object of type $\mathcal{O}_T$. Now, the safety property of such an object is a trace that describes the correct response to a sequence of INV 's when all the transactions are executed sequentially. The *strict serializability* of $\mathcal{O}_T$ says that each trace produced by an execution of $\mathcal{O}_T$ with concurrent transactions appears as some trace of $\mathcal{O}_T$. We describe this below in more detail.

Definition 7.1 (Strict-serializability). *Let us consider an execution β of an object of type $\mathcal{O}_T$, such that the invocations of any transaction at any client respects the well-formedness property. Let Π denote the set of complete transactions in β then we say β satisfies the strict-serializability property for $\mathcal{O}_T$ if the following are possible:*

- (i) *For every complete READ or WRITE transaction π we insert a point (serialization point) π_* between the actions $INV(\pi)$ and $RESP(\pi)$.*
- (ii) *We select a set Φ of incomplete transactions in β such that for each $\pi \in \Phi$ we select a response $RESP(\pi)$.*
- (iii) *For each $\pi \in \Phi$ we insert π_* somewhere after $INV(\pi)$ in β , and remove the INV for the rest of the incomplete transactions in β .*
- (iv) *If we assume for each $\pi \in \Pi \cup \Phi$ both $INV(\pi)$ and $RESP(\pi)$ to occur consecutively at π_* , with the interval of the transaction shrunk to π_* , then the sequence of transactions in this new trace is a trace of an object of data type $\mathcal{O}_T$.*

Now, we consider any automaton $\mathcal{B}$ that implements an object of type $\mathcal{O}_T$, and prove a result that serves us an equivalent condition for proving the strict serializability property of $\mathcal{B}$. Any trace property P of an automaton is a *safety property* if the set of executions in P is non-empty; *prefix-closed*, meaning any prefix of an execution in P is also in P ; and *limit-closed*, i.e., if $\beta_1, \beta_2, \dots$ is any infinite sequence of executions in P is such that β_i is prefix of β_{i+1} for any i , then the limit β of the sequence of executions $\{\beta_i\}_{i=0}^\infty$ is also in P . From Theorem 13.1 in [16], we know that the trace property, which we denote by P_{SC} , of any well-formed execution of $\mathcal{B}$ that satisfies the strict-serializability property is a safety property. Moreover, from Lemma 13.10 in [16] we can deduce that if every execution of $\mathcal{B}$ that is well-formed and failure-free, and also contains no incomplete transactions, satisfies P_{SC} , then any well-formed execution of $\mathcal{B}$ that can possibly have incomplete transactions is also in P_{SC} . Therefore, in the following lemma, which gives us an equivalent condition for the strict serializability property of an execution β , we consider only

executions without any incomplete transactions. The lemma is proved in a manner similar to Lemma 13.16 in [16], for atomicity guarantee of a single multi-reader multi-writer object.

Lemma 20. *Let β be an execution (finite or infinite) of an automaton $\mathcal{B}$ that implements an object of type $\mathcal{O}_T$, which consists of a set of k sub-objects. Suppose all clients in β behave in an well-formed manner. Suppose β contains no incomplete transactions and let Π be the set of transactions in β . Suppose there exists an irreflexive partial ordering ($\prec$) among the transactions in Π , such that,*

- P1 For any transaction $\pi \in \Pi$ there are only a finite number of transactions $\phi \in \Pi$ such that $\phi \prec \pi$;*
- P2 If the response event for π precedes the invocation event for ϕ in β , then it cannot be that $\phi \prec \pi$;*
- P3 If π is a WRITE transaction in Π and ϕ is any transaction in Π , then either $\pi \prec \phi$ or $\phi \prec \pi$; and*
- P4 A tuple $\mathbf{v} \equiv (v_{i_1}, v_{i_2}, \dots, v_{i_q})$ returned by a $READ(o_{i_1}, o_{i_2}, \dots, o_{i_q})$, where q is any positive integer, $1 \leq q \leq k$, is such that v_{i_j} $j \in \{1, \dots, q\}$ is written in β by the last preceding (w.r.t. $\prec$) WRITE transaction that contains a $write(o_{i_j}, *)$, or the initial value $v_{i_j}^0$ if no such WRITE exists in β .*

Then execution β is strictly serializable.

