

Scalable Agreement Protocols with Optimal Optimistic Efficiency

Yuval Gelles*

Ilan Komargodski†

Abstract

Designing efficient distributed protocols for various agreement tasks such as Byzantine Agreement, Broadcast, and Committee Election is a fundamental problem. We are interested in *scalable* protocols for these tasks, where each (honest) party communicates a number of bits which is sublinear in n , the number of parties. The first major step towards this goal is due to King et al. (SODA 2006) who showed a protocol where each party sends only $\tilde{O}(1)$ bits throughout $\tilde{O}(1)$ rounds, but guarantees only that $1 - o(1)$ fraction of honest parties end up agreeing on a consistent output, assuming constant $< 1/3$ fraction of static corruptions. Few years later, King et al. (ICDCN 2011) managed to get a full agreement protocol in the same model but where each party sends $\tilde{O}(\sqrt{n})$ bits throughout $\tilde{O}(1)$ rounds. Getting a full agreement protocol with $o(\sqrt{n})$ communication per party has been a major challenge ever since.

In light of this barrier, we propose a new framework for designing efficient agreement protocols. Specifically, we design $\tilde{O}(1)$ -round protocols for all of the above tasks (assuming constant $< 1/3$ fraction of static corruptions) with optimistic and pessimistic guarantees:

- **Optimistic complexity:** In an honest execution, (honest) parties send only $\tilde{O}(1)$ bits.
- **Pessimistic complexity:** In any other case, (honest) parties send $\tilde{O}(\sqrt{n})$ bits.

Thus, all an adversary can gain from deviating from the honest execution is that honest parties will need to work harder (i.e., transmit more bits) to reach agreement and terminate. Besides the above agreement tasks, we also use our new framework to get a scalable secure multiparty computation (MPC) protocol with optimistic and pessimistic complexities.

Technically, we identify a relaxation of Byzantine Agreement (of independent interest) that allows us to fall-back to a pessimistic execution in a coordinated way by all parties. We implement this relaxation with $\tilde{O}(1)$ communication bits per party and within $\tilde{O}(1)$ rounds.

*School of Computer Science and Engineering, Hebrew University of Jerusalem, Israel. Email: yuval.gelles@mail.huji.ac.il.

†School of Computer Science and Engineering, Hebrew University of Jerusalem, Israel, and NTT Research. Email: ilank@cs.huji.ac.il. Incumbent of the Harry & Abe Sherman Senior Lectureship at the School of Computer Science and Engineering at the Hebrew University.

Contents

1	Introduction	1
1.1	Related Work	3
2	Technical Overview	5
2.1	Optimistic/Pessimistic Agreement	5
2.2	A Best-Possible Optimistic/Pessimistic Framework	6
2.3	Secure MPC	7
3	The Model and Problem Definitions	7
3.1	The Model	7
3.2	Problem Definitions	7
4	Preliminaries	8
4.1	Secure Computation	8
4.2	(Attacker-Dependent) Complexity Measures of Protocols	10
5	Propagating the Detection of Failures	11
5.1	Failure Detection Functionalities	11
5.2	Propagating Knowledge of Failures	12
5.3	Byzantine Agreement for Failures	12
5.4	Proof of Theorem 5.2	16
5.4.1	Realizing \mathcal{F}^{gFD} Against \mathcal{A}^*	17
5.4.2	Realizing \mathcal{F} Against \mathcal{A}	18
6	Optimistic/Pessimistic Combiner Framework	18
7	Applications	20
7.1	A Quorum	20
7.2	Byzantine Agreement, Broadcast, and Committee Election	21
7.3	Secure Multiparty Computation	21
7.3.1	Security	23
7.3.2	Efficiency	25
7.3.3	The Gate Wrapper	26
	References	26
A	Distribute Sub-Protocol	30
B	Quorum Building Protocol	30

1 Introduction

We propose a new framework for designing efficient fault-tolerant distributed algorithms for large scale networks. Fault-tolerance means that the functionality of the protocol should not be compromised even if some of the participants in the protocol collude and arbitrarily deviate from the protocol’s specification. Basic and very well-studied abstractions in this context are Byzantine Agreement (BA) [PSL80, LSP82], Broadcast, and Committee Election. These serve as building blocks in many distributed protocols. Perhaps most notably, these protocols underlie essentially all secure (cryptographic) multiparty computation protocols [BGW88, GMW87, CCD87].

