

Lattice Agreement in Message Passing Systems

Xiong Zheng

University of Texas at Austin, Austin, TX 78712, USA
zhengxiongty@utexas.edu

Changyong Hu

University of Texas at Austin, Austin, TX 78712, USA
colinhu9@utexas.edu

Vijay K. Garg

University of Texas at Austin, Austin, TX 78712, USA
garg@ece.utexas.edu

Abstract

This paper studies the lattice agreement problem and the generalized lattice agreement problem in distributed message passing systems. In the lattice agreement problem, given input values from a lattice, processes have to non-trivially decide output values that lie on a chain. We consider the lattice agreement problem in both synchronous and asynchronous systems. For synchronous lattice agreement, we present two algorithms which run in $\log(f)$ and $\min\{O(\log^2 h(L)), O(\log^2 f)\}$ rounds, respectively, where $h(L)$ denotes the height of the input sublattice L , $f < n$ is the number of crash failures the system can tolerate, and n is the number of processes in the system. These algorithms have significant better round complexity than previously known algorithms. The algorithm by Attiya et al. [Attiya et al. DISC, 1995] takes $\log(n)$ synchronous rounds, and the algorithm by Mavronicolasa [Mavronicolasa, 2018] takes $\min\{O(h(L)), O(\sqrt{f})\}$ rounds. For asynchronous lattice agreement, we propose an algorithm which has time complexity of $2 * \min\{h(L), f + 1\}$ message delays which improves on the previously known time complexity of $O(n)$ message delays.

The generalized lattice agreement problem defined by Faleiro et al in [Faleiro et al. PODC, 2012] is a generalization of the lattice agreement problem where it is applied for the replicated state machine. We propose an algorithm which guarantees liveness when a majority of the processes are correct in asynchronous systems. Our algorithm requires $\min\{O(h(L)), O(f)\}$ units of time in the worst case which is better than $O(n)$ units of time required by the algorithm in [Faleiro et al. PODC, 2012].

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Lattice Agreement, Replicated State Machine, Consensus

Digital Object Identifier 10.4230/LIPIcs.DISC.2018.41

Related Version A full version of the paper is available at <https://arxiv.org/abs/1807.11557>.

Funding Supported by NSF CNS-1812349, NSF CNS-1563544, NSF CNS-1346245, Huawei Inc., and the Cullen Trust for Higher Education Endowed Professorship.

Acknowledgements We want to thank John Kaippallimalil for providing some useful application cases for CRDT and generalized lattice agreement.



© Xiong Zheng, Changyong Hu, and Vijay K. Garg;
licensed under Creative Commons License CC-BY

32nd International Symposium on Distributed Computing (DISC 2018).

Editors: Ulrich Schmid and Josef Widder; Article No. 41; pp. 41:1–41:17

Leibniz International Proceedings in Informatics



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

1 Introduction

Lattice agreement, introduced in [2] to solve the atomic snapshot problem [1] in shared memory, is an important decision problem in distributed systems. In this problem, processes start with input values from a lattice and need to decide values which are comparable to each other. Lattice agreement problem is a weaker decision problem than consensus. In synchronous systems, consensus cannot be solved in fewer than $f + 1$ rounds [6], but lattice agreement can be solved in $\log f$ rounds (shown by an algorithm we propose). In asynchronous systems, the consensus problem cannot be solved even with one failure [8], whereas the lattice agreement problem can be solved in asynchronous systems when a majority of processes is correct [7].

In synchronous message passing systems, a $\log n$ rounds recursive algorithm based on “branch-and-bound” approach is proposed in [2] to solve the lattice agreement problem with message complexity of $O(n^2)$. It can tolerate at most $n - 1$ process failures. Later, [12] gave an algorithm with round complexity of $\min\{1 + h(L), \lfloor (3 + \sqrt{8f + 1})/2 \rfloor\}$, for any execution where at most $f < n$ processes may crash. Their algorithm has the early-stopping property and is the first algorithm with round complexity that depends on the actual height of the input lattice. Our first algorithm, for synchronous lattice agreement, LA_α , requires $\log h(L)$ rounds. It assumes that the height of the input lattice is known to all processes. By applying this algorithm as a building block, we give an algorithm, LA_β , which requires only $\log f$ rounds without the height assumption in LA_α . Instead of directly trying to decide on the comparable output values which are from a lattice with an unknown height, this algorithm first performs lattice agreement on the failure set known by each process by using LA_α . Then each process removes values from *faulty* processes they know and outputs the join of all the remaining values. Our third algorithm, LA_γ , has round complexity of $\min\{O(\log^2 h(L)), O(\log^2 f)\}$, which depends on the height of the input lattice but does not assume that the height is known. This algorithm iteratively guesses the actual height of the input lattice and applies LA_α with the guessed height as input, until all processes terminate.

Lattice agreement in asynchronous message passing systems is useful due to its applications in atomic snapshot objects and fault-tolerant replicated state machines. Efficient implementation of atomic snapshot objects in crash-prone asynchronous message passing systems is important because they can make design of algorithms in such systems easier (examples of algorithms in message passing systems based on snapshot objects can be found in [18], [13] and [4]). As shown in [2], any algorithm for lattice agreement can be applied to solve the atomic snapshot problem in a shared memory system. We note that [3] does not directly use lattice agreement to solve the atomic snapshot problem, but their idea of producing comparable *views* for processes is essentially lattice agreement. Thus, by using the same transformation techniques in [2] and [3], algorithms for lattice agreement problem can be directly applied to implement atomic snapshot objects in crash-prone message passing systems. We give an algorithm for asynchronous lattice agreement problem which requires $\min\{O(h(L)), O(f)\}$ message delays. Then, by applying the technique in [3], our algorithm can be used to implement atomic snapshot objects on top of crash-prone asynchronous message passing systems and achieve time complexity of $O(f)$ message delays in the worst case. Our result significantly improves the message delays in the previous work by Delporte-Gallet, Fauconnier et al [5]. The algorithm in [5] directly implements an atomic snapshot object on top of crash-prone message passing systems and requires $O(n)$ message delays in the worst case.

Another related work for lattice agreement in asynchronous systems is by Faleiro et al. [7]. They solve the lattice agreement problem in asynchronous systems by giving a Paxos style protocol [10, 11], in which each proposer keeps proposing a value until it gets *accept* messages from a majority of acceptors. The acceptor only *accepts* a proposal when the proposal has a bigger value than its accepted value. Their algorithm requires $O(n)$ message delays. Our asynchronous lattice agreement algorithm does not have Paxos style. Instead, it runs in *round-trips*. Each round-trip is composed of sending a message to all and getting $n - f$ acknowledgements back. Our algorithm guarantees termination in $\min\{O(h(L)), O(f)\}$ message delays which is a significant improvement over $O(n)$ message delays.