Proof. We discuss how to insert a serialization point $*_\pi$ in β for every transaction $\pi \in \Pi$. First, we add $*_\pi$ immediately after the latest of the invocations of π or $\phi \in \Pi$ such that $\phi \prec \pi$. We want to stress that any complete operation π in Π refers to an operation, with the invocation and response events, whereas π_* refers to a point in the executions. Note that according to condition P1 for π there are only finite number of such invocations in β , therefore, π_* is well-defined for any $\pi \in \Pi$. Now, since the order of the invocation events of the transactions in Π are already defined, therefore, the order of the corresponding set of serialization points are well-defined, except for the case when more than one serialization points are placed immediately after an invocation. In the case such multiple serialization points corresponding to an invocation we order these serialization points in accordance with the $\prec$ relation of the underlying transactions.

Next, we show that for any pair of transactions $\phi, \pi \in \Pi$ if $\phi \prec \pi$ then in β we have $*_\phi$ precedes $*_\pi$. Suppose $\phi \prec \pi$. By construction, each of π_* and ϕ_* appear immediately after some invocation, in β , of some transaction in Π . If both π_* and ϕ_* appear immediately after the same invocation, then since $\phi \prec \pi$, by construction of π_* , the serialization point π_* appears in β after ϕ_* . Also, if the invocations after which π_* and ϕ_* appear are distinct, then by construction of π_* the serialization point π_* appear after ϕ_* in β since $\phi \prec \pi$.

Next we argue that each $*_\pi$ serialization point for any $\pi \in \Pi$ is placed between the invocation $INV(\pi)$ and responses $RESP(\pi)$. By construction, $*_\pi$ is after $INV(\pi)$ in β . To show that $*_\pi$ is before $RESP(\pi)$ for the sake of contradiction assume that $*_\pi$ appears after $RESP(\pi)$. By construction, $*_\pi$ must be after $INV(\psi)$ for some $\psi \in \Pi$ and $\psi \neq \pi$, then by the condition of construction of π_* we have $\psi \prec \pi$. But from above $INV(\psi)$ occurs after $RESP(\pi)$, i.e., π completes before ψ is invoked which means, by property P2, we cannot have $\psi \prec \pi$, which is a contradiction.

Next, we show that if we were to shrink the transactions intervals to their corresponding serialization points, the resulting trace would be a trace of the underlying data type $\mathcal{O}_T$. In other

words, we show any READ $READ(o_{i_1}, o_{i_2}, \dots, o_{i_q})$ returns the values $(v_{i_1}, v_{i_2}, \dots, v_{i_q})$, such that each value v_{i_j} , $j \in [q]$, was written by the immediately preceding (w.r.t. the serialization points) WRITE that contained $write(o_{i_j}, v_{i_j})$ or the initial values if no such previous WRITE exists. Let us denote the set of WRITES that precedes (w.r.t. $\prec$) π by $\Pi_W^{\prec\pi}$, i.e., $\phi \in \Pi_W^{\prec\pi}$ ϕ is a write and $\phi \prec \pi$. By property *P3*, all transactions in $\Pi_W^{\prec\pi}$ are totally-ordered. By property *P4*, v_{i_j} must be the value updated by the most recent WRITE in $\Pi_W^{\prec\pi}$. Since the total order of serialization points are consistent with $\prec$ and hence the v_{i_j} corresponds to the write operation of a WRITE transaction with the most recent serialization point and contains a operation of type $write(o_{i_j}, *)$. $\square$

8 SNW + One Version, MWMR setting

Here present algorithm *B*, which satisfies SNW and "one-version" properties, in MWMR setting where a READ transaction must consist of one version of the data but, possibly, multiple communication trips between the reader and the servers. In *B*, the steps for the writers are shown in Pseudocode 5 and for readers and the servers are presented in Pseudocode 6. We assume a set of writers $\mathcal{W}$, a set of readers $\mathcal{R}$ and a set of $k \geq 1$ servers, $\mathcal{S}$, with ids $s_1, s_2 \dots s_k$ that stores the objects $o_1, o_2, \dots, o_k$, respectively. A key κ is defined as a pair (z, w) , where $z \in \mathbb{N}$ and $w \in \mathcal{W}$ the id of a writer. We use $\mathcal{K}$ to denote the set of all possible keys. In *B*, a key uniquely identifies some transaction. Also, with each transaction we associate a tag $t \in \mathbb{N}$.

In *B*, we designate one of the servers as coordinator, denote as s^* , for the transactions. The s^* maintains the order of the WRITE transactions and the objects that are updated during the WRITE transaction in the variable *List*.