We are interested in designing protocols for all of the above tasks that can be used in large scale settings, where the number of parties could potentially be really large. Let n be the number of parties involved. For protocols to be scalable, we want the amount of work or communication required by each party to be sub-linear in n . Even for the very basic tasks, such as BA, this question was out of reach for many years. The situation changed with the groundbreaking work of [KSSV06] that gave a technique for “almost everywhere” *scalable* agreement. Using their technique, it is possible to solve BA, Broadcast, and Committee Election in a scalable fashion with the caveat that only $1 - o(1)$ fraction of the honest parties agree on the output of the protocol. In terms of efficiency though, their protocol is essentially optimal: it requires only $\tilde{O}(1)$ communication rounds and $\tilde{O}(1)$ communication per party.¹

The model that is considered in [KSSV06] is point-to-point synchronous communication and with a computationally unbounded Byzantine (a.k.a malicious) adversary that controls a $(1/3 - \epsilon)$ -fraction of the parties for any $\epsilon > 0$. Corruptions occur statically, namely, before the protocol begins (but after it is specified). Further, the adversary has full-information, namely, it sees all messages sent, even messages sent between two honest party, and it is rushing, namely, it gets to send its messages after seeing the honest party’s messages for that round. This is the model that we consider in this work as well; see Section 3.1 for the full specification of the model.

Since [KSSV06]’s work, there has been an effort to boost their almost everywhere agreement property into full agreement in various different settings [KS09, KS10, KLST11, BCG21]. In the classical information theoretic setting, the work of [KLST11] managed to get full agreement protocols for all of the above tasks at the cost of increased overhead: each party’s communication is $\tilde{O}(\sqrt{n})$ bits (the round complexity remains $\tilde{O}(1)$). More than a decade since this work, it is still essentially the state of the art (excluding protocols relying on cryptographic assumptions). The work of [HKK08] proved a lower bound saying that $\tilde{\Omega}(\sqrt[3]{n})$ communication is necessary for at least one party for a certain (non-trivial) class of protocols. See Section 1.1 for more details. Thus, it is still major open problem to get a full agreement protocol with $o(\sqrt{n})$ communication per party.

Optimistic/pessimistic efficiency. Given the lack of progress in getting better agreement protocols in the worst-case, we suggest a new approach for designing protocols. Specifically, we consider protocols that could have different complexities, depending on the attacker’s actions. It is important to emphasize that we want protocols that are correct/secure no matter what; i.e., all honest parties terminate and full agreement is reached. Thus, all an adversary can gain from deviating from the honest execution is that honest parties will need to work harder (i.e., transmit more bits) to reach agreement and terminate. So, it is conceivable to assume that in some applications an attacker will have no incentive to execute an attack. That is, if all the adversary cares about is breaking correctness/security (and gains nothing from a delay), then there is no point to deviate from the protocol and much better efficiency is obtained in this case. We state our main

¹We use the notation $\tilde{O}(\cdot), \tilde{\Omega}(\cdot)$ to hide poly-logarithmic factors in n .

result next.

Theorem 1.1 (Scalable agreement compiler; Informal). *Let $X \in \{\text{ByzantineAgreement}, \text{Broadcast}, \text{CommitteeElection}\}$ be a task. Assume that there is a protocol Π for X tolerating $1/3 - \epsilon$ fraction of static corruptions for any $\epsilon > 0$. Then, there is another protocol Π' for X (tolerating $1/3 - \epsilon$ fraction of static corruptions for any $\epsilon > 0$) with the following features:*

Optimistic Complexity: *In an honest execution, each party sends $\tilde{O}(1)$ bits and the protocol terminates within $\tilde{O}(1)$ rounds.*

Pessimistic Complexity: *In the worst-case, the communication and round complexities of Π' is the same as that of Π plus $\tilde{O}(1)$.*

We remark that while we mentioned above that our protocols obtain optimistic efficiency ($\tilde{O}(1)$ communication overall per party) in honest executions, in fact, this is the case even for slightly misbehaving attackers. Concretely, we provide the same efficiency guarantees even for attackers that follow the protocol but can omit some of their outgoing messages (this includes crash failures).

As a corollary, using the protocols of [KSSV06, KLST11] as Π , we get the following result.

Corollary 1.2 (Scalable agreement instantiation; Informal). *There are $\tilde{O}(1)$ -round Byzantine Agreement, Committee Election, and Broadcast protocols tolerating $1/3 - \epsilon$ fraction of static corruptions for any $\epsilon > 0$ with the following features:*

Optimistic Complexity: *In an honest execution, each party sends $\tilde{O}(1)$ bits.*

Pessimistic Complexity: *In the worst-case, each honest party sends $\tilde{O}(\sqrt{n})$ bits.*

The $\tilde{O}(\sqrt{n})$ term in the pessimistic complexity bullet comes from the cost of [KLST11]’s protocol. Any improvement on the latter will immediately translate to improved pessimistic complexity.

Technical highlight: agreement with error detection. The main building block of the above theorem is a new protocol for a relaxed variant of an agreement functionality. Recall that in an agreement functionality the goal is to guarantee that all parties terminate and agree on some specified value. We introduce a relaxation of the above requirement by requiring that (1) in an honest execution, indeed all parties terminate and reach agreement, but (2) in all other cases, all parties should agree on a special Fail symbol. With this abstraction, we can identify failure and then invoke another (expensive) agreement protocol. To get Theorem 1.2, we implement such a relaxed agreement protocol using only $\tilde{O}(1)$ communication and rounds.

Generic optimistic/pessimistic framework. More generally, we use our agreement with error detection protocol to obtain a generic framework for combining efficient optimistic protocols with less efficient pessimistic protocols. Specifically, we manage to combine an efficient protocol Π_{light} (solving a given task X) where it is guaranteed that a failure was noticed by at least one party with a less efficient protocol Π_{heavy} (solving the same task X) that is executed only if everyone knows that some failure occurred. The complexity of the combined protocol is inherited from Π_{light} in an honest execution and from Π_{heavy} in any other case. We refer to Section 6 for full details.

Application to Scalable MPC. The above agreement tasks can be thought of as special cases of secure multi-party computation (MPC) [GMW87, BGW88, CCD87]. MPC protocols enable a set of mutually distrusting parties to compute a function on their private inputs, while guaranteeing various properties such as correctness, privacy, independence of inputs, and more. We consider the problem of scalable MPC in the peer-to-peer synchronous communication model with private channels.

Feasibility results for (non-scalable) MPC have been long known, e.g., the BGW [BGW88] protocol gives a method for computing an arbitrary function with communication cost that grows multiplicatively with the circuit size of the function and some polynomial in the number of parties. That is, the communication complexity of each of the n parties is $s \cdot \text{poly}(n)$ when computing a function represented as a circuit of size s .² The question of scalable MPC, i.e., protocols where the dominant term in the complexity is just the circuit size, is still an active topic and modern results achieve MPC protocols with strong security guarantees and communication complexity $\tilde{O}(s + \text{poly}(n))$ (see Section 1.1 for an overview and references).

We use our scalable agreement protocols to obtain a new generic scalable MPC. The properties of the resulting protocol are summarized in the next theorem.

Theorem 1.3 (Scalable MPC; Informal). *There is a statistically maliciously secure MPC protocol tolerating $1/3 - \epsilon$ fraction of static corruptions for any $\epsilon > 0$ with the following features. Given a circuit of size s and depth d over n inputs, the protocol has the following complexity:*

Optimistic Complexity: *In an honest execution, each party sends $\tilde{O}(s/n)$ bits.*

Pessimistic Complexity: *In the worst-case, each party sends $\tilde{O}(s/n + \sqrt{n})$ bits.*

Round Complexity: *All parties terminate after $\tilde{O}(d)$ rounds.*

The additive $\tilde{O}(\sqrt{n})$ term in the pessimistic complexity bullet comes generically from Corollary 1.2. At a high level, the above MPC is obtained by using our agreement protocol to generate a *quorum*: assign to each party its own representative (small and balanced) committee where there is a strong majority of honest parties. Then, we distribute the gates of the circuit to these committees. Each gate is evaluated by its assigned committees using some standard MPC (e.g., BGW). We emphasize that our protocol has the appealing feature that in an honest execution the total communication complexity is essentially equal to the circuit size, for any circuit, and it is split in a balanced manner across parties. This idea largely appeared in the scalable MPC protocol of [DKM⁺17]. The main difference is that we obtain optimistic/pessimistic complexity, while they only had pessimistic complexity, and this is due to the use of Corollary 1.2 instead of the protocol of [KLST11].

1.1 Related Work

Scalable agreement. The BA problem was introduced in the landmark work of Lamport, Shostak, and Pease [LSP82]. In the following couple of decades, several protocols were presented (e.g., [DS83, DLS88, CL99]) but they all had quadratic total overhead, that is, every party had to essentially communicate with every other party. The protocol of [KSSV06] was the first to break this barrier, but it had the caveat of almost-everywhere agreement (rather than full agreement). Extending their almost-everywhere agreement to full agreement in a scalable manner and with

²Interestingly, in BGW it was already observed that their protocol has an optimistic/pessimistic flavor where in the former the polynomial in n is slightly better than in the pessimistic case.

minimal cost is still an exciting challenge. We mention some of the key papers addressing this challenge.

First, [KS09] presented a protocol that satisfies full agreement but it is not balanced. That is, while most parties do communicate a sublinear amount of bits in n overall, there are few parties that communicate essentially with everyone. Several follow up works (e.g., [BGH13, ACD⁺19]) suffer from the same issue. Then, [KLST11] presented a protocol that satisfies full agreement and has balanced communication, but this comes with an extra $\tilde{O}(\sqrt{n})$ term in communication cost.

The extra cost in efficiency is partially explained by an $\Omega(\sqrt[3]{n})$ lower bound on the communication complexity of at least one party in any BA protocol with full agreement, due to [HKK08]. This lower bound, however, applies only to protocols with *static filtering*. In static filtering, every party decides on the set of parties it will listen to before the beginning of each round (as a function of its internal view at the end of the previous round). It is an intriguing open problem to extend the lower bound beyond protocol with static filtering rules.

Lastly, we mention a work of [BCG21]. They assume cryptographic and setup assumptions and further that the adversary is computationally bounded. Also, they assume dynamic filter—namely, the decision of which message is received can be based on the content of received messages (in their case, every message is checked if it contains a valid digital signature). With these relaxations of the model, no lower bound is known. They showed a communication-optimal protocol: only $\tilde{O}(1)$ bits of communication per party are needed to reach full agreement.³

We remark that all of the above works, as well as ours, assume near-optimal resilience, i.e., up to $(1/3 - \epsilon)$ fraction of corruptions (i.e., near-optimal resilience range). Less than $n/3 - 1$ corruptions is strictly necessary due to lower bounds of [LSP82, FLM86, Bor96] (unless further assumptions are made such as a trusted setup or a computationally bounded adversary).

Scalable MPC. There has been a rich line of work on scalable MPC protocol. The main goal is to design protocols where the total communication complexity scales like $O(s + \text{poly}(n))$ for securely computing a size s circuit by n parties. This was studied in the context of perfect or statistical security and optimal resilience (up to $n/3 - 1$ or $n/2 - 1$ corrupted parties) [HM01, DN07, BH08, GIP⁺14, IKP⁺16, CCXY18, CGH⁺18], or with perfect or statistical security and near-optimal resilience (up to $(1/3 - \epsilon)$ or $(1/2 - \epsilon)$ fraction of corrupted parties) [DI06, IPS09, DIK⁺08, DIK10]. In all of these works a broadcast channel is assumed but its usage is limited to a number of times which is independent of the circuit size. All of these works obtain somewhat stronger notions of security than what we obtain in Theorem 1.3 (e.g., they often tolerate adaptive corruptions while we handle only static ones). Note that we can use our broadcast protocol from Corollary 1.2 to instantiate the broadcast channel in the above works, achieving statistical security for $(1/3 - \epsilon)$ fraction of static corruptions.

Most related to us are the works [DKMS12, DKM⁺17]. In these works the authors used the everywhere agreement protocol of [KLST11] to get a scalable MPC with complexity $\tilde{O}(s/n + \sqrt{n})$ to compute a size s circuit by n parties, per party. The corruption model is $(1/3 - \epsilon)$ fraction static corruptions, same as ours. Our MPC protocol is very similar to theirs (associating the wires of the circuit to quorum members); but, our description is somewhat simpler because we use generic maliciously secure MPC as black-box whereas they sometime use internal building blocks such as verifiable secret sharing. The optimistic/pessimistic aspect is new to our work. Lastly, we mention the work of [DKMS14] who studied scalable MPC protocols in an asynchronous setting.

³[BCG21] have two variants presenting tradeoffs between the cryptographic assumptions and the trusted setup assumptions. Either a weaker trusted setup assumption (a public-key infrastructure and a common random string) and a stronger cryptographic assumption (SNARKs with linear-time extraction and a collision resistant hash).

2 Technical Overview

In this section we give a high level overview of our protocols. We start by explaining the underlying idea of Theorem 1.1 and Corollary 1.2, an agreement protocol with optimistic/pessimistic efficiency features. We then continue to explain our optimistic/pessimistic framework for general protocols, and finally we explain our application for secure multiparty computation, Theorem 1.3.

2.1 Optimistic/Pessimistic Agreement

Recall that in an agreement protocol, abstractly, the goal is for all honest parties to terminate and agree on some value. Specifically, we will consider the task of agreeing on a long enough string with high enough min-entropy—this will be sufficient to obtain all of our applications, including Byzantine Agreement, Broadcast, and Committee Election.

Our starting point is the almost-everywhere agreement protocol of [KSSV06]. In their protocol, the above task is met by almost all parties. That is, after poly-logarithmically many rounds of communication, a poly-logarithmically long random string str with poly-logarithmic min-entropy is known to $1 - o(1)$ fraction of parties. The goal is to “boost” the knowledge of this string to all other honest parties as well. At this point, a naive attempt would be to let every party poll sufficiently many random parties and ask their opinion about str . Then, each party will decide on a string that it obtained from more than $1/2$ of the polled parties. By a simple application of Chernoff’s bound, it suffices to poll independently poly-logarithmically many random parties.

This solution, however, is flawed because an adversary can easily prevent an honest party from learning the right value. Specifically, assume that honest party P that *does not know* str polled parties P_1, \dots, P_ℓ . The adversary can *flood* all parties in P_1, \dots, P_ℓ and thereby prevent them from replying to P . In such case, P will not learn str . At this point, previous works, e.g., [KLST11], suggested various techniques limiting the adversary’s ability to perform such flooding attacks, but this came with a significant cost in efficiency, requiring $\tilde{O}(\sqrt{n})$ communication.

We deviate from all of these previous works by making the following critical observation: While the flooding attack does prevent P from learning str , it does provide P (and some of its poll list members) with important knowledge: *they were attacked!* At this point we would have been happy to merely fall-back to one of the known agreement protocols (such as [KLST11] itself) and pay the extra cost in efficiency. But, we are in a state of inconsistency: some of the parties know str and think they are done while others do not know str and want to fall back to a heavy protocol.

It may seem, at first sight, that we have not gained much because we are again in a situation where some of the parties are not aware of the right state of the protocol and so we may need to perform some heavy agreement protocol. But, a closer look reveals that we have actually made a lot of progress:

1. A party that did not know str after the almost-everywhere agreement either knows str or knows the polling phase failed.
2. If the polling phase failed for at least one party, there must exist at least one party that did know str *and* learned that the protocol failed.

Our goal now is to propagate the knowledge of failure to everyone so that everyone will consistently fall back to a heavy agreement protocol. This task is much simpler than agreeing on str to begin with because (a) all the parties that know that the protocol failed, we do not care about them, and (b) all other parