

Optimal Resilience in Systems That Mix Shared Memory and Message Passing

Hagit Attiya 

Department of Computer Science, Technion, Haifa, Israel
hagit@cs.technion.ac.il

Sweta Kumari

Department of Computer Science, Technion, Haifa, Israel
sweta@cs.technion.ac.il

Noa Schiller

Department of Computer Science, Technion, Haifa, Israel
noa.schiller@cs.technion.ac.il

Abstract

We investigate the minimal number of failures that can *partition* a system where processes communicate both through shared memory and by message passing. We prove that this number precisely captures the resilience that can be achieved by algorithms that implement a variety of shared objects, like registers and atomic snapshots, and solve common tasks, like randomized consensus, approximate agreement and renaming. This has implications for the *m&m-model* of [5] and for the hybrid, cluster-based model of [28,31].

2012 ACM Subject Classification Theory of computation → Distributed computing models; Theory of computation → Concurrent algorithms; Computing methodologies → Distributed algorithms

Keywords and phrases fault resilience, m&m model, cluster-based model, randomized consensus, approximate agreement, renaming, register implementations, atomic snapshots

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2020.16

Related Version A full version of the paper is available at <http://arxiv.org/abs/2012.10846>.

Funding This research was supported by ISF grant 380/18.

Acknowledgements We thank Vassos Hadzilacos, Xing Hu, Sam Toueg and the anonymous reviewers for helpful comments.

1 Introduction

Some distributed systems combine more than one mode of communication among processes, allowing them both to send messages among themselves and to access shared memory. Examples include recent technologies such as *remote direct memory access (RDMA)* [2–4], *disaggregated memory* [30], and *Gen-Z* [1]. In these technologies, the crash of a process does not prevent access to its shared memory by other processes. Under these technologies, it is infeasible to share memory among a large set of processes, so memories are shared by smaller, strict subsets of processes.

Systems mixing shared memory and message passing offer a major opportunity since information stored in shared variables remains available even after the failure of the process who stored it. Mixed systems are expected to withstand more process failures than pure message-passing systems, as captured by the *resilience* of a problem – the maximal number of failures that an algorithm solving this problem can tolerate. This is particularly the case in an *asynchronous* system. At one extreme, when all processes can access the same shared memory, many problems can be solved even when all processes but one fail. Such



© Hagit Attiya, Sweta Kumari, and Noa Schiller;
licensed under Creative Commons License CC-BY

24th International Conference on Principles of Distributed Systems (OPODIS 2020).

Editors: Quentin Bramas, Rotem Oshman, and Paolo Romano; Article No. 16; pp. 16:1–16:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

wait-free algorithms exist for implementing shared objects and solving tasks like randomized consensus, approximate agreement and renaming. At the other extreme, when processes only communicate by message passing, the same problems require that at least a majority of processes do not fail [7,8,18]. Thus, typically, shared-memory systems are $(n-1)$ -resilient, and pure message-passing systems are $\lfloor (n-1)/2 \rfloor$ -resilient, where n is the number of processes.

The resilience in systems that mix shared memory and message passing falls in the intermediate range, between $\lfloor (n-1)/2 \rfloor$ and $n-1$. It is, however, challenging to solve specific problems with the best-possible resilience in a particular system organization: the algorithm has to coordinate between non-disjoint sets of processes that have access to different regions of the shared memory. On the other hand, bounding the resilience requires to take into account the fact that processes might be able to communicate *indirectly* through shared memory accesses of third-party processes.

This paper explores the *optimal* resilience in systems that provide message-passing support between all pairs of processes, and access to shared memory between subsets of processes. We do this by studying the minimal number of failures that can *partition* the system, depending on its structure, i.e., how processes share memory with each other. We show that the partitioning number exactly characterizes the resilience, that is, a host of problems can be solved in the presence of $< f$ crash failures, if and only if f is the minimal number of failures that partition the system.

A key step is to focus on the implementation of a *single-writer multi-reader register* shared among all processes, in the presence of f crash failures. A read or a write operation takes $O(1)$ round-trips, and requires $O(n)$ messages. Armed with this implementation, well-known shared-memory algorithms can be employed to implement other shared objects, like *multi-writer multi-reader registers* and *atomic snapshots*, or to solve fundamental problems, such as *randomized consensus*, *approximate agreement* and *renaming*. Because the register implementation is efficient, these algorithms inherit the good efficiency of the best-known shared-memory algorithm for each of these problems.

Going through a register simulation, instead of solving consensus, approximate agreement or renaming from scratch, does not deteriorate their resilience. One of our key contributions is to show that the resilience achieved in this way is optimal, by proving that these problems cannot be solved in the presence of f crash failures, if f failures can partition the system.

We consider memories with access restrictions and model mixed systems by stating which processes can read from or write to each memory. (Note that every pair of processes can communicate using messages.) Based on this concept, we define f_{opt} to be the largest number of failures that do not partition the system. We prove that f -resilient registers and snapshot implementations, and f -resilient solutions to randomized consensus, approximate agreement and renaming, exist if and only if $f \leq f_{opt}$.

One example of a mixed model is the *message-and-memory* model [5], in short, the *m&m* model. In the *general* m&m model [5], the shared-memory connections are defined by (not necessarily disjoint) subsets of processes, where each subset of processes share a memory. Most of their results, however, are for the *uniform* m&m model, where shared-memory connections can be induced by an undirected graph, whose vertices are the processes. Each process has an associated shared memory that can be accessed by all its neighbors in the *shared-memory graph* (see Section 5). They present bounds on the resilience for solving randomized consensus in the uniform model. Their algorithm is based on Ben-Or's exponential algorithm for the pure message-passing model [15]. The algorithm terminates if the nonfaulty processes and their neighbors (in the shared-memory graph) are a majority of the processes. They also prove an upper bound on the number of failures a randomized consensus algorithm can

tolerate in the uniform m&m model. We show that in the uniform m&m model, this bound is equivalent to the partitioning bound (f_{opt}) proved in our paper (Theorem 17 in Section 5). We further show that this bound *does not match the resilience of their algorithm*, whose resilience is strictly smaller than f_{opt} , for some shared-memory graphs.

In the special case where the shared memory *has no access restrictions*, our model is dual to the general m&m model, i.e., it captures the same systems as the general m&m model. However, rather than listing which processes can access a memory, we consider the flipped view: we consider for each process, the memories it can access. We believe this makes it easier to obtain some extensions, for example, for memories with access restrictions.

Hadzilacos, Hu and Toueg [23] present an implementation of a SWMR register in the general m&m model. The resilience of their algorithm is shown to match the maximum resilience of an SWMR register implementation in the m&m model. Our results for register implementations are adaptations of their results. For the general m&m model specified by the set of process subsets L , they define a parameter f_L and show that it is the maximum number of failures tolerated by an algorithm implementing a SWMR register [23] or solving randomized consensus [24]. For memories without access restrictions, f_L is equal to f_{opt} . Their randomized consensus algorithm is based on the simple algorithm of [6] and inherits its exponential expected step complexity.