State variables: Each of the writers and servers maintain a set of state variables as follows:

(i) At any writer w , there is a counter z to keep track of the number of WRITE transaction the writer has invoked, initially 0. (ii) At any server, s_i , for $i \in [k]$, there is a set variable *Vals* with elements that are key-value pairs $(\kappa, v_i) \in \mathcal{K} \times \mathcal{V}_i$. Initially, $Vals = \{(\kappa^0, v_i^0)\}$. A server also contains an ordered list variable *List* of elements as $(\kappa, (b_1, \dots, b_k))$, where $\kappa \in \mathcal{K}$ and $(b_1, \dots, b_k) \in \{0, 1\}^k$. Initially, $List = [(\kappa^0, (1, \dots, 1))]$, where $\kappa^0 \equiv (0, w_0)$, where w_0 is any place holder identifier string for writer id. The elements in *List* can be identified with an index, e.g., $List[0] = (\kappa^0, (1, \dots, 1))$. Essentially, a $(k + 1)$ -tuple $(\kappa, (b_1, \dots, b_k))$ in *List* corresponds to a WRITE transaction and identifies the set of objects that are updated during the WRITE transaction, i.e., if $b_i = 1$ then object o_i was updated during the WRITE transaction, otherwise $b_i = 0$.

Writer steps: A WRITE transaction updates a list of p objects $o_{i_1}, o_{i_2}, \dots, o_{i_p}$ with values $v_{i_1}, v_{i_2}, \dots, v_{i_p}$, respectively, is invoked at w via the procedure $W((o_{i_1}, v_{i_1}), \dots, (o_{i_p}, v_{i_p}))$. We use the notations: $I \triangleq \{i_1, i_2, \dots, i_p\}$ and $S_I \triangleq \{s_{i_1}, s_{i_2}, \dots, s_{i_p}\}$. This procedure consists of two phases: *write-value* and *update-coor*. During the *write-value* phase, w creates a new key κ as $\kappa \equiv (z + 1, w)$, where w identifies the writer; and also increments the local counter z by one. Then w sends $(WRITE-VAL, (\kappa, v_i))$ to each server in S_I , and awaits ACK from all servers in S_I . After receiving ACK from all servers in S_I , w initiates the *update-coor* phase where it sends $(UPDATE-COOR, (\kappa, (b_1, \dots, b_k)))$ to s^* , where for any $i \in [k]$, $b_i = 1$ if $s_i \in S_I$, otherwise $b_i = 0$, and completes then WRITE transaction after it receive a (ACK, t_w) from s^* .

Reader steps: We use the same notations for I and S_I as above but the indices can vary across transactions. The procedure $R(o_{i_1}, o_{i_2}, \dots, o_{i_p})$ can be initiated by some reader r , as a READ transaction, intending to read the values of subset $o_{i_1}, o_{i_2}, \dots, o_{i_p}$ of the objects. The

procedure consists of two consecutively executed phases of communication rounds between the r and the servers, viz., *get-tag-array* and *read-value*. During the phase *get-tag-array*, r sends s^* the message GET-TAG-ARR requesting the list of the latest added keys for each object. Once r receives a list of tags, such as, $(t_r, (\kappa_1, \kappa_2, \dots, \kappa_k))$ from s^* the phase completes. In the subsequent phase, *read-value*, r requests each server s_i in S_I by sending the message (READ-VAL, κ_i). After receiving the values $v_{i_1}, v_{i_2}, \dots, v_{i_p}$ from the servers in S_I , r completes the transaction by returning the tuple of values $(v_{i_1}, \dots, v_{i_p})$.

Server steps: When a server s_i receives a message of type (WRITE-VAL, (κ, v_i)) from a writer w then it adds (κ, v_i) to its set variable $Vals$ and sends ACK back to w .

If the coordinator s^* receives (UPDATE-COOR, $(\kappa, (b_1, \dots, b_k))$) from writer w , then it appends $(\kappa, (b_1, \dots, b_k))$ to its $List$, and responds with ACK and t_w (set to be the number of elements in the local list $List$) to w . The order of the elements in $List$ corresponds to the order the WRITE transactions, the order of the incoming UPDATE-COOR updates, as seen by s^* .

When s^* receives the message GET-TAG-ARR from r it responds with $(\kappa_1, \dots, \kappa_k)$ such that for each $i \in [k]$, κ_i is the key part of the $(k+1)$ -tuple that was modified last, i.e., $\kappa_i = List[j^*].\kappa$ such that $j^* \triangleq \max\{j : List[j].b_i = 1\}$, and $t_r, t_r \triangleq \max_{1 \leq j \leq |List|} \{j : List[j].b_i = 1 \wedge i \in I\}$. If any server s_i receives a message (READ-VAL, κ) from a reader r then it responds to r with the value v_i corresponding to key with value κ in $Vals$.

Theorem 4. *Any well-formed and fair execution of algorithm B satisfies the SNW and "one-version" properties.*

Proof. Below we show that algorithm B satisfies the SNoW properties.

S property: Let β be any fair execution of B and suppose all clients in β behave in a well-formed manner. Suppose β contains no incomplete transactions and let Π be the set of transactions in β . We define an irreflexive partial ordering ($\prec$) in Π as follows: if ϕ and π are any two distinct transactions in Π then we say $\phi \prec \pi$ if either (i) $tag(\phi) < tag(\pi)$ or (ii) $tag(\phi) = tag(\pi)$ and ϕ is a WRITE transaction and π is a READ transaction. Below we prove the S property of B by showing that properties P1, P2, P3 and P4 of Lemma 20 hold for β .

P1: Clearly, from an inspection of the algorithm, $tag(\pi) \in \mathbb{N}$. From inspection of the algorithm, each WRITE transaction increases the size of $List$, and the value of the tags are defined by the size of $List$. Therefore, there can be at most a finite number of WRITE transactions such that can precede π (w.r.t. $\prec$) in β . On the other hand, if π is a READ transaction then since all READ transactions are invoked by readers in a well-formed manner, and there are only finite number of readers therefore, there cannot be an infinite number of READ transactions such that they all precede π (w.r.t. $\prec$).

P2: Suppose ϕ and π are any two transactions in Π , such that, π begins after ϕ completes. Then we show that we cannot have $\pi \prec \phi$. Now, we consider four cases, depending on whether ϕ and π are READ transactions or WRITE transactions.

- (a) π and ϕ are WRITE transactions invoked by writers w_π and w_ϕ , respectively. Since the size of $List$, in s^* grows monotonically due to each WRITE transaction hence w_π receives the tag from s^* at least as high as $tag(\phi)$, so $\pi \not\prec \phi$.
- (b) π is a READ transaction, ϕ is a WRITE transaction invoked by reader r_π and writer w_ϕ , respectively. Since the size of $List$, in s^* , grows monotonically, because r_π invokes π after ϕ completes hence $tag(\pi)$ is at least as high as $tag(\phi)$, so $\pi \not\prec \phi$.

- (c) π and ϕ are both READ transactions invoked by readers r_π and r_ϕ , respectively. Since the size of $List$, in s^* , grows monotonically, because w_π invokes π after ϕ completes hence $tag(\pi)$ is at least as high as $tag(\phi)$, so $\pi \not\prec \phi$.
- (d) π is a WRITE transaction, ϕ is a READ transaction invoked by writer w_π and reader r_ϕ , respectively. This case is simple because new values are added to $List$, in s^* , only by writers, and $tag(\pi)$ has to be larger than the tag of ϕ and hence $\pi \not\prec \phi$.

P3: This is from the fact that any WRITE transaction always creates a unique tag and all tags are totally ordered since they all belong to $\mathbb{N}$

P4: Consider a READ transaction ρ as $READ(o_{i_1}, o_{i_2}, \dots, o_{i_q})$, in β . Let the returned value from ρ be $\mathbf{v} \equiv (v_{i_1}, v_{i_2}, \dots, v_{i_q})$ such that $1 \leq i_1 < i_2 < \dots < i_q \leq k$, where value v_{i_j} corresponds to o_{i_j} . Suppose $tag(\rho) \in \mathbb{N}$ was created during some WRITE transaction, say ϕ , i.e., ϕ is the WRITE transaction that added the elements in index $(tag(\rho) - 1)$ of $List$ at the coordinator s^* . Note that element in index 0 contains the initial value. Now we consider two cases:

Case $tag(\rho) = 1$. We know that it corresponds the initial default value v_i^0 at each sub-object o_i , and this equates to ρ returning the default initial value for each sub-object.