Generalized lattice agreement problem defined in [7] is a generalization of the lattice agreement problem in asynchronous systems. It is applied to implement a specific class of replicated state machines. In conventional replicated state machine approach [14], consensus based mechanism is used to implement strong consistency. Due to performance reasons, many systems relax the strong consistency requirement and support eventual consistency [17], i.e, all copies are eventually consistent. However, there is no guarantee on when this eventual consistency happens. Also, different copies could be in an inconsistent state before this eventual situation happens. Conflict-free replicated data types (CRDT) [15, 16] is a data structure which supports such eventual consistency. In CRDT, all operations are designed to be commutative such that they can be concurrently executed without coordination. As shown in [7] by applying generalized lattice agreement on top of CRDT, the states of any two copies can be made comparable and thus provide linearizability guarantee [9] for CRDT.

The following example from [7] motivates generalized lattice agreement. Consider a replicated set data structure which supports adds and reads. Suppose there are two concurrent updates, $add(a)$ and $add(b)$, and two concurrent reads on copy one and two respectively. By using CRDT, it could happen that the two reads return $\{a\}$ and $\{b\}$ respectively. This execution is not linearizable [9], because if $add(a)$ appears before $add(b)$ in the linear order, then no read can return $\{b\}$. On the other hand, if we use conventional consensus replicated state machine technique, then all operations would be coordinated including the two reads. This greatly impacts the throughput of the system. By applying generalized lattice agreement on top of CRDT, all operations can be concurrently executed and any two reads always return comparable views of the system. In the above example, the two reads return either (i) $\{a\}$ and $\{a, b\}$ or (ii) $\{b\}$ and $\{a, b\}$ which is linearizable. Therefore, generalized lattice agreement can be applied on top of CRDT to provide better consistency guarantee than CRDT and better availability than conventional replicated state machine technique.

Since the generalized lattice agreement problem has applications in building replicated state machines, it is important to reduce the message delays for a value to be learned. Faleiro et al. [7] propose an algorithm for the generalized lattice agreement by using their algorithm for the lattice agreement problem as a building block. Their generalized lattice agreement algorithm satisfies safety and liveness assuming $f < \lceil \frac{n}{2} \rceil$. A value is eventually learned in their algorithm after $O(n)$ message delays in the worst case. Our algorithm guarantees that a value is learned in $\min\{O(h(L)), O(f)\}$ message delays.

In summary, this paper makes the following contributions:

- We present an algorithm, LA_α to solve the lattice agreement in synchronous system in $\log h(L)$ rounds assuming $h(L)$ is known. Using LA_α , we propose an algorithm, LA_β to solve the standard lattice agreement problem in $\log f$ rounds. This bound is significantly better than the previously known upper bounds of $\log n$ by [3] and $\min\{1 + h(L), \lceil (3 + \sqrt{8f + 1})/2 \rceil\}$ by [12] (and solves the open problem posed there). We also give an algorithm, LA_γ which runs in $\min\{O(\log^2 h(L)), O(\log^2 f)\}$ rounds.

■ **Table 1** Previous Work and Our Results.

Problem	Protocol	Time	Message
LA sync	[3]	$O(\log n)$	$O(n^2)$
	[12]	$\min\{O(h(L)), O(\sqrt{f})\}$	$n^2 \cdot \min\{O(h(L)), O(\sqrt{f})\}$
	LA_α	$O(\log h(L))$	$O(n^2 \log h(L))$
	LA_β	$O(\log f)$	$O(n^2 \log f)$
	LA_γ	$\min\{O(\log^2 h(L)), O(\log^2 f)\}$	$n^2 \cdot \min\{O(\log^2 h(L)), O(\log^2 f)\}$
LA async	[7]	$O(n)$	$O(n^3)$
	LA_δ	$\min\{O(h(L)), O(f)\}$	$n^2 \cdot \min\{O(h(L)), O(f)\}$
GLA async	[7]	$O(n)$	$O(n^3)$
	GLA_α	$\min\{O(h(L)), O(f)\}$	$n^2 \cdot \min\{O(h(L)), O(f)\}$

- For the lattice agreement problem in asynchronous systems, we give an algorithm, LA_δ which requires $2 \cdot \min\{h(L), f + 1\}$ message delays which improves the $O(n)$ bound by [7].
- Based on the asynchronous lattice agreement algorithm, we present an algorithm, GLA_α , to solve the generalized lattice agreement with time complexity $\min\{O(h(L)), O(f)\}$ message delays which improves the $O(n)$ bound by [7].

Related previous work and our results are summarized in Table 1. LA_{sync} and LA_{async} represent lattice agreement in synchronous systems and asynchronous systems, respectively. GLA_{async} represents generalized lattice agreement in asynchronous systems. LA_α is designed to solve the lattice agreement problem with the assumption that the height of the input lattice is given. It serves as a building block for LA_β and LA_γ . For synchronous systems, the time complexity is given in terms of synchronous rounds. For asynchronous system, the time complexity is given in terms of message delays. The message column represents the total number of messages sent by all processes in one execution. For generalized lattice agreement problem, the message complexity is given in terms of the number of messages needed for a value to be learned.

2 System Model and Problem Definitions

2.1 System Model

We assume a distributed message passing system with n processes in a completely connected topology, denoted as p_1, \dots, p_n . We consider both synchronous or asynchronous systems. Synchronous means that message delays and the duration of the operations performed by the process have an upper bound on the time. Asynchronous means that there is no upper bound on the time for a message to reach its destination. The model assumes that processes may have crash failures but no Byzantine failures. The model parameter f denotes the maximum number of processes that may crash in a run. We assume that the underlying communication system is reliable but the message channel may not be FIFO. We say a process is *faulty* in a run if it crashes and *correct* or *non-faulty* otherwise. In our following algorithms, when a process sends a message to all, it also sends this message to itself.

2.2 Lattice Agreement

Let (X, \leq, \sqcup) be a finite join semi-lattice with a partial order \leq and join \sqcup . Two values u and v in X are comparable iff $u \leq v$ or $v \leq u$. The join of u and v is denoted as $\sqcup\{u, v\}$. X is a *join semi-lattice* if a join exists for every nonempty finite subset of X . As customary in this area, we use the term *lattice* instead of *join semi-lattice* in this paper for simplicity.

In the lattice agreement problem [2], each process p_i can propose a value x_i in X and must decide on some output y_i also in X . An algorithm is said to solve the lattice agreement problem if the following properties are satisfied:

Downward-Validity: For all $i \in [1..n]$, $x_i \leq y_i$.

Upward-Validity: For all $i \in [1..n]$, $y_i \leq \sqcup\{x_1, \dots, x_n\}$.

Comparability: For all $i \in [1..n]$ and $j \in [1..n]$, either $y_i \leq y_j$ or $y_j \leq y_i$.