Another example of a model that mixes shared memory and message passing is the hybrid model of [28,31]. In this model, which we call *cluster-based*, processes are partitioned into disjoint *clusters*, each with an associated shared memory; all processes in the cluster (and only them) can read from and write to this shared memory. Two randomized consensus algorithms are presented for the cluster-based model [31]. Their resilience is stated as an operational property of executions: the algorithm terminates if the clusters of responsive processes contain a majority of the processes. We prove (Lemma 19 in Section 6) that the optimal resilience we state in a closed form for the cluster-based model is equal to their operational property.

Our model is general and captures all these models within a single framework, by precisely specifying the shared-memory layout. The tight bounds in this general model provide the exact resilience of any system that mix shared memory and message passing.

2 Modelling Systems that Mix Shared Memory and Message Passing

We consider n asynchronous processes p_1, \dots, p_n , which communicate with each other by sending and receiving messages, over a complete communication network of asynchronous reliable links. In addition, there are m shared memories $M = \{\mu_1, \dots, \mu_m\}$, which can be accessed by subsets of the processes. A memory $\mu \in M$ has access restrictions, where R_μ denotes all the processes that can read from the memory and W_μ denotes all the processes that can write to the memory. The set of memories a process p can read from is denoted R_p , i.e., $R_p = \{\mu \in M : p \in R_\mu\}$. The set of memories p can write to is denoted W_p , i.e., $W_p = \{\mu \in M : p \in W_\mu\}$. We assume the network allows nodes to send the same message to all nodes; message delivery is FIFO. A process p can *crash*, in which case it stops taking steps; messages sent by a crashed process may not be delivered at their recipients. We assume that the shared memory does not fail, as done in prior work [5, 23, 28, 31].

A *configuration* C is a tuple with a state for each process, a value for each shared register, and a set of messages in transit (sent but not received) between any pair of processes. A *schedule* is a sequence of process identifiers. For a set of processes P , a schedule is *P-free* if no process from P appears in the schedule; a schedule is *P-only* if only processes from P appear in the schedule. An *execution* α is an alternating sequence of configurations and

events, where each event is a step by a single process that takes the system from the preceding configuration to the following configuration. In a step, a process either accesses the shared memory (read or write) or receives and sends messages. Additionally, a step may involve the invocation of a higher-level operation. A schedule is associated with the execution in a natural way; this induces notions of P -free and P -only executions.

If there is a shared memory $\mu \in M$ that p can read from and q can write to, then we denote $p \rightarrow q$. If $p \rightarrow q$ and $q \rightarrow p$, then we denote $p \leftrightarrow q$. Since a process can read what it writes to its local memory, this relation is *reflexive*, i.e., for every process p , $p \rightarrow p$. Let P and Q be two sets of processes. Denote $P \rightarrow Q$ if some process $p \in P$ can read what a process $q \in Q$ writes, i.e., $p \rightarrow q$. If $P \rightarrow Q$ and $Q \rightarrow P$, then we denote $P \leftrightarrow Q$.

► **Definition 1.** *A system is f -partitionable if there are two sets of processes P and Q , both of size $n - f$, such that $P \not\leftrightarrow Q$. Namely, the failure of f processes can partition (disconnect) two sets of $n - f$ processes. Denote by f_{opt} the largest integer f such that $P \leftrightarrow Q$, for every pair of sets of processes P and Q , each of size $n - f$.*

Clearly, a system is f -partitionable if and only if $f > f_{opt}$. Note that $f_{opt} \geq \lfloor (n - 1)/2 \rfloor$. In the pure message-passing model, $p \rightarrow q$ if and only if $p = q$; hence, $f_{opt} = \lfloor (n - 1)/2 \rfloor$.

The special case of shared memory *without access restrictions* is when for every memory $\mu \in M$, $R_\mu = W_\mu$, and all processes that can read from a memory can also write to it. In this case, the \rightarrow relation is *symmetric*, i.e., for every pair of processes p and q , if $p \rightarrow q$ then $q \rightarrow p$. Therefore, for every two processes p and q , $p \leftrightarrow q$. Later, we discuss two models without access restrictions, the m&m model and the cluster-based model.

For a set of processes P , \vec{P} are the processes that some process in P can read what they write to the shared memory, i.e., $\vec{P} = \{q : \exists p \in P, p \rightarrow q\}$. f_{maj} is the largest integer f such that for every set P of $n - f$ processes, $|\vec{P}| > \lfloor n/2 \rfloor$. That is, f_{maj} is the largest number of failures that still allows the remaining (nonfaulty) processes to communicate with a majority of the processes. It is simple to see that $f_{opt} \leq f_{maj}$. The converse direction does not necessarily hold, as discussed for the m&m model and the cluster-based model.

3 Necessary and Sufficient Condition for Implementing a Register

This section shows that a register can be implemented in the presence of f failures, if and only if the system is not f -partitionable, that is, $f \leq f_{opt}$. This is an adaptation of the register implementation of [23] in the m&m model. A *single-writer multi-reader* (SWMR) register R can be written by a single writer process w , using a procedure WRITE, and can be read by all processes p_1, \dots, p_n , using a procedure READ. A register is *atomic* [29] if any execution of READ and WRITE operations can be linearized [27]. This means that there is a total order of all completed operations and some incomplete operations, that respects the real-time order of non-overlapping operations, in which each READ operation returns the value of the last preceding WRITE operation (or the initial value of the register, if there is no such WRITE).

The algorithm appears in Algorithm 1; for simplicity of presentation, a process sends each message also to itself and responds with the appropriate response. All the message communication between the processes is done in `msg_exchange()`, where we simply send a message and wait for $n - f$ acknowledgement. This modular approach allows us to replace the communication pattern according to the specific shared-memory layout. For example, Section 6 shows that in the cluster-based model this communication pattern can be changed to wait for less than $n - f$ processes.

■ **Algorithm 1** Atomic SWMR register implementation (w is the single writer).