Case $tag(\rho) > 1$. Then we argue that there exists no WRITE transaction, say π , that updated object o_{i_j} , in β , such that, $\pi \neq \phi$ and ρ returns values written by π and $\phi \prec \pi \prec \rho$. Suppose we assume the contrary, which means $tag(\phi) < tag(\pi) < tag(\rho)$. The latter implies $tag(\phi) = tag(\pi)$ which is not possible because this contradicts the fact that for any two distinct WRITE transactions $tag(\phi) \neq tag(\pi)$ in any execution of B .

N, o and W properties: Evident from an inspection of the algorithm. □

Pseudocode. 5 Protocol for writer w in algorithms B and C .

At writer w <i>State Variables:</i> $z \in \mathbb{N}$, initially 0 $\mathbf{W}((i_1, v_{i_1}), \dots, (i_p, v_{i_p}))$ $I \triangleq \{i_1, i_2, \dots, i_p\}$ 3: <u><i>write-value:</i></u> $\kappa \leftarrow (z + 1, w)$ $z \leftarrow z + 1$ 6: for $i \in I$ do Send (WRITE-VAL, (κ, v_{s_i})) to s_i 9: Await ACK from servers in S_I.	<u><i>update-coor:</i></u> for $i \in [k]$ do 12: if $i \in I$ then $b_i \leftarrow 1$ else 15: $b_i \leftarrow 0$ Send (UPDATE-COOR, $(\kappa,$ $(b_1, \dots, b_k))$) to s^* 18: Receive (ACK, t_w) from s^*
---	---

9 SNW + One Round, MWMR setting

Here, we present algorithm C which satisfies SNW and "one-round" properties in the MWMR setting, where a *READ* consists one round of communications between the reader and the servers but servers may respond with multiple versions of the data. The notation for the writers, servers and tag are similar to algorithm B . Pseudocodes 5 and 7 show the steps for the writers, and the readers and the servers, respectively. We designate a server as the coordinator, denote as s^* .

State variables: The state variables are similar to B .

Pseudocode. 6 Protocols reader r and server s_i in alg. B .

At reader r $R(o_{i_1}, o_{i_2}, \dots, o_{i_p})$ $I \triangleq \{i_1, i_2, \dots, i_p\}$ <u>get-tag-array:</u> 3: Send (GET-TAG-ARR) to s^* Receive $(t_r, (\kappa_1, \kappa_2, \dots, \kappa_k))$ from s^*		<u>read-value:</u> 6: for $i \in I$ do Send (READ-VAL, κ_i) to s_i Await responses as $v_i \forall s_i \in S$ 9: Return $(v_{i_1}, v_{i_2}, \dots, v_{i_p})$	
At server s_i for any $i \in [k]$ State Variables: $Vals \subset \mathcal{K} \times \mathcal{V}_i$, initially $\{(\kappa^0, v_i^0)\}$ $List$, list of $\mathcal{K} \times \{0, 1\}^k$, initially $[(\kappa^0, (1, \dots, 1))]$		Send (ACK, tag) to w 24: On recv (READ-VAL, κ) from r : Send v_i s.t. $(\kappa, v) \in Vals$ to r	
12: On recv (WRITE-VAL, (κ, v)) from w : $Vals \leftarrow Vals \cup \{(\kappa, v)\}$ 15: Send ACK to w . On recv (UPDATE-COOR, $(\kappa, (b_1, \dots, b_k))$) from w : 18: $List \leftarrow List \oplus (\kappa, (b_1, \dots, b_k))$ // $\oplus$ for append 21: $tag \leftarrow List // \cdot $ list size		/* used only by s^* */ On recv GET-TAG-ARR from r : for $i \in [k]$ do 27: $j^* \leftarrow \max\{j : List[j].b_i = 1\}$ $\kappa_i \leftarrow List[j^*].\kappa$ 30: $t_r \triangleq \max_{1 \leq j \leq List } \{j : List[j].b_i = 1 \wedge i \in I\}$ Send $(t_r, (\kappa_1, \kappa_2, \dots, \kappa_k))$ to r	

Pseudocode. 7 Protocols for reader r & server s_i in alg. C .