In this paper, all the algorithms that we propose apply join operation to some subset of input values. Therefore, it is sufficient to focus on the join-closed subset of X that includes all input values. Let L be the join-closed subset of X that includes all input values. L is also a join semi-lattice. We call L the *input sublattice* of X . All algorithms proposed in this paper are based on L . Since the complexity of our algorithms depend on the height of lattice L , we give the formal definitions as below:

► **Definition 1.** The height of a value v in a lattice X is the length of longest path from any minimal value to v , denoted as $h_X(v)$ or $h(v)$ when it is clear.

► **Definition 2.** The height of a lattice X is the height of its largest value, denoted as $h(X)$.

Each process proposes a value from a boolean lattice. Thus, the largest value in this lattice is the set consists of all the n values. From the definition 2, we have $h(L) \leq n$.

2.3 Generalized Lattice Agreement

In generalized lattice agreement problem, each process may receive a possibly infinite sequence of values as inputs that belong to a lattice at any point of time. Let x_i^p denote the i th value received by process p . The aim is for each process p to learn a sequence of output values y_j^p which satisfies the following conditions:

Validity: Any learned value y_j^p is a join of some set of received input values.

Stability: The value learned by any process p is non-decreasing: $j < k \implies y_j^p \leq y_k^p$.

Comparability: Any two values y_j^p and y_k^q learned by any two process p and q are comparable.

Liveness: Every value x_i^p received by a correct process p is eventually included in some learned value y_k^q of every correct process q : i.e, $x_i^p \leq y_k^q$.

3 Lattice Agreement in Synchronous Systems

3.1 Lattice Agreement with Known Height

In this section, we first consider a simpler version of the standard lattice agreement problem by assuming that the height of the input sublattice L is known in advance, i.e, $h(L)$ is given. We propose an algorithm, LA_α , to solve this problem in $\log h(L)$ synchronous rounds. In section 3.2, we give an algorithm to solve the lattice agreement problem when the height is not given using this algorithm.

Algorithm LA_α runs in synchronous rounds. At each round, by calling a *Classifier* procedure (described below), processes within a same group (to be defined later) are classified into different groups. The algorithm guarantees that any two processes within the same group have equal values and any two processes in different groups have comparable values at the end. Thus, values of all processes are comparable to each other at the end. We present the algorithm by first introducing the fundamental *Classifier* procedure.

3.1.1 The Classifier Procedure

The *Classifier* procedure is inspired by the Classifier procedure given by Attiya and Rachman in [3], called *AR-Classifier*, where it is applied to solve the atomic snapshot problem in the shared memory system. The intuition behind the *Classifier* procedure is to classify processes to *master* or *slave* and ensure all *master* processes have values greater than all *slave* processes.

The pseudo-code for *Classifier* is given in Figure 1. It takes two parameters: the input value v and the threshold value k . The output is composed of three items: the output value, the classification result and the decision status. The process which calls the *Classifier* procedure should update their value to be the output value. The classification result is either *master* or *slave*. The decision status is a Boolean value which is used to inform whether the invoking process can decide on the output value or not. The main functionality of the *Classifier* procedure is either to tell the invoking process to decide, or to classify the invoking process as a *master* or a *slave*. Details of the *Classifier* procedure are shown below:

Line 1-3: The invoking process sends a message with its input value v and the threshold value k to all. It then collects all the received values associated with the threshold value k in a set U .

Line 5-6: It checks whether all values in U are comparable to the input value. If they are comparable, it terminates the *Classifier* procedure and returns the input value as the output value and *true* as the decision status.

Line 8-12: It performs classification based on received values. Let w be the join of all received values associated with the threshold value k . If the height of w in lattice L is greater than the threshold value k , then the *Classifier* returns w as the output value, *master* as the classification result and *true* as the decision status. Otherwise, it returns the input value as the output value, *slave* as the classification result and *false* as the decision status. From the classification steps, it is easy to see that the processes classified as *master* have values greater than those classified as *slave* because w is the join of all values in U .

There are four main differences between the *AR-Classifier* and our *Classifier*: 1) The *AR-Classifier* is based on the shared memory model whereas our algorithm is based on synchronous message passing. 2) The *AR-Classifier* does not allow early termination. 3) Each process in the *AR-Classifier* needs values from all processes whereas our *Classifier* uses values only from processes within its group. 4) The *AR-Classifier* procedure requires the invoking process to read values of all processes again if the invoking process is classified as *master* whereas our algorithm needs to receive values from all processes only once.

3.1.2 Algorithm LA_α

Algorithm LA_α (shown in Figure 2) runs in at most $\log h(L)$ rounds. It assumes knowledge of $H = h(L)$, the height of the input lattice. Let x_i denote the initial input value of process i , v_i^r denote the value held by process i at the beginning of round r , and $class$ denote the classification result of the *Classifier* procedure. The $class$ indicates whether the process is classified as a *master* or a *slave*. The $decided$ variable shows whether the process has decided or not. Each process i has a label denoted as l_i . This label is updated at each round. Processes which have the same label l are said to be in the same group with label l . The definitions of *label* and *group* are formally given as:

► **Definition 3 (label).** Each process has a *label*, which serves as a knowledge threshold and is passed as the threshold value k whenever the process calls the *Classifier* procedure.

► **Definition 4 (group).** A *group* is a set of processes which have the same label. The label of a group is the label of the processes in this group.


```

Classifier( $v, k$ ):
 $v$ : input value     $k$ : threshold value
1: Send ( $v, k$ ) to all
2: Receive messages of the form ( $-, k$ )
3: Let  $U$  be values contained in received messages
3:
4: /* Early Termination */
5: if  $|U| = 0$  or  $\forall u \in U : v \leq u \vee u \leq v$ 
6:   return ( $v, -, true$ )
6:
7: /* Classification */
8: Let  $w := \sqcup\{u : u \in U\}$ 
9: if  $h(w) > k$ 
10:  return ( $w, master, false$ )
11: else
12:  return ( $v, slave, false$ )

```

■ **Figure 1** *Classifier*.

```

LA $_{\alpha}$ ( $H, x_i$ ) for  $p_i$ :
 $H$ : given height     $x_i$ : input value
1:  $v_i^1 := x_i$  // value at round 1
2:  $l_i := \frac{H}{2}$  // label
3:  $decided := false$ 
4:
5: for  $r := 1$  to  $\log H + 1$ 
6:   ( $v_i^{r+1}, class, decided$ )
   := Classifier( $l_i, v_i^r$ )
7:   if  $decided$ 
8:     return  $v_i^{r+1}$ 
9:   else if  $class = master$ 
10:     $l_i := l_i + \frac{H}{2^{r+1}}$ 
11:   else
12:     $l_i := l_i - \frac{H}{2^{r+1}}$ 
13: end for

```

■ **Figure 2** Algorithm *LA* $_{\alpha}$.