Local Variables:

$w\text{-sqno}$: int, initially 0 ▷ write sequence number
 $r\text{-sqno}$: int, initially 0 ▷ read sequence number
 $last\text{-sqno}$: int, initially 0 ▷ last write sequence number observed
 $counter$: int, initially 0 ▷ number of replies/acks received so far

Shared Variables: for every process p and every $\mu \in W_p$:

$R_\mu[p]$: $\langle \text{int}, \text{int} \rangle$, initially $\langle 0, v_0 \rangle$ ▷ writable by p and readable by all processes that can read from μ , i.e., all the processes in R_μ

WRITE(v) – Code for the writer w :

1: $w\text{-sqno} = w\text{-sqno} + 1$ ▷ increment the write sequence number
2: $acks = \text{msg_exchange}\langle W, w\text{-sqno}, v \rangle$
3: return

Code for any process p :

4: Upon receipt of a $\langle W/WB, sqno, v \rangle$ message from process w/q :
5: **if** ($sqno > last\text{-sqno}$) **then**
6: $last\text{-sqno} = sqno$
7: **for each** $\mu \in W_p$ **do** ▷ write value and sequence number to every register p can write to
8: $R_\mu[p] = \langle sqno, v \rangle$
9: send $\langle \text{Ack-W/Ack-WB}, sqno \rangle$ to process w/q

READ() – Code for the reader q :

10: $r\text{-sqno} = r\text{-sqno} + 1$ ▷ increment the read sequence number
11: $set_of_tuples = \text{msg_exchange}\langle R, r\text{-sqno}, \perp \rangle$
12: $\langle seq, val \rangle = \max(set_of_tuples)$ ▷ maximum $\langle seq, val \rangle$
13: $acks = \text{msg_exchange}\langle WB, seq, val \rangle$ ▷ write back
14: return val

Code for any process p :

15: Upon receipt of a $\langle R, r\text{-sqno}, - \rangle$ message from process q :
16: $\langle w\text{-seq}, w\text{-val} \rangle = \max\{\langle seq, val \rangle : \mu \in R_p \cap W_q \text{ and } R_\mu[q] = \langle seq, val \rangle\}$ ▷ find val with maximum seq
17: send $\langle \text{Ack-R}, r\text{-sqno}, \langle w\text{-seq}, w\text{-val} \rangle \rangle$ to process q

msg_exchange(m, seq, val): returns set of responses

18: send $\langle m, seq, val \rangle$ to all processes
19: $responses = \emptyset$
20: **repeat**
21: wait to receive a message m of the form $\langle \text{Ack-}m, seq, - \rangle$
22: $counter = counter + 1$
23: $responses = responses \cup \{m\}$
24: **until** $counter \geq n - f$
25: return($responses$)

For each process p and memory $\mu \in W_p$ there is a shared SWMR register $R_\mu[p]$, writable by p and readable by every process that can read from μ , i.e., every process in R_μ . In **WRITE(v)**, the writer w increments its local write sequence number $w\text{-sqno}$ and calls **msg_exchange()**.

This procedure sends a message of type W with value v and w -sqno to all processes. On receiving a write message from w , p checks if the write value is more up-to-date than the last value it has observed, by checking if w -sqno is larger than $last$ -sqno. If so, p updates $last$ -sqno to be w -sqno and writes the value v and sequence number w -sqno to all the registers it can write to. When done, the process sends an acknowledgment to the writer w . Once w receives $n - f$ acknowledgments, it returns successfully.

In READ, a reader process q increments its local read sequence number r -sqno and calls `msg_exchange()`. This procedure sends a message of type R and r -sqno to all processes. On receiving a read message from q , a process p reads all the registers it can read and finds the maximum sequence number and value stored in them and sends this pair to the reader q . Once q receives $n - f$ acknowledgments, it finds the value val with maximum sequence number seq among the responses (i.e., it selects the most up-to-date value). Then, q calls `msg_exchange()`, with message type WB (write back) and value val and seq to update other readers. On receiving a write back message from q , each process p handles WB like W message, checking if w -sqno is larger than $last$ -sqno and if so updating $last$ -sqno and all the registers it can write to. When done, the process sends an acknowledgment to q . Once q receives $n - f$ acknowledgments, it returns val successfully.

The communication complexities of read and write operations are dominated by the cost of a `msg_exchange()`, invoked once in a write and twice in a read. This procedure takes one round-trip and $O(n)$ messages, like the algorithm for the pure message-passing model [7]. The number of shared SWMR registers depends on the shared-memory topology and is $\rho = \sum_{\text{process } p} |W_p|$, as every process has a single register in each memory it can write to. The number of accesses to the shared memory is $\sigma = \sum_{\text{process } p} \sum_{\mu \in R_p} |W_\mu|$, as every process reads all the registers it can read from. Note that $\sigma \leq n\rho$.

The only statement that could prevent the completion of a WRITE or a READ is waiting for $n - f$ responses (Line 24). Since at most f processes may crash, the wait statement eventually completes, implying that a WRITE or READ invoked by a process that does not crash completes.

► **Lemma 2.** *Let t_2 be the largest sequence number returned in a read `msg_exchange` by reader p_j , and assume that the `msg_exchange` starts after the completion of a write `msg_exchange`, either by the writer w or in a write back by reader p_i , with sequence number t_1 , then, $t_1 \leq t_2$.*

We explicitly order all completed reads and all invoked writes (even if they are incomplete). Note that values written by the writer w have distinct write sequence numbers, and are different from the initial value of the register, denoted v_0 ; the value of the k^{th} write operation is denoted v_k , $k \geq 1$. Writes are ordered by the order they are invoked by process w ; if the last write is incomplete, we place this write at the end. Since only one process invokes write, this ordering is well-defined and furthermore, the values written appear in the order v_1, v_2, \dots

Next, we consider reads in the order they complete; note that this means that non-overlapping operations are considered in their order in the execution. A read that returns the value v_{k-1} , $k \geq 0$, is placed before the k -th write in the ordering, if this write exists, and at the end of the ordering, otherwise. For $k = 0$, this means that the read is placed before the first write, which may be at the end of the order, if there is no write.

Lemma 2 implies that this order respects the real-time order of non-overlapping operations.

► **Theorem 3.** *If a system is not f -partitionable then Algorithm 1 implements an atomic SWMR register, in the presence of f failures.*

The impossibility proof holds even if only *regular* register [29] is implemented. In a regular register, a read should return the value of a WRITE operation that either overlaps it, or immediately precedes it. The proof is similar to the one in [23], where they show that a SWMR register cannot be implemented in the m&m model if more than f_L processes may fail.

► **Theorem 4.** *If a system is f -partitionable then there is no implementation of a regular SWMR register in the presence of f failures.*

4 Solving Other Problems in Non-Partitionable Systems

4.1 Constructing Other Read/Write Registers

The atomic SWMR register presented in the previous section can be used as a basic building block for implementing other shared-memory objects. Recall that if a system is not f -partitionable (i.e., $f \leq f_{opt}$), a SWMR register can be implemented so that each operation takes $O(1)$ time, $O(n)$ messages, $O(\rho)$ SWMR shared-memory registers and $O(\sigma)$ SWMR shared-memory accesses. Given a shared-memory algorithm that uses $O(r)$ SWMR registers and has $O(s)$ step complexity, it can be simulated with $O(s)$ round-trips, $O(ns)$ messages, and $O(\sigma s)$ shared-memory accesses. The simulation requires $O(\rho r)$ SWMR shared-memory registers. (Recall that $\rho = \sum_{\text{process } p} |W_p|$ and $\sigma = \sum_{\text{process } p} \sum_{\mu \in R_p} |W_\mu|$.)