At reader r $R(o_{i_1}, o_{i_2}, \dots, o_{i_p})$ $I \triangleq \{i_1, i_2, \dots, i_p\}$ <u>read-values-and-tags:</u> 3: Send (GET-TAG-ARR) to s^* for $i \in I$ do		Send (READ-VALS) to s_i 6: Recv $(t_r, (\kappa_1, \kappa_2, \dots, \kappa_k))$ from s^* Recv. $Vals_i$ from $\forall s_i \in S_I$ Return $(v_{i_1}, v_{i_2}, \dots, v_{i_p})$ 9: s.t. $(\kappa_j, v_j) \in Vals_j, j \in I$	
At server s_i for any $i \in [k]$ State Variables: $Vals \subset \mathcal{K} \times \mathcal{V}_i$, initially $\{(\kappa^0, v_i^0)\}$ $List$, a list of $\mathcal{K} \times \{0, 1\}^k$, initially $[(\kappa^0, (1, \dots, 1))]$		Send (ACK, tag) to w 24: On recv (READ-VALS) from r : Send $Vals$ to r	
12: On recv (WRITE-VAL, (κ, v)) from w : $Vals \leftarrow Vals \cup \{(\kappa, v)\}$ 15: Send ACK to writer w . On recv (UPDATE-COOR, $(\kappa, (b_1, \dots, b_k))$) from w : 18: $List \leftarrow List \oplus (\kappa, (b_1, \dots, b_k))$ // $\oplus$ for append 21: $tag \leftarrow List /* \cdot $ list size */		/* used only by s^* */ On recv GET-TAG-ARR from r : for $i \in [k]$ do 27: $j^* \leftarrow \max\{j : List[j].b_i = 1\}$ $\kappa_i \leftarrow List[j^*].\kappa$ 30: $t_r \triangleq \max_{1 \leq j \leq List } \{j : List[j].b_i = 1 \wedge i \in I\}$ Send $(t_r, (\kappa_1, \kappa_2, \dots, \kappa_k))$ to r	

Writer steps: WRITE transaction is similar to algorithm B .

Reader steps: The step $R(o_{i_1}, o_{i_2}, \dots, o_{i_p})$ can be initiated by some reader r intending to read the values of subset $o_{i_1}, o_{i_2}, \dots, o_{i_p}$ of the objects. Denote $I \triangleq \{i_1, i_2, \dots, i_p\}$ and $S_I \triangleq$

$\{s_{i_1}, s_{i_2}, \dots, s_{i_p}\}$. The procedure consists of only one phase of communication round between the r and the servers, called *read-values-and-tags*. During *read-values-and-tags*, r sends s^* the message GET-TAG-ARR requesting the list of the latest added keys for each object, and also sends requests (READ-VALS) each server s_i in S_I . Note that if s^* is also one of the servers in S_I then the GET-TAG-ARR and READ-VALS messages to s^* can be combined to create one message; however, we keep them separate for clarity of presentation. Once r receives a list of tags, such as, $(t_r, (\kappa_1, \kappa_2, \dots, \kappa_k))$ from s^* and the set of $Vals_i$ from each $s_i \in S_I$ then r returns the values $v_{i_1}, v_{i_2}, \dots, v_{i_p}$ such that $(\kappa_j, v_j) \in Vals_j, j \in \{1, \dots, p\}$, and completes the *READ*.

Server steps: When a server s_i receives a message (WRITE-VAL, (κ, v_i)) from a writer w or s^* , receives (UPDATE-COOR, $(\kappa, (b_1, \dots, b_k))$) from writer w or receives a message as GET-TAG-ARR r the steps are similar to those in B . On the other hand, if any server s_i receives a message (READ-VALS) from a reader r then it responds to r with $Vals$. The following result states that C respects SNW and “one-round” properties.

Theorem 5. *Any well-formed and fair execution of C , in the MWMM setting satisfies the SNW and “one-round” properties.*

Proof. Below we show that algorithm C satisfies the SN $\bar{o}$ W properties.

S property: Let β be any fair execution of B and suppose all clients in β behave in a well-formed manner. Suppose β contains no incomplete transactions and let Π be the set of transactions in β . We define an irreflexive partial ordering ($\prec$) in Π as follows: if ϕ and π are any two distinct transactions in Π then we say $\phi \prec \pi$ if either (i) $tag(\phi) < tag(\pi)$ or (ii) $tag(\phi) = tag(\pi)$ and ϕ is a *WRITE* and π is a *READ*. Below we prove the *S* property of B by showing that properties $P1$, $P2$, $P3$ and $P4$ of Lemma 20 hold for β . The properties $P1$ - $P4$ can be proved to hold in a manner very similar to algorithm B (Section 8). Therefore, we avoid repeating them.