A process has *decided* if it has set its decision status to true. Otherwise, it is undecided. At each round r , an undecided process invokes the *Classifier* procedure with its current value and its current label l_i as parameters v and k , respectively. Since each process passes its label as the threshold value k when invoking the *Classifier* procedure, line 2 of the *Classifier* is equivalent to receiving messages from processes within the same group; that is, at each round, a process performs the *Classifier* procedure within its group. Processes which are in different groups do not affect each other. At round r , by invoking the *Classifier* procedure, each process i sets v_i^{r+1} , *class* and *decided* to the returned output value, the classification result and the decision status. Each process first checks the value of *decided*. If it is true, process i decides on v_i^{r+1} and terminates the algorithm. Otherwise, if it is classified as a *master*, it increases its label by $\frac{H}{2^{r+1}}$. If it is classified as a *slave*, it decreases its label by $\frac{H}{2^{r+1}}$. Now we show how the *Classifier* procedure combined with this label update mechanism makes any two processes have comparable values at the end.

Let G be a group of processes at round r . Let $M(G)$ and $S(G)$ be the group of processes which are classified as *master* and *slave*, respectively, when they run the *Classifier* procedure in group G . We say that G is the parent of $M(G)$ and $S(G)$. Thus, $M(G)$ and $S(G)$ are both groups at round $r + 1$. Process $i \in M(G)$ or $i \in S(G)$ indicates that i does not decide in group G at round r . Initially, all process have the same label $\frac{H}{2}$ and are in the same group with label $\frac{H}{2}$. When they execute the *Classifier*, they will be classified into different groups. We can view the execution as processes traversing through a binary tree. Initially, all of them are at the root of the tree. As the program executes, if they are classified as *master*, then they go to the right child. Otherwise, they go to the left child.

Before we prove the correctness of the given algorithm, we first give some useful properties satisfied by the *Classifier* procedure. Although Lemma 5 is similar to a lemma given in [5], it is discussed here in message passing systems and the proofs are different.

► **Lemma 5.** *Let G be a group at round r with label k . Let L and R be two nonnegative integers such that $L \leq k \leq R$. If $L < h(v_i^r) \leq R$ for every process $i \in G$, and $h(\sqcup\{v_i^r : i \in G\}) \leq R$, then*

- (p1) for each process $i \in M(G)$, $k < h(v_i^{r+1}) \leq R$
 (p2) for each process $i \in S(G)$, $L < h(v_i^{r+1}) \leq k$
 (p3) $h(\sqcup\{v_i^{r+1} : i \in M(G)\}) \leq R$
 (p4) $h(\sqcup\{v_i^{r+1} : i \in S(G)\}) \leq k$, and
 (p5) for each process $i \in M(G)$, $v_i^{r+1} \geq \sqcup\{v_i^{r+1} : i \in S(G)\}$

Proof.

(p1)–(p3): Immediate from the *Classifier* procedure.

(p4): Since $S(G)$ is a group of processes which are at round $r + 1$, all processes in $S(G)$ are correct (non-faulty) at round r . So, all processes in $S(G)$ must have received values of each other in the *Classifier* procedure at round r in group G . Thus, $h(\sqcup\{v_i^{r+1} : i \in S(G)\}) \leq k$, otherwise all of them should be in group $M(G)$ instead of $S(G)$, according to the condition at *line 9* of the *Classifier* procedure.

(p5): Since all processes in $S(G)$ are correct at round r , all processes in $M(G)$ must have received values of all processes in $S(G)$ in the *Classifier* procedure at round r . Any process which proceeds to group $M(G)$ takes the join of all received values at round r , according to *line 10*. Thus, for every process $i \in M(G)$, $v_i^{r+1} \geq \sqcup\{v_i^{r+1} : i \in S(G)\}$. ◀

► **Lemma 6.** *Let x be a value from a lattice L , and V be a set of values from L . Let U be any subset of V . If x is comparable with $\forall v \in V$, then x is comparable with $\sqcup\{u \mid u \in U\}$.*

Proof. If $\forall u \in U : u \leq x$, then $\sqcup\{u \mid u \in U\} \leq x$. Otherwise, $\exists y \in U : x \leq y$. Since $y \leq \sqcup\{u \mid u \in U\}$, so $x \leq \sqcup\{u \mid u \in U\}$. ◀

► **Lemma 7.** *If process i decides at round r on value y_i , then y_i is comparable with v_j^r for any correct process j .*

Proof. Let process i decide in group G at round r . Consider the two cases below:

Case 1: $j \notin G$. Let G' be a group at the maximum round r' such that both i and j belong to G' . Then, either $i \in M(G') \wedge j \in S(G')$ or $j \in M(G') \wedge i \in S(G')$. We only consider the case $i \in M(G') \wedge j \in S(G')$. The other case is similar. From (p5) of Lemma 5, we have $\sqcup\{v_p^r : p \in S(G')\} \leq y_i$. Since $j \in S(G')$, then $v_j^r \leq \sqcup\{v_p^r : p \in S(G')\}$. Thus, $v_j^r \leq y_i$. For the other case, we have $y_i \leq v_j^r$. Therefore, y_i is comparable with v_j^r .

Case 2: $j \in G$, since process j is correct, then i must have received v_j^r at round r . Thus, by *line 5* of the *Classifier* procedure, we have that y_i is comparable with y_j^r . ◀

Now we show that any two processes decide on comparable values.

► **Lemma 8.** (*Comparability*) *Let process i and j decide on y_i and y_j , respectively. Then y_i and y_j are comparable.*

Proof. Let process i and j decide at round r_i and r_j , respectively. Without loss of generality, assume $r_i \leq r_j$. At round r_i , from Lemma 7 we have y_i is comparable with v_k^r for any correct undecided process k . Let $V = \{v_k^{r_i} \mid \text{process } k \text{ undecided and correct}\}$. Since $r_j \geq r_i$, y_j is at most the join of a subset of V . Thus, from Lemma 6 we have y_i and y_j are comparable. ◀

Now we prove that all processes decide within $\log H + 1$ rounds by showing all processes in the same group at the beginning of round $\log H + 1$ have equal values, given by Lemma 9 and Lemma 10. Since Lemma 9 and Lemma 10 and the corresponding proofs are similar to the ones given in [3], the proofs are omitted here and can be found in the full paper. Proof of Lemma 9 is based on (p1-p4) of Lemma 5 by induction. Proof of Lemma 10 is based on Lemma 9.