An atomic *multi-writer multi-reader* (MWMR) register can be built from atomic SWMR registers [33]; each read or write requires $O(n)$ round-trips, $O(n^2)$ messages, $O(\rho n)$ SWMR shared registers and $O(\alpha n)$ shared-memory accesses.

Atomic snapshots can also be implemented using SWMR registers [13]; each scan or update takes $O(n \log n)$ round-trips, $O(n^2 \log n)$ messages, $O(\rho n)$ SWMR shared registers and $O(\sigma n \log n)$ shared-memory accesses.

4.2 Batching

A simple optimization is *batching* of read requests, namely reading the registers of several processes simultaneously. Batching is useful when each process replicates a register for each other process – not for just one writer. A process p can send read requests for all these registers together, instead of sending n separate read requests (for the registers of all processes), one after the other. When a process q receives the batched request from p , it replies with a vector containing the values of all registers in a single message, rather than sending them separately. Process p waits for vectors from $n - f$ processes, and picks from them the latest value for each other process. Finally, the reader does a write-back of this vector.

Batching reduces the number of round-trips and messages, and shared-memory registers and accesses, but increases the size of messages and registers. With batching, an operation on a MWMR register requires $O(1)$ round-trips, $O(n)$ messages, $O(\rho)$ SWMR shared registers and $O(\alpha)$ shared-memory accesses, when each process saves all the writers values in a single SWMR register. Batching can also be applied to atomic snapshots, so that each scan or update takes $O(\log n)$ round-trips, $O(n \log n)$ messages, $O(\rho)$ SWMR shared registers and $O(\sigma \log n)$ shared-memory accesses.

Batching provides a *regular collect*, as defined in [11]. Regular collects can be used in the following building block, where a process repeatedly call collect, and returns a vector of values if it has received it *twice* (in two consecutive collects). Two vectors are the same if they contain the same sequence numbers in each component. Process p can write the value v using

procedure $\text{WRITE}_p(v)$, and repeatedly double collect all the processes current values using the procedure $\text{BUILDINGBLOCK}()$. An invocation of $\text{BUILDINGBLOCK}()$ returns a vector V with n components, one for each process. Each component contains a pair of a value with a sequence number. For every process p_i , $V[i]$ is the entry in the vector corresponding to p_i 's value. A vector V_1 *precedes* a vector V_2 if the sequence number of each component of V_1 is smaller than or equal to the corresponding component of V_2 . Although the writes are not atomic, it can be shown that if V_1 and V_2 are vectors returned by two pairs of successful double collects then either V_1 precedes V_2 or V_2 precedes V_1 .

This building block may not terminate (even if the system is not f -partitionable), due to continuous writes. However, if two consecutive collects are not equal then some sequence number was incremented, i.e., a write by some process is in progress.

4.3 Consensus

In the *consensus* problem, a process starts with an input value and decides on an output value, so that all processes decide on the same value (*agreement*), which is the input value of some process (*validity*). With a standard *termination* requirement, it is well known that consensus cannot be solved in an asynchronous system [21]. This result holds whether processes communicate through shared memory or by message passing, and even if only a single process fails. However, consensus can be solved if the termination condition is weakened, either to be required only with high probability (*randomized consensus*), or to hold when it is possible to eventually detect failures (using a *failure detector*), or to happen only under fortunate situations.

There are numerous shared-memory randomized consensus algorithms, which rely on read / write registers, or objects constructed out of them. Using these algorithms together with linearizable register implementations is not obvious since linearizability does not preserve *hyperproperties* [9, 22]. It has been shown [24] that the ABD register implementation [7] is not *strongly linearizable* [22]. This extends to the mixed-model register implementations, as ABD is a special case of them.

Hadzilacos et al. [25] have proved that the simple randomized consensus algorithm of [6] works correctly with *regular* registers, and used it to obtain consensus in m&m systems [24]. Their algorithm inherits exponential complexity from the simple algorithm of [6], which employs independent coin flips by the processes.

Here, we explain how to use $\text{BUILDINGBLOCK}()$ to emulate the weak shared coin of [6], following [14]. This holds with f failures, if the system is not f -partitionable.

In Algorithm 2, a process flips a coin using a local function $\text{flip}()$, which returns the value 1 or -1, each with probably 1/2. Invoking $\text{flip}()$ is a single atomic step. After each flip, a process writes its outcome in an individual cumulative sum. Then it calls $\text{BUILDINGBLOCK}()$ to obtain a vector V with the individual cumulative sums of all processes. (We assume that the initial value in each component is 0.) The process then checks the absolute value of the total sum of the individual cumulative sums, denoted $\text{sum}(V)$. If it is at least $c \cdot n$ for some constant $c > 1$, then the process returns its sign.

Intuitively, the only way the adversary can create disagreement on the outcome of the shared coin is by preventing as many processors as possible to move the counter in the unwanted direction. We will show that the adversary cannot “hide” more than $n - 1$ coin flips. (This was originally proved when processes use atomic writes [6]; here, we show it holds even when writes are not atomic.) Therefore, after the cumulative sum is big or small enough the adversary can no longer affect the outcome of the shared coin, and cannot prevent the processes from terminating.

■ **Algorithm 2** Weak shared coin [6].

Local Variables:

my-counter: int, initially 0

V: vector of size n , with all entries initially 0

COIN() – Code for process p :

```

1: while true do
2:   my-counter = my-counter + flip()
3:   WRITEp(my-counter)
4:   V = BUILDINGBLOCK()
5:   if sum(V) ≥ c · n then return 1
6:   else if sum(V) ≤ -c · n then return -1

```

Let H and T be the number of 1 and -1 (respectively) flipped by all processes at some point in the execution. These numbers are well-defined since the local coin flips are atomic.

► **Lemma 5.** *If $H - T < -(c + 1) \cdot n$ (respectively, $H - T > (c + 1) \cdot n$) at some point in the execution, then a process that invokes BUILDINGBLOCK() after this point returns -1 (respectively, 1).*

Proof. (Sketch) We consider the first case; the other case is symmetric. Consider the set of processes that invoked BUILDINGBLOCK() after the point in the execution when $H - T < -(c + 1) \cdot n$, in the order their BUILDINGBLOCK() returns. Let p_{j_i} , $i \geq 1$, be the i th process in this order, and let V_i be the vector returned by its BUILDINGBLOCK(). We prove, by induction on i , that $\text{sum}(V_i) \leq -c \cdot n$, and hence, p_{j_i} returns -1.