N, $\bar{o}$ and W properties: Evident from an inspection of the algorithm. □

10 Conclusion

We revisited the SNOW Theorem and when it is possible for READ transactions to have the same latency as simple reads. We provided a new and more rigorous proof of the original result. We also closed several open questions that were either explicitly posed by the original work or that emerged from our careful analysis. We found that READ transactions can match the latency of simple reads when client-to-client communication is allowed in MWSR setting. We found that they cannot and must have higher worst-case latency when client-to-client communication is disallowed or there are at least two readers. We also presented the first algorithms that provide bounded worst-case latency for read-only transactions in strictly serializable systems with WRITE transactions.

Acknowledgment

This work was supported by NSF awards CNS-1824130, CCF-2003830, CCF-0939370 and NSF CCF-1461559.

References

- [1] Hagit Attiya and Jennifer L. Welch. Sequential consistency versus linearizability. *ACM Trans. Comput. Syst.*, 12(2):91–122, 1994.
- [2] P.A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley Longman Publishing, 1987.
- [3] Eric A. Brewer. Towards robust distributed systems. In *Proc. Principles of Distributed Computing*, Jul 2000.
- [4] Nathan Bronson and et. al. TAO: Facebook’s distributed data store for the social graph. In *USENIX Annual Technical Conference (USENIX ATC 13)*., pages 49–60, 2013.
- [5] Phil Dixon. Shopzilla site redesign: We get what we measure. Velocity Conf. Talk, 2009.
- [6] James C. Corbett et. al. Spanner: Google’s globally-distributed database. In *Proc. OSDI*, Oct 2012.
- [7] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. In *Proc. Prin. of Database Sys*, 1983.
- [8] Seth Gilbert and Nancy Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. In *ACM SIGACT News*, Jun 2002.
- [9] J. Gray and A. Reuter. *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann, San Mateo, CA, 1993.
- [10] M. P. Herlihy and Jeannette M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 12(3):463–492, 1990.
- [11] Kishori M Konwar, Wyatt Lloyd, Haonan Lu, and Nancy Lynch. The snow theorem revisited, 2018.
- [12] Greg Linden. Make data useful. Stanford CS345 Talk, 2006.
- [13] Richard J. Lipton and Jonathan S. Sandberg. PRAM: A scalable shared memory. Technical Report TR-180-88, Princeton Univ., Dept. Comp. Sci., 1988.
- [14] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Anderson. Stronger Semantics for Low-Latency Geo-Replicated Storage. In *Proc. NSDI*, Apr 2013.
- [15] H. Lu, C. Hodsdon, K. Ngo, S Mu, and W. Lloyd. The SNOW theorem and latency-optimal read-only transactions. In *12th USENIX Symp. on Operating Sys. Design and Implementation (OSDI 16)*, pages 135–150, Savannah, GA, 2016.
- [16] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Pub., 1996.
- [17] Christos H. Papadimitriou. The serializability of concurrent database updates. *Journal of ACM*, 26(4):631–653, 1979.
- [18] E. Schurman and J. Brutlag. The user and business impact of server delays, additional bytes, and HTTP chunking in web search. Velocity Conf. Talk, 2009.

A Algorithm Specification with I/O Automata

We model a distributed algorithm using the I/O Automata [16]. Here we limit our discussion of I/O Automata to the relevant concepts, but for a detailed account the reader should refer to [16]. An algorithm is a composition $\mathcal{A}$ of a set of *automata* where each automaton A_i corresponds to a process (e.g., client or server) or a communication channel in the system. A_i is defined in terms of a set of deterministic transition functions $trans(A_i)$ (also called *actions*), which can be thought of as the algorithmic steps of A_i ; and a set of states $states(A_i)$. An execution of $\mathcal{A}$ is a sequence of alternating states and actions of A , $\sigma_0, a_1, \sigma_1, a_2, \sigma_2, \dots, \sigma_n$. A state change, called a *step*, is a 3-tuple $(\sigma_i, a_i, \sigma_{i+1})$, with $\sigma_i, \sigma_{i+1} \in states(A_i)$ and $a_i \in trans(A_i)$. The set of input actions is denoted by $in(A_i)$, e.g., $a_i \in in(A_i)$ is an input action if it receives a message. The set of output actions is denoted by $out(A_i)$. The input and output actions are also called external actions, denoted by $ext(A_i)$ and $in(A_i) \cup out(A_i) = ext(A_i)$. If an action $a_i \notin ext(A_i)$, then a_i is an internal action. Communications between any two automata A_i and A_j is modeled by using channel automata $Channel_{i,j}$ for sending messages from A_i to A_j ; and $Channel_{j,i}$ for sending message from A_j to A_i . When A_i sends some message m to A_j the following sequence of actions occur: $send(m)_{i,j}$ occurs at A_i , then $send(m)_{i,j}$ followed by $recv(m)_{i,j}$ occur at $Channel_{i,j}$; then finally, A_j receives m via the action $recv_{i,j}(m)$. In our model, the communication channels are simple because we assume reliable communication between each pair of processes. Therefore, we ignore the actions in the $Channel_{i,j}$ and instead say $send(m)_{i,j}$ occurs at A_i and then A_j receives m via the action $recv_{i,j}(m)$. An execution fragment α can be either finite, i.e., having finite states, or infinite. If α is a finite execution and β is an execution fragment, such that β starts with the final state of α then we use $\alpha \circ \beta$ to denote the concatenation of α and β . If ϵ and ϵ' are two execution fragments, such that they have the same sequence of states at automaton A_i , i.e., $\epsilon|A_i = \epsilon'|A_i$, then ϵ and ϵ' are indistinguishable at A_i , denoted by $\epsilon \stackrel{A_i}{\sim} \epsilon'$. When the context is clear, we simply use $\epsilon \sim \epsilon'$.

For any execution of $\mathcal{A}$, $\sigma_0, a_1, \dots, a_k, \sigma_k, \dots$, where σ 's and a 's are states and actions, we use the notation $a_1, \dots, a_k, \dots$ that shows only the actions while leaving out the states to simplify notation.

B Some Useful I/OA results

Below we add some useful theorems are useful related to executions of a composed I/O Automata [16].

Theorem 6. *Let $\{A_i\}_{i \in I}$ be a compatible collection of automata and let $A = \Pi_{i \in I} A_i$. Suppose α_i is an execution of A_i for every $i \in I$, and suppose β is a sequence of actions in $ext(A)$ such that $\beta|A_i = trace(\alpha_i)$ for every $i \in I$. Then there is an execution α of A such that $\beta = trace(\alpha)$ and $\alpha_i = \alpha|A_i$, for every $i \in I$.*

Theorem 7. *Let A be any I/O automaton.*

1. *If α is a finite execution of A , then there is a fair execution of A that starts with α .*
2. *If β is a finite trace of A , then there is a fair trace of A that starts with β .*
3. *If α is a finite execution of A and β is any finite or infinite sequence of input actions of A , then there is a fair execution $\alpha \circ \alpha'$ of A such that the sequence of input actions in α' is exactly β .*

4. If β is a finite trace of A and β' is any finite or infinite sequence of input actions of A , then there is a fair execution $\alpha \circ \alpha'$ of A such that $\text{trace}(\alpha) = \beta$ and such that the sequence of input actions in α' is exactly β' .

The following useful claim is adopted from Chapter 16 of [16].

Claim. Suppose we have an automaton $A = \Pi_{i=1}^k A_i$ where A is composed of the compatible collection of automata A_i , where $i \in \{1, \dots, k\}$. Let β be a fair trace of A then we define an irreflexive partial order $\rightarrow_\beta$ on the actions of β as follows. If π and ϕ are events in β , with π preceding ϕ , then we say ϕ depends on π , which we denote as $\pi \rightarrow_\beta \phi$, if one of the following holds:

1. π and ϕ are actions at the same automaton;
2. π is some $\text{send}(\cdot)_{j,i}$ at some A_j and ϕ is some $\text{recv}(\cdot)_{j,i}$ at A_i ; and
3. π and ϕ are related by a chain of the relations of items 1. and 2.

Then if γ is a sequence obtained by reordering the events in β while preserving the $\rightarrow_\beta$, then γ is also a fair trace of A .

Theorem 8 ([16]). Let $\{A_i\}_{i \in I}$ be a compatible collection of automata and let $A = \Pi_{i=1}^k A_i$. Suppose α_i is a fair execution of A_i for every $i \in I$, and suppose β is a sequence of actions in $\text{ext}(A)$ such that $\beta|_{A_i} = \text{trace}(\alpha_i)$ for every $i \in I$. Then there is a fair execution α of A such that $\beta = \text{trace}(\alpha)$ and $\alpha_i = \alpha|_{A_i}$ for every $i \in I$.