► **Lemma 9.** *Let G be a group of processes at round r with label k . Then*

- (1) *for each process $i \in G$, $k - \frac{H}{2^r} < h(v_i^r) \leq k + \frac{H}{2^r}$*
 (2) *$h(\sqcup\{v_i^r : i \in G\}) \leq k + \frac{H}{2^r}$*

► **Lemma 10.** *Let i and j be two processes that are within the same group G at the beginning of round $r = \log H + 1$. Then v_i^r and v_j^r are equal.*

► **Lemma 11.** *All processes decide within $\log H + 1$ rounds.*

Proof. From Lemma 10, we know any two processes which are in the same group at the beginning of round $\log H + 1$ have equal values. Then, the condition in *line 5* of *Classifier* procedure is satisfied. Thus, all undecided processes decide at round $\log H + 1$. ◀

► **Remark 12.** Since at the beginning of round $\log H + 1$ all undecided processes have comparable values, LA_α only needs $\log H$ rounds. For simplicity, one more round is executed to make all processes decide at *line 5* of the *Classifier* procedure.

► **Theorem 13.** *Algorithm LA_α solves lattice agreement problem in $\log H$ rounds and can tolerate $f < n$ failures.*

Proof. *Downward-Validity* follows from the fact that the value of each process is non-decreasing at each round. For *Upward-Validity*, according to the *Classifier* procedure, each process either keeps its value unchanged or takes the join of the values proposed by other processes which could never be greater than $\sqcup\{x_1, \dots, x_n\}$. For *Comparability*, from Lemma 8, we know for any two process i and j , if they decide, then their decision values must be comparable. From Lemma 11, we know all processes decide. Thus, comparability holds. ◀

Complexity. Time complexity is $\log H$ rounds. For message complexity, since each process sends n messages per round, $\log H$ rounds results in $n^2 \log H$ messages in total. Notice that the number of messages can be further reduced by keeping a set of processes which are not in its group. If a process p receives a message from process q with a threshold value different from its own threshold value, it knows that q is not in its group. Each process does not send messages to the processes in this set.

Algorithm LA_β runs in $\log \text{height}(L)$ rounds by assuming that $\text{height}(L)$ is given. However, in order to know that actual height of input lattice, we need to know how many distinct values all process propose which needs extra effort. For this reason, in following sections, we introduce algorithms to solve the lattice agreement problem without this assumption.

3.2 Lattice Agreement with Unknown Height

In this section, we consider the standard lattice agreement in which the height of the lattice is not known to any process. We propose algorithm, LA_β , (shown in Figure 3) based on algorithm LA_α .

3.2.1 Algorithm LA_β

Algorithm LA_β runs in $\log f + 1$ synchronous rounds. It makes use of algorithm LA_α as a building block. Instead of directly agreeing on input values which are taken from a lattice with unknown height, we first do lattice agreement on the failure set that each process knows *after one round of broadcast*. The set of all failure sets forms a boolean lattice with union as the join operation and with height equal to f (since there are at most f failures). The algorithm consists of two phases. At *Phase A*, all processes exchange their values. Process i

<p>LA_β for p_i</p> <pre> 1: $V_i := \{x_i\}$ // set of values, initially x_i 2: $F_i := \emptyset$ // set of known failure processes 3: $f :=$ the maximum number of failures 4: /* Exchange Values and Record Failures */ 5: Phase A: 6: Send V_i to all 7: for $j := 1$ to n 8: if V_j is received from process j 9: $V_i := V_i \cup V_j$ 10: else 11: $F_i := F_i \cup j$ 12: end for 13: /* LA with Known Height f */ 14: Phase B: 15: $F'_i := LA_\alpha(f, F_i)$ 16: $U_i :=$ values from processes in F'_i in <i>Phase A</i> 17: $C_i := V_i - U_i$ // set of correct values 18: $y_i := \sqcup \{v : v \in C_i\}$ </pre>	<p>LA_γ for p_i</p> <pre> 1: $v_i := x_i$ // input value 2: $decided := false$ 3: /* Exchange Values */ 4: Phase A: 5: Send v_i to all 6: for $j := 1$ to n 7: receive v_j from p_j 8: $v_i := v_i \sqcup v_j$ 9: end for 10: /* Guessing Height */ 11: Phase B: 12: $guess := 2$ // guess height 13: while ($\neg decided$) 14: $v_i := LA_\alpha(guess, v_i)$ 15: $guess := 2 * guess$ 16: end while 17: $y_i := v_i$ </pre>
---	---

■ **Figure 3** Algorithm LA_β .

■ **Figure 4** Algorithm LA_γ .

includes j into its failure set if it does not receive value from process j at the first phase. After the first phase, each process has a failure set which contains failed processes it knows. Then in *phase B*, they invoke algorithm LA_α with f as the height and its failure set as input. After that, each process *decides* on a failure set which satisfies lattice agreement properties. The new failure set of any two process i and j are comparable to each other, i.e., F'_i is comparable to F'_j . Equipped with this comparable failure set, each process removes values it received from processes which are in its failure set and decides on the join of the remaining values.

The following lemma shows that any two processes decide on comparable values. We only give the sketch of proof, and the detailed proof is available in the full paper.

► **Lemma 14.** (*Comparability*) *Let process i and j decide on y_i and y_j , respectively. Then y_i and y_j are comparable.*

Proof sketch. According to comparability of LA_α , all processes have comparable failure sets. Then, the set of values they received at *Phase A* from correct processes must be comparable, i.e., C_i is comparable with C_j . Therefore, y_i and y_j are comparable. ◀

► **Theorem 15.** *LA_β solves lattice agreement problem in $\log f + 1$ rounds, where $f < n$ is the maximum number of failures the system can tolerate.*

Proof. *Downward-Validity.* Initially, for correct process i , $v_i = x_i$. After *Phase A*, since i is correct, so i is not in any failure set of any process. At *Phase B*, process i invokes algorithm LA_α with failure set as the input value. Thus, according to the *Upward-Validity* of LA_α , i

is not included in F'_i . So, $x_i \in C_i$. Therefore, $x_i \leq y_i$. *Upward-Validity* is immediate from the fact that each process receives at most all values by all processes. *Comparability* follows from Lemma 14. ◀

3.2.2 Algorithm LA_γ

Algorithm LA_β solves lattice agreement in $\log f + 1$ rounds whereas Algorithm LA_α solves lattice agreement in $\log h(L)$ rounds assuming $h(L)$ is given. We now propose an algorithm to solve lattice agreement which has round complexity related to $h(L)$ even when $h(L)$ is not known. This algorithm called LA_γ (shown in Figure 4), solves the standard lattice agreement in $O(\min\{\log^2 h(L), \log^2 f\})$ rounds. The basic idea is to “guess” the height of L and apply algorithm LA_α using the guessed height as input. The algorithm is composed of two phases. At *Phase A*, each process simply broadcasts its value and takes the join of all received values. *Phase B* is the guessing phase which invokes algorithm LA_α repeatedly. Notice that *decided* variable is updated at *line 6* of LA_α .

Let w_i denote the value of v_i after *Phase A*. Let Ψ denote the sublattice formed by values of all correct processes after *Phase A*, i.e., $\Psi = \{u \mid (u \in L) \wedge (\exists i : w_i \leq u)\}$. Since there are at most f failures, we have $h(\Psi) \leq f$. Now we show that *Phase B* terminates in at most $\lceil \log h(\Psi) \rceil$ executions of LA_α . We call the i -th execution of LA_α as iteration i . Notice that the guessed height of iteration i is 2^i .

► **Lemma 16.** *After iteration $\lceil \log h(\Psi) \rceil$ of LA_α at *Phase B*, all processes decide.*

Proof. Since $2^{\lceil \log h(\Psi) \rceil} \geq h(\Psi)$, Lemma 9 still holds which implies Lemma 10. Thus, all undecided processes have equal values at the last round of iteration $\lceil \log h(\Psi) \rceil$. Therefore, all undecided processes decide after iteration $\lceil \log h(\Psi) \rceil$. ◀

We now show that two processes decide on comparable values irrespective of whether they both decide on the same iteration of LA_α .

► **Lemma 17.** (*Comparability*) *Let i and j be any two processes that decide on value y_i and y_j , respectively. Then y_i and y_j are comparable.*

Proof. Assume process i decides on G_i at round r_i of execution e_i of LA_α and process j decides on G_j at round r_j of execution e_j of LA_α . If $e_i = e_j$, then y_i and y_j are comparable by Lemma 8. Otherwise, $e_i \neq e_j$. Without loss of generality, suppose $e_i < e_j$. Consider round r_i of execution e_i of LA_α . Since i decides on value y_i at this round, then from Lemma 7, we have that y_i is comparable with v_k^r for any correct process k . Let $V = \{v_k^r \mid k \text{ is correct}\}$. Then, y_i is at most the join of a subset of V . From Lemma 6, it follows that y_i is comparable with y_j . ◀

► **Theorem 18.** *LA_γ solves the lattice agreement problem and can tolerate $f < n$ failures.*

Proof. *Downward-Validity* follows from that fact that the value of each process is non-decreasing along the execution. *Upward-Validity* follows since each process can receive at most all values from all processes. *Comparability* holds by Lemma 17. ◀

Complexity. From Lemma 16, we know *Phase B* terminates in at most $\lceil \log h(\Psi) \rceil$ executions of LA_α . Thus, *Phase B* takes $\log 2 + \log 4 + \dots + \lceil \log h(\Psi) \rceil = \frac{(\lceil \log h(\Psi) \rceil + 1) * (\lceil \log h(\Psi) \rceil)}{2}$ rounds in worst case. Since $h(\Psi) \leq f$ and $h(\Psi) \leq h(L)$, LA_γ has round complexity of $\min\{O(\log^2 h(L)), O(\log^2 f)\}$. Each process sends n messages at each round, thus message complexity is $n^2 \cdot \min\{O(\log^2 h(L)), O(\log^2 f)\}$.

<p>LA_δ for p_i</p> <p>$acceptVal := x_i$ // accept value $learnedVal := \perp$ // learned value</p> <p>on receiving $prop(v_j, r)$ from p_j:</p> <p>if $v_j \geq acceptVal$ Send $ACK("accept", -, r)$ $acceptVal := v_j$</p> <p>else Send $ACK("reject", acceptVal, r)$</p>	<p>for $r := 1$ to $f + 1$</p> <p>$val := acceptVal$ Send $prop(val, r)$ to all wait for $n - f$ $ACK(-, -, r)$ messages let V_r be values contained in <i>reject</i> $ACKs$ let $tally$ be number of <i>accept</i> $ACKs$ if $tally > \frac{n}{2}$ $learnedVal := val$ break</p> <p>else $acceptVal := acceptVal \sqcup \{v \mid v \in V_r\}$</p> <p>end for</p>
--	--

■ **Figure 5** Algorithm LA_δ .

4 Lattice Agreement in Asynchronous Systems

In this section, we discuss the lattice agreement problem in asynchronous systems. The algorithm proposed in [7] requires $O(n)$ units of time, whereas our algorithm (LA_δ shown in Figure 5) requires only $O(f)$ units of time. We first note that

► **Theorem 19.** *The lattice agreement problem cannot be solved in asynchronous message systems if $f \geq \frac{n}{2}$.*

Proof. The proof follows from the standard partition argument. If two partitions have incomparable values then they can never decide on comparable values. ◀

4.1 Algorithm LA_δ

On account of Theorem 19, we assume that $f < \frac{n}{2}$. The algorithm proceeds in *round-trips*. A single round-trip is composed of sending messages to all and getting $n - f$ acknowledgement messages back. At each round-trip, a process sends a *prop* message to all, with its current accepted value as the proposal value, and waits for $n - f$ ACK messages. If majority of these ACK messages are *accept*, then it decides on its current proposed value. Otherwise, it updates its current accept value to be the join of all values received and starts next round-trip. Whenever a process receives a proposal, i.e., a *prop* message, if the proposal has a value at least as big as its current value, then it sends back an ACK message with *accept* and updates its current accept value to be the received proposal value. Otherwise, it sends back an ACK message with *reject*.

Let $acceptVal_i^r$ denote the accept value (variable $acceptVal$) held by p_i at the beginning of round-trip r . Let $L^{(r)} = \{u \mid (u \in L) \wedge (\exists i : acceptVal_i^r \leq u)\}$, i.e., $L^{(r)}$ denotes the join-closed subset of L that includes the accept values held by all undecided processes at the beginning of the round-trip r . Notice that $L^{(1)} = L$.

► **Lemma 20.** *For any round-trip r , $h(L^{(r+1)}) < h(L^{(r)})$.*

Proof. If a process decides at round-trip r , its value is not in $L^{(r+1)}$. So, we only need to prove that $h(acceptVal_i^r) < h(acceptVal_i^{r+1})$ for any process i which does not decide at round-trip r . The fact that process i does not decide at round-trip r implies that i must have received at least one *reject* ACK with a greater value. Since $acceptVal_i^{r+1}$ is the join of all values received at round-trip r , $acceptVal_i^r < acceptVal_i^{r+1}$. Hence, $h(acceptVal_i^r) < h(acceptVal_i^{r+1})$ for any undecided process i . Therefore, $h(L^{(r+1)}) < h(L^{(r)})$. ◀

► **Lemma 21.** *All process decide within $\min\{h(L), f + 1\}$ asynchronous round-trips.*

Proof. We first show that $h(L^{(2)}) \leq f$. At the first round-trip, each process receives $n - f$ ACKs, which is equivalent to receiving $n - f$ values. Therefore, $h(L^{(2)}) \leq f$. Let $r_{min} = \min\{h(L), f + 1\}$. Combining the fact that $h(L^{(2)}) \leq f$ with Lemma 20, we have $h(L^{(r_{min})}) \leq 1$. This means that undecided correct processes have the same value. Thus, all of them receive $n - f$ ACK messages with *accept* and decide. Therefore, all processes decide within $\min\{h(L), f + 1\}$ round-trips. ◀

We note here that the algorithm in [7] takes $O(n)$ message delays for a value to be learned in the worst case. A crucial difference between LA_δ and the algorithm in [7] is that LA_δ starts with the accepted value as the input value. Hence, after the first round-trip, there is a significant reduction in the height of the sublattice, from n initially (in the worst case) to f . In [7], acceptors start with the accepted value as null. Hence, there is reduction of height by only 1 in the worst case. Since in their algorithm, acceptors are different from proposers (in the style of Paxos), acceptors do not have access to the proposed values.

► **Theorem 22.** *Algorithm LA_δ solves the lattice agreement problem in $\min\{h(L), f + 1\}$ round-trips.*

Proof. *Down-Validity* holds since the accept value is non-decreasing for any process i . *Upward-Validity* follows because each learned value must be the join of a subset of all initial values which is at most $\sqcup\{x_1, \dots, x_n\}$. For *Comparability*, suppose process i and j decide on values y_i and y_j . There must be at least one process that has accepted both y_i and y_j . Since each process can only *accept* comparable values. Thus, we have either $y_i \leq y_j$ or $y_j \leq y_i$. ◀

Complexity. From Lemma 21, we know that LA_δ takes at most $\min\{h(L), f + 1\}$ round-trips, which results in $2 \cdot \min\{h(L), f + 1\}$ message delays, since one round-trip takes two message delays. At each round-trip, each process sends out at most $2n$ messages. Thus, the number of messages for all processes is at most $2 \cdot n^2 \cdot \min\{h(L), f + 1\}$.

5 Generalized Lattice Agreement

In this section, we discuss the generalized lattice agreement problem as defined in Section 2.3. Since it is easy to adapt algorithms for lattice agreement in synchronous systems to solve generalized lattice agreement problem, we only consider asynchronous systems. We show how to adapt LA_δ to solve the generalized lattice agreement problem (algorithm GLA_α shown in Figure 6) in $\min\{O(h(L)), O(f)\}$ units of time.

5.1 Algorithm GLA_α

GLA_α invokes the `Agree()` procedure to learn a new value multiple times. The `Agree()` procedure is an execution of LA_δ with some modifications (to be given later). A sequence number is associated with each execution of the `Agree()` procedure, thus each correct process has a learned value for each sequence number. The basic idea of GLA_α is to let all processes sequentially execute LA_δ to learn values, and make sure: 1) any two learned values for the same sequence number are comparable, 2) any learned value for a bigger sequence number is at least as big as any learned value for a smaller sequence number. The first goal can be simply achieved by invoking LA_δ with the sequence number. In order to achieve the second goal, the key idea is to make any proposal for sequence number $s + 1$ to be at least as big

as the largest learned value for sequence number s . Notice that at each round-trip of LA_δ execution, a process waits for $n - f$ *ACKs*, and any two set of $n - f$ processes have at least one process in common. Thus, the second goal can be achieved by making sure at least $n - f$ processes know the largest learned value after execution of LA_δ for a sequence number.

Upon receiving a value v from client in a message tagged with *ClientValue*, a process adds v into its buffer and sends a *ServerValue* message with v to all other processes. The process can start to learn new values only when it succeeds at its current proposal. Otherwise, LA_α may not terminate, as shown by an example in [7]. Upon receiving a *ServerValue* message with value v , a process simply adds v to its buffer.

The *Agree()* procedure is automatically executed when the guard condition is satisfied; that is, it is not currently proposing a value and it has some value in its buffer or it has seen a sequence number bigger than its current sequence number. Inside the *Agree()* procedure, a process first updates its *acceptVal* to be the join of current *acceptVal* and *buffVal*. Then, it starts an adapted LA_δ execution. The original LA_δ and adapted LA_δ differ in the following ways: 1) Each message in the adapted LA_δ is associated with a sequence number. 2) A process can also decide on a value for a sequence number if it receives any *decide ACK* message for that sequence number. 3) On receiving a *prop* message associated with a sequence number s' , if s' is smaller than its current sequence number which means it has learned a value for s' , then it simply sends *ACK* message with its learned value for s' back. If s' is greater than its current sequence number, it updates its *maxSeq* and waits until its current sequence number matches s' . After that it sends back *ACK* message with *accept* or *reject* based on whether the proposal value is bigger than its current *accept* value or not. The reason a process keeps track of the maximal sequence number it has ever seen, is to make sure each process has a learned value for each sequence number. When the maximum sequence number is bigger than its current sequence number, it has to invoke *Agree()* procedure even if it does not have any new value to propose. After execution of adapted LA_δ , a process increases its current sequence number.

We next show the correctness of GLA_α . Let $acceptVal_s^p$ denote the *acceptVal* of process p at the end of *Agree()* procedure for sequence number s . Let LV_p denote the map of sequence number to learned value (variable *LV*) for process p and $m_s = \sqcup\{LV_p[s] : p \in [1..n]\}$, i.e., m_s denotes the join of all learned values for sequence number s . Let $LP_s = \{p \mid (p \in [1..n]) \wedge (m_s \leq acceptVal_s^p)\}$, i.e., LP_s is the set of processes which have *acceptVal* greater than the join of all learned values for the sequence number s . Notice that a process has two ways to learn a value for its current sequence number in the *Agree()* procedure: 1) by receiving a majority of *accept ACKs*. 2) by receiving some *decide ACKs*.

The following lemma proves that the adapted LA_α satisfies the first goal.

► **Lemma 23.** *For any sequence number s , $LV_p[s]$ is comparable with $LV_q[s]$ for any two processes p and q .*

Proof. We only need to show that any two processes which learn by the first way must learn comparable values, since processes which learn by the second way simply learn values from processes which learn by the first way. By the same reasoning as *Comparability* of Theorem 22, we know this is true. ◀

From Lemma 23, we know that m_s is the largest learned value for sequence number s .

► **Lemma 24.** *For any sequence number s , $|LP_s| > \frac{n}{2}$.*

<p>GLA_α for p_i</p> <p>$s := 0$ // sequence number $maxSeq := -1$ // max seq number seen $buffVal := \perp$ // received values /* map from seq to learned value */ $LV := \perp$ $acceptVal := \perp$ $active := false$</p> <p>on receiving $ClientValue(v)$: $buffVal := buffVal \sqcup v$ Send $ServerValue(v)$ to all</p> <p>on receiving $ServerValue(v)$: $buffVal := buffVal \sqcup v$</p> <p>on receiving $prop(v_j, r, s')$ from p_j: if $s' < s$ Send $ACK("decide", LV[s'], r, s')$ break if $s' > s$ $maxSeq := \max\{s', maxSeq\}$ wait until $s = s'$ if $v_j \geq acceptVal$ Send $ACK("accept", -, r, s')$ $acceptVal := v_j$ else Send $ACK("reject", acceptVal, r, s')$</p>	<p>Procedure Agree(): guard: $(active = false)$ $\wedge (buffVal \neq \perp \vee maxSeq \geq s)$</p> <p>effect: $active := true$ $acceptVal := buffVal \sqcup acceptVal$ $buffVal := \perp$</p> <p>/* LA_δ with sequence number */ for $r := 1$ to $f + 1$ $val := acceptVal$ Send $prop(val, r, s)$ to all wait for $n - f$ $ACK(-, -, r, s)$ let V be values in <i>reject ACKs</i> let D be values in <i>decide ACKs</i> let $tally$ be number of <i>accept ACKs</i> if $D > 0$ $val := \sqcup\{d \mid d \in D\}$ break else if $tally > \frac{n}{2}$ break else $acceptVal := acceptVal$ $\sqcup\{v \mid v \in V\}$ end for $LV[s] := val$ $s := s + 1$ $active := false$</p>
--	--

■ **Figure 6** Algorithm GLA_α .

Proof. Consider Agree() procedure for s . Since m_s is the largest learned value for sequence number s , there must exist a process p which learns m_s by the first way. Thus, p must have received a majority of *accept ACKs*, which means at least a majority of processes have *acceptVal* greater than m_s after Agree() procedure for s . Therefore, $|LP_s| > \frac{n}{2}$. ◀

The lemma below shows that GLA_α achieves the second goal.

► **Lemma 25.** $m_s \leq LV_p[s + 1]$ for any process p and any sequence number s .

Proof. From Lemma 24, we know for sequence number s at least a majority of processes have *acceptVal* greater than m_s . To decide on $LV_p[s + 1]$, process p must get majority *accept*. Since any two majority has at least one process in common, $m_s \leq LV_p[s + 1]$. ◀

► **Theorem 26.** Algorithm GLA_α solves generalized lattice agreement when a majority of processes is correct.

Proof. *Validity* holds since any learned value is the join of a subset of values received.

Stability. From Lemma 25 and the fact that $LV_p[s] \leq m_s$, we have that $LV_p[s] \leq LV_p[s+1]$ for any process p and any sequence number s , which implies *Stability*.

Comparability. We need to show that $LV_p[s]$ and $LV_q[s']$ are comparable for any two processes p and q , and for any two sequence number s and s' . If $s = s'$, this is immediate from Lemma 23. Now consider the case when $s \neq s'$. Without loss of generality, assume $s < s'$. From Lemma 25, we can conclude that $LV_p[s] \leq LV_q[s']$. Thus, *comparability* holds.

Liveness. Any received value v is eventually included in some proposal, i.e. *prop* message. From Theorem 22, we know that in at most $2 \cdot \min\{h(L), f + 1\}$ message delays that proposal value will be included in some learned value. Thus, v will be learned eventually. ◀

Complexity. For time complexity, from the analysis for *liveness* in Theorem 26, we know that a received value is learned in at most $2 \cdot \min\{h(L), f + 1\}$ message delays. For message complexity, since each process sends out n messages per *round-trip*, the total number of messages needed to learn a value is $2 \cdot n^2 \cdot \min\{h(L), f + 1\}$.

6 Conclusions

We have presented algorithms for the lattice agreement problem and the generalized lattice agreement problem. These algorithms achieve significantly better time complexity than previous algorithms. For future work, we would like to know the answers to the following two questions: 1) Is $\log f$ rounds a lower bound for lattice agreement in synchronous message passing systems? 2) Is $O(f)$ message delays optimal for the lattice agreement and generalized lattice agreement problem in asynchronous message passing systems?

References

- 1 Yehuda Afek, Hagit Attiya, Danny Dolev, Eli Gafni, Michael Merritt, and Nir Shavit. Atomic snapshots of shared memory. *Journal of the ACM (JACM)*, 40(4):873–890, 1993.
- 2 Hagit Attiya, Maurice Herlihy, and Ophir Rachman. Atomic snapshots using lattice agreement. *Distributed Computing*, 8(3):121–132, 1995.
- 3 Hagit Attiya and Ophir Rachman. Atomic snapshots in $\mathcal{O}(n \log n)$ operations. *SIAM Journal on Computing*, 27(2):319–340, 1998.
- 4 Hagit Attiya and Jennifer Welch. *Distributed computing: fundamentals, simulations, and advanced topics*, volume 19. John Wiley & Sons, 2004.
- 5 Carole Delporte-Gallet, Hugues Fauconnier, Sergio Rajsbaum, and Michel Raynal. Implementing snapshot objects on top of crash-prone asynchronous message-passing systems. In *International Conference on Algorithms and Architectures for Parallel Processing*, pages 341–355. Springer, 2016.
- 6 Danny Dolev and H Raymond Strong. Authenticated algorithms for byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983.
- 7 Jose M Faleiro, Sriram Rajamani, Kaushik Rajan, G Ramalingam, and Kapil Vaswani. Generalized lattice agreement. In *Proceedings of the 2012 ACM symposium on Principles of distributed computing*, pages 125–134. ACM, 2012.
- 8 Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2):374–382, 1985.
- 9 Maurice 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.

- 10 Leslie Lamport. The part-time parliament. *ACM Transactions on Computer Systems (TOCS)*, 16(2):133–169, 1998.
- 11 Leslie Lamport et al. Paxos made simple. *ACM Sigact News*, 32(4):18–25, 2001.
- 12 Marios Mavronicolas. A bound on the rounds to reach lattice agreement, 2000. URL: <http://www.cs.ucy.ac.cy/~mavronic/pdf/lattice.pdf>.
- 13 Michel Raynal. *Concurrent programming: algorithms, principles, and foundations*. Springer Science & Business Media, 2012.
- 14 Fred B Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys (CSUR)*, 22(4):299–319, 1990.
- 15 Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. Conflict-free replicated data types. In *Symposium on Self-Stabilizing Systems*, pages 386–400. Springer, 2011.
- 16 Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. Convergent and commutative replicated data types. *Bulletin-European Association for Theoretical Computer Science*, 104:67–88, 2011.
- 17 Andrew S Tanenbaum and Maarten Van Steen. *Distributed systems: principles and paradigms*. Prentice-Hall, 2007.
- 18 Gadi Taubenfeld. *Synchronization algorithms and concurrent programming*. Pearson Education, 2006.