In the base case, $i = 1$. Since a process invokes BUILDINGBLOCK() after every write, there can be at most n writes (either pending or finished) after the point $H - T < -(c + 1) \cdot n$ and the return of BUILDINGBLOCK() by p_{j_1} . Therefore, $\text{sum}(V_1) < -c \cdot n$, and p_{j_1} decides -1 in Line 6.

Inductive step: Assume that for $i > 1$, processes $p_{j_1}, \dots, p_{j_{i-1}}$ decide after their BUILDINGBLOCK() invocation returns. Therefore, there are no additional writes in the execution, and p_{j_i} will observe at most n additional values from $H - T$ and will return -1. ◀

► **Lemma 6.** *If process p returns 1 (respectively, -1) from the shared coin, then $H - T > (c - 1) \cdot n$ (respectively, $H - T < -(c - 1) \cdot n$) at some point during its last call to BUILDINGBLOCK().*

Proof. (Sketch) We consider the first case; the other case is symmetric. Consider the last pair of collects in the last BUILDINGBLOCK() invocation before process p returns, and assume they return a vector V . Assume p misses a write by some process q that overlaps the first collect, i.e., the sequence number of this write is smaller than the corresponding sequence number in V . Then q 's write overlaps p 's first collect, and it returns after the second collect starts. (Otherwise, the regularity of collect implies that the second collect returns this write by q , or a later one, contradicting the fact it is equal to the first collect.) Therefore, each process has at most one write that overlaps the first collect and can be missed by the first collect. So, the sum of V differs by at most $n - 1$ values from $H - T$ at the point when the first collect completes. Since p returns 1, $\text{sum}(V) \geq c \cdot n$, and it holds that $H - T > (c - 1) \cdot n$ when the first collect completes. ◀

The next lemma can be proved along the lines of [6, Theorem 17], using the fact (see proof of Lemma 5) that there are at most n additional writes after $H - T$ drops below $-(c + 1) \cdot n$.

► **Lemma 7.** *The adversary can force the weak shared coin procedure of a process to return 1 (respectively, -1) with probability at most $(c + 1)/2c$.*

It follows that the adversary can force the processes to disagree with probability at most $(c - 1)/2c$. The next theorem has the same proof as in [6].

► **Theorem 8.** *For a constant $c > 1$, the expected number of coin flips in an execution of the weak shared coin is $O(n^2)$.*

Since the expected number of coin flips is $O(n^2)$, the expected number of write and building block invocations is also $O(n^2)$. The total number of collect operations in these building block invocations for all the processes is $O(n^3)$ in expectation, this is because a double collect fails only when another coin is written. Therefore, the complexity of the weak shared coin is $O(n^3)$ round-trips, $O(n^4)$ messages, $O(\rho)$ registers and $O(\alpha n^3)$ shared-memory accesses. Plugging the weak shared coin in the overall algorithm of [6], proved to be correct by [25], yields a randomized consensus algorithm with the same expected complexities as the weak shared coin.

Next, we prove that randomized consensus cannot be solved in a partitionable system, by considering the more general problem of *non-deterministic f -terminating* consensus, an extension of *nondeterministic solo termination* [20]. This variant of consensus has the usual validity and agreement properties, with the following termination property:

Non-deterministic f -termination: For every configuration C , process p and set F of at most f processes, such that $p \notin F$, there is an F -free execution in which process p terminates.

► **Theorem 9.** *If a system is f -partitionable then non-deterministic f -terminating consensus is unsolvable.*

Proof. Assume, by way of contradiction, that there is a non-deterministic f -terminating consensus algorithm. Since the system is f -partitionable, there are two disjoint sets of processes P and P' , each of size $n - f$, such that $P' \not\rightarrow P$. Therefore, there are no two processes $p \in P$ and $p' \in P'$ so that p' can read from a memory and p can write to that same memory. Let Q be the processes not in $P \cup P'$. Since $|P|, |P'| = n - f$, it follows that $|P \cup Q| = |P' \cup Q| = f$.

To prove the theorem, we construct three executions. Consider an initial configuration, in which all processes in P have initial value 0. Since $|P' \cup Q| = f$, non-deterministic f -termination implies there is a $(P' \cup Q)$ -free execution, in which some process $p \in P$ terminates, say by time t_1 . Call this execution α_1 , and note that only processes in P take steps in α_1 . By validity, p decides 0.

In a similar manner, we can get a $(P \cup Q)$ -free execution, α_2 , in which initial values of all the processes in P' are 1, and by non-deterministic f -termination, some process $p' \in P'$ decides on 1, say by time t_2 . Note that only processes in P' take steps in α_2 .

Finally, the third execution α_3 combines α_1 and α_2 . The initial value of processes in P is 0, and the initial value of processes in P' is 1. Processes in Q have arbitrary initial values, and they take no steps in α_3 . The execution is identical to α_1 from time 0 until time t_1 , and to α_2 from this time until time $t_1 + t_2$. All messages sent between processes in P and processes in P' are delivered after time $t_1 + t_2$. Since processes in P' do not take steps in α_3 until time t_1 , all processes in P decides 0, as in α_1 . Processes in P' cannot receive messages from processes in P or read what processes in P write to the shared memory, therefore all processes in P' decides 1, as in execution α_2 , violating the agreement property. ◀

4.4 Approximate Agreement

In the *approximate agreement* problem with parameter $\epsilon > 0$, all processes start with a real-valued input and must decide on an output value, so any two decision values are in distance at most ϵ from each other (*agreement*), and any decision value is in the range of all initial values (*validity*).

There is a wait-free algorithm for the approximate agreement problem in the shared-memory model, which uses only SWMR registers [12]. This algorithm can be simulated if the system is not f -partitionable, and at most f processes fail. Similarly to randomized consensus, it can be shown that this problem is unsolvable in partitionable systems.

► **Theorem 10.** *If a system is f -partitionable then approximate agreement is unsolvable in the presence of f failures.*

4.5 Renaming

In the M -*renaming* problem, processes start with unique *original* names from a large namespace $\{1, \dots, N\}$, and the processes pick distinct *new* names from a smaller namespace $\{1, \dots, M\}$ ($M < N$). To avoid a trivial solution, in which a process p_i picks its index i as the new name, we require *anonymity*: a process p_i with original name m performs the same as process p_j with original name m .

Employing the SWMR register simulation in a $(2n - 1)$ -renaming algorithm [10] yields an algorithm that requires $O(n \log n)$ round-trips, $O(n^2 \log n)$ messages, $O(\rho n^4)$ shared registers and $O(\sigma n \log n)$ shared-memory accesses. The number of registers can reduce to $O(\rho)$, at the cost of increasing their size.

This algorithm assumes that the system is not f -partitionable and at most f processes fail. The next theorem shows that this is a necessary condition.

► **Theorem 11.** *If a system is f -partitionable then renaming is unsolvable in the presence of f failures.*

Proof. Assume, by way of contradiction, that there is a renaming algorithm. Since the system is f -partitionable, there are two disjoint sets of processes P and P' , each of size $n - f$, such that $P' \not\rightarrow P$. Denote $P = \{p_{i_1}, \dots, p_{i_{n-f}}\}$ and $P' = \{p'_{i_1}, \dots, p'_{i_{n-f}}\}$. Let Q be the set of processes not in $P \cup P'$. Since $|P|, |P'| = n - f$, we have that $|P \cup Q| = |P' \cup Q| = f$.

Given a vector I of $n - f$ original names, denote by $\alpha(I, P)$ the P -only execution in which processes in P have original names I : processes in $(P' \cup Q)$ crash and take no step, and processes in P are scheduled in round-robin. Since at most f processes fail in $\alpha(I, P)$, eventually all processes in P pick distinct new names, say by time $t(I)$. Note that by anonymity, the same names are picked in the execution $\alpha(I, P')$, in which p'_{i_j} starts with the same original name as p_{i_j} and takes analogous steps.

Consider $\alpha(I_i, P)$, for any possible set of original names. The original name space can be picked to be big enough to ensure that for two *disjoint* name assignments, I_1 and I_2 , some process $p_{i_j} \in P$ decides the same new name r in the executions $\alpha(I_1, P)$ and $\alpha(I_2, P)$.

Denote $\alpha_1 = \alpha(I_1, P)$ and $\alpha_2 = \alpha(I_2, P')$, namely, the execution in which processes in P' replace the corresponding processes from P . The anonymity assumption ensures that p'_{i_j} decides on r , just as p_{i_j} decides on r in $\alpha(I_1, P)$ and $\alpha(I_2, P)$.

The execution α_3 combines α_1 and α_2 , as follows. Processes in Q take no steps in α_3 . The original names of processes in P are I_1 , and original names of processes in P' are I_2 . The execution is identical to α_1 from time 0 until time $t(I_1)$, and to α_2 from this time until time $t(I_1) + t(I_2)$. All messages sent from processes in P to processes in P' and from processes in P' to processes in P are delivered after time $t(I_1) + t(I_2)$.

In α_3 , processes in P do not receive messages from processes in $P' \cup Q$. Furthermore, $P' \not\rightarrow P$; i.e., processes in P' cannot read what processes in P wrote to the shared memory. Hence, α_3 is indistinguishable to p_{i_j} from α_1 , and hence, it picks new name r . Similarly, α_3 is indistinguishable to p'_{i_j} from α_2 , and hence, it also picks new name r , which contradicts the uniqueness of new names. \blacktriangleleft

5 The M&M Model

In the m&m model [5, 23], the shared memory connections are defined by a *shared-memory domain* L , which is a collection of sets of processes. For each set $S \in L$, all the processes in the set may share any number of registers among them. Our model when the shared memory has no access restrictions is a dual of the general m&m model, and they both capture the same systems. We say that L is *uniform* if it is induced by an undirected *shared-memory graph* $G = (V, E)$, where each vertex in V represents a process p . For every process p , $S_p = \{p\} \cup \{q : (p, q) \in E\}$, then $L = \{S_p : p \text{ is a process}\}$. In the uniform m&m model each memory is associated with a process p , and all the processes in S_p may access it. That is, a process can access its own memory and the memories of its neighbors.

In the m&m model, there are no access restrictions on the shared memory. Hence, for every process p , $|R_p| = |W_p| = |S_p|$. Therefore, $\rho = \sum_{\text{process } p} |S_p| = \sum_{\text{process } p} d(p) + 1 = 2|E| + n = O(n^2)$ and $\sigma = O(n^3)$, where $d(p)$ is the degree of process p in the graph. Substituting into the algorithms presented in Section 4, we obtain polynomial complexity for all of them, including a polynomial randomized consensus algorithm. In the general m&m model, ρ and σ are unbounded.

► **Definition 12** ([23]). *Given a shared-memory domain L , f_L is the largest integer f such that for all process subsets P and P' of size $n - f$ each, either $P \cap P' \neq \emptyset$ or there is a set $S \in L$ that contains both a process from P and a process from P' .*

Hadzilacos, Hu and Toueg [23] show that an SWMR register can be implemented in the m&m model if and only if at most f_L process may fail. Therefore in the m&m model, $f_{opt} = f_L$. We can see the connection between the two definitions by observing that, in this model, $p \leftrightarrow q$ if $p = q$ or there is a set $S \in L$ such that $p, q \in S$. We simply write \leftrightarrow , since the shared memory has no access restrictions.


The *square* of a graph $G = (V, E)$ is the graph $G^2 = (V, E^2)$, where $E^2 = E \cup \{(u, v) : \exists w \in E \text{ such that } (u, w) \in E \text{ and } (w, v) \in E\}$. I.e., there is an edge in G^2 between every two vertices that are in distance at most 2 in the graph G .

► **Definition 13** ([23]). *Given an undirected graph $G = (V, E)$, f_G is the largest integer f such that for all subsets P and P' of V of size $n - f$ each, either $P \cap P' \neq \emptyset$ or G^2 has an edge (u, v) such that $u \in P$ and $v \in P'$.*


In the uniform m&m model, $f_L = f_G = f_{opt}$ [23], and $p \leftrightarrow q$ if $p = q$ or (p, q) is an edge in G^2 .

We have seen that $f_{opt} \leq f_{maj}$. Figure 1 shows a graph where $f_{opt} < f_{maj}$. Thus, the converse inequality does not hold in the (uniform or general) m&m model.

► **Definition 14** ([5]). *A process p represents itself and all its neighbors, that is, $\{p\} \cup \{q : (p, q) \in E\}$. A set of processes P represents the union of all the processes represented by processes in P .*



■ **Figure 1** Counter example.



■ **Figure 2** Counter example for $n = 5$.

Aguilera et al. [5] present a randomized consensus algorithm, called HBO, which is based on Ben-Or's algorithm [15]. Like Ben-Or's algorithm, HBO has exponential time and message complexities. HBO assumes that the nonfaulty processes represent a majority of the processes. Below, we show that the resilience of the HBO algorithm is not optimal. We first capture the condition required for the correctness of the HBO algorithm, with the next definition.

► **Definition 15.** $f_{m\&m}$ is the largest integer f such that every set P of $n - f$ processes represents a majority of the processes.

It can be shown that $f_{m\&m} \leq f_{opt}$. On the other hand, for every $n > 4$, there is a shared-memory graph, such that $f_{m\&m} < f_{opt}$ in the uniform m&m model. The graph is the star graph over n vertices, and has edges $\{(p_1, p_2)\} \cup \{(p_2, p_i) : 3 \leq i \leq n\}$. (See Figure 2, for $n = 5$.) Thus, requiring at least $n - f_{m\&m}$ nonfaulty processes is strictly stronger than requiring $n - f_{opt}$ nonfaulty processes. Therefore, the HBO algorithm does not have optimal resilience. Intuitively this happens since HBO does not utilize all the shared-memory connections that are embodied in G^2 . Thus, our algorithm (Section 4.3), has better resilience than HBO, which we show is optimal, in addition to having polynomial complexity.

Aguilera et al. [5] also present a lower bound on the number of failures any consensus algorithm can tolerate in the m&m model. To state their bound, consider a graph $G = (V, E)$, and let B, S and T be a partition of V . (B, S, T) is an *SM-cut* in G if B can be partitioned into two disjoint sets B_1 and B_2 , such that for every $b_1 \in B_1, b_2 \in B_2, s \in S$ and $t \in T$, we have that $(s, t), (b_1, t), (b_2, s) \notin E$.

► **Theorem 16** ([5]). *Consensus cannot be solved in the uniform m&m model in the presence of f failures if there is a SM-cut (B, S, T) such that $|S| \geq n - f$ and $|T| \geq n - f$.*

Although the resilience of HBO is not optimal, we show that this lower bound on resilience is optimal, by proving that if a system is f -partitionable then the condition in Theorem 16 holds. By Theorem 9, these two conditions are equal in the m&m model.

► **Theorem 17.** *In the uniform m&m model, if the system is f -partitionable then there is an SM-cut (B, S, T) with $|S| \geq n - f$ and $|T| \geq n - f$.*

6 The Cluster-based Model

In the hybrid, *cluster-based* model of [28, 31], processes are partitioned into $m, 1 \leq m \leq n$, non-empty and disjoint subsets P_1, \dots, P_m , called *clusters*. Each cluster has an associated shared memory; only processes of this cluster can (atomically) read from and write to this shared memory. The set of processes in the cluster of p is denoted $cluster[p]$. As in the m&m model, there are no access restrictions on the shared memory. Hence, $|R_p| = |W_p| = 1$ for every process p , and therefore, $\rho = n$ and $\sigma = O(n^2)$.

In the cluster-based model, $p \leftrightarrow q$ if and only if p and q are in the same cluster.

If $p \leftrightarrow q$ and $q \leftrightarrow w$, for some processes p, q and w , then p and q are in the same cluster and q and w are in the same cluster. Since clusters are disjoint, it follows that p and w are in the same cluster, implying that \leftrightarrow is transitive.

► **Definition 18.** $f_{cluster}$ is the largest integer f such that for all sets of processes P and P' , each of size $(n - f)$, either $P \cap P' \neq \emptyset$ or some cluster contains a process in P and a process in P' .

► **Observation 1.** In the cluster-based model $f_{opt} = f_{cluster}$.

► **Lemma 19.** In the cluster-based model, $f_{opt} = f_{maj}$.

► **Lemma 20.** In the cluster-based model, for every two sets of processes, P and Q , and $f \leq f_{opt}$, if $|\vec{P}| \geq n - f$ and $|\vec{Q}| \geq n - f$ then $P \leftrightarrow Q$.

Raynal and Cao [31] present two randomized consensus algorithms for the cluster-based model. One is also based on Ben Or's algorithm [15], using local coins, and the other is based on an external common coin (whose implementation is left unspecified). These algorithms terminate in an execution if there are distinct clusters whose total size is (strictly) larger than $n/2$, each containing at least one nonfaulty process. Clearly, if $f \leq f_{maj}$, this condition holds for every execution with at most f failures. Since $f_{opt} \leq f_{maj}$, the condition holds if there are at most $f \leq f_{opt}$ failures. Lemma 19 implies that these two definitions are equivalent by proving that $f_{opt} = f_{maj}$. This means that the maximum resilience guaranteeing that every two sets of nonfaulty processes can communicate is equal to the one guaranteeing that every set of nonfaulty processes can communicate with a majority of the processes.

In the cluster-based model, if a process $p \in P_i$ does not crash then all other processes receive the information from all the processes of P_i , as if none of them crashed. For this reason, we say that p represents all processes in P_i (note that this definition is different than Definition 14). If a process q receives messages from processes representing k clusters P_1, \dots, P_k , such that $|P_1| + \dots + |P_k| > n/2$, then it has received information from a majority of the processes. This observation does not change the resilience threshold, i.e., the maximal number of failures that can be tolerated, but allows to wait for a smaller number of messages, thereby, making the algorithm execute faster. Lemma 20 proves that every two sets of processes representing at least $n - f_{opt}$ processes can communicate. Therefore, instead of waiting for a majority of represented processes, as is done in [31], it suffices to wait for $n - f_{opt}$ represented processes. Since $n - f_{opt} \leq \lfloor n/2 \rfloor + 1$, this means that in some cases it suffices to wait for fewer than a majority of represented processes.

This is not the case in the m&m model. For example, in the graph of Figure 1, $f_{opt} = 6$. For $P = \{p_7, p_9\}$, $\vec{P} = \{p_1, p_6, p_7, p_8, p_9\}$, and for $Q = \{p_3, p_5\}$, $\vec{Q} = \{p_1, p_2, p_3, p_4, p_5\}$, so $|\vec{P}| = |\vec{Q}| = 5 > n/2$, but $P \not\leftrightarrow Q$. Therefore, even though the system is not f -partitionable, and the set of non-faulty processes can communicate with a majority of the processes, it does not suffice to wait for more than $n/2$ represented processes.

7 Discussion

This paper studies the optimal resilience for various problems in mixed models. Our approach builds on simulating a SWMR register, which allows to investigate the resilience of many problems, like implementing MWMR registers and atomic snapshots, or solving randomized consensus, approximate agreement and renaming. Prior consensus algorithms for mixed models [5, 31] start from a pure message-passing algorithm and then try to exploit the

added power of shared memory. In contrast, we start with a shared-memory consensus algorithm and systematically simulate it in the mixed model. This simplifies the algorithms and improves their complexity, while still achieving optimal resilience.

It would be interesting to investigate additional tasks and objects. An interesting example is *k-set consensus* [19], in which processes must decide on at most k different values. This is trivial for $k = n$ and reduces to consensus, for $k = 1$. For the pure message-passing model, there is a k -set consensus algorithm [19], when the number of failures $f < k$. This bound is necessary for solving the problem in shared memory systems [17, 26, 32]. Since resilience in a mixed system cannot be better than in the shared-memory model, it follows that $f < k$ is necessary and sufficient for any mixed model. Thus, when $f_{opt} < k - 1$, a system can be f -partitionable and still offer f -resilience for k -set consensus.¹

The weakest failure detector needed for implementing a register in the cluster-based model is strictly weaker than the weakest failure detector needed in the pure message-passing model [28]. This aligns with the improved resilience we can achieve in a mixed model compared to the pure message-passing model. It is interesting to explore the precise improvement in resilience achieved with specific failure detectors and other mixed models.

We would also like to study systems where the message-passing network is not a clique.

References

- 1 Gen-Z draft core specification. <https://genzconsortium.org/specification/gen-z-core-specification-1-1-draft/>. Accessed: 2020-08-26.
- 2 InfiniBand. <https://www.infinibandta.org/about-infiniband/>. Accessed: 2020-08-26.
- 3 iWARP. <https://en.wikipedia.org/wiki/IWARP>. Accessed: 2020-08-26.
- 4 RDMA over converged ethernet. https://en.wikipedia.org/wiki/RDMA_over_Converged_Ethernet. Accessed: 2020-08-26.
- 5 Marcos K. Aguilera, Naama Ben-David, Irina Calciu, Rachid Guerraoui, Erez Petrank, and Sam Toueg. Passing messages while sharing memory. In *PODC*, page 51–60, 2018.
- 6 James Aspnes and Maurice Herlihy. Fast randomized consensus using shared memory. *Journal of Algorithms*, 11(3):441–461, September 1990.
- 7 Hagit Attiya, Amotz Bar-Noy, and Danny Dolev. Sharing memory robustly in message-passing systems. *Journal of the ACM*, 42(1):124–142, January 1995.
- 8 Hagit Attiya, Amotz Bar-Noy, Danny Dolev, David Peleg, and Rüdiger Reischuk. Renaming in an asynchronous environment. *Journal of the ACM*, 37(3):524–548, 1990.
- 9 Hagit Attiya and Constantin Enea. Putting strong linearizability in context: Preserving hyperproperties in programsthat use concurrent objects. In *DISC*, pages 2:1–2:17, 2019.
- 10 Hagit Attiya and Arie Fouren. Adaptive and efficient algorithms for lattice agreement and renaming. *SIAM J. Comput.*, 31(2):642–664, 2001. doi:10.1137/S0097539700366000.
- 11 Hagit Attiya, Arie Fouren, and Eli Gafni. An adaptive collect algorithm with applications. *Distributed Computing*, 15(2):87–96, 2002.
- 12 Hagit Attiya, Nancy A. Lynch, and Nir Shavit. Are wait-free algorithms fast? *J. ACM*, 41(4):725–763, 1994. doi:10.1145/179812.179902.
- 13 Hagit Attiya and Ophir Rachman. Atomic snapshots in $O(n \log n)$ operations. *SIAM J. Comput.*, 27(2):319–340, 1998.

¹ There is a lower bound of $k > \frac{n-1}{n-f}$ for pure message-passing systems, proved using a partitioning argument [16]. It might seem that adding shared memory will allow to reduce this bound, however, this is not the case, since for the relevant ranges of k ($1 < k < n$), the bound on the number of failures implied from this bound is at least k .

- 14 Amotz Bar-Noy and Danny Dolev. A partial equivalence between shared-memory and message-passing in an asynchronous fail-stop distributed environment. *Math. Syst. Theory*, 26(1):21–39, 1993. doi:10.1007/BF01187073.
- 15 Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols. In *PODC*, pages 27–30, 1983.
- 16 Martin Biely, Peter Robinson, and Ulrich Schmid. Easy impossibility proofs for k -set agreement in message passing systems. In *OPODIS*, pages 299–312, 2011.
- 17 Elizabeth Borowsky and Eli Gafni. Generalized FLP impossibility result for t -resilient asynchronous computations. In *STOC*, pages 91–100, 1993.
- 18 Gabriel Bracha and Sam Toueg. Asynchronous consensus and broadcast protocols. *Journal of the ACM*, 32(4):824–840, 1985.
- 19 Soma Chaudhuri. More choices allow more faults: Set consensus problems in totally asynchronous systems. *Inf. Comput.*, 105(1):132–158, 1993. doi:10.1006/inco.1993.1043.
- 20 Faith E. Fich, Maurice Herlihy, and Nir Shavit. On the space complexity of randomized synchronization. *Journal of the ACM*, 45(5):843–862, 1998. doi:10.1145/290179.290183.
- 21 Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985. doi:10.1145/3149.214121.
- 22 Wojciech Golab, Lisa Higham, and Philipp Woelfel. Linearizable implementations do not suffice for randomized distributed computation. In *STOC*, page 373–382, 2011.
- 23 Vassos Hadzilacos, Xing Hu, and Sam Toueg. Optimal register construction in m&m systems. In *OPODIS*, pages 28:1–28:16, 2019.
- 24 Vassos Hadzilacos, Xing Hu, and Sam Toueg. Optimal register construction in m&m systems (version 3). *CoRR*, abs/1906.00298, 2020. URL: <http://arxiv.org/abs/1906.00298>.
- 25 Vassos Hadzilacos, Xing Hu, and Sam Toueg. Randomized consensus with regular registers. *CoRR*, abs/2006.06771, 2020. URL: <http://arxiv.org/abs/2006.06771>.
- 26 Maurice Herlihy and Nir Shavit. The asynchronous computability theorem for t -resilient tasks. In *STOC*, pages 111–120, 1993.
- 27 Maurice P. Herlihy and Jeannette M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, July 1990.
- 28 Damien Imbs and Michel Raynal. The weakest failure detector to implement a register in asynchronous systems with hybrid communication. *Theor. Comput. Sci.*, 512:130–142, 2013. doi:10.1016/j.tcs.2012.06.030.
- 29 Leslie Lamport. On interprocess communication—part I: Basic formalism. *Distributed Computing*, pages 77–85, 1986.
- 30 Kevin Lim, Jichuan Chang, Trevor Mudge, Parthasarathy Ranganathan, Steven K. Reinhardt, and Thomas F. Wenisch. Disaggregated memory for expansion and sharing in blade servers. *SIGARCH Comput. Archit. News*, 37(3):267–278, June 2009. doi:10.1145/1555815.1555789.
- 31 Michel Raynal and Jiannong Cao. One for all and all for one: Scalable consensus in a hybrid communication model. In *ICDCS*, pages 464–471, 2019.
- 32 Michael E. Saks and Fotios Zaharoglou. Wait-free k -set agreement is impossible: the topology of public knowledge. In *STOC*, pages 101–110, 1993.
- 33 Paul M. B. Vitányi and Baruch Awerbuch. Atomic shared register access by asynchronous hardware. In *FOCS*, pages 233–243, 1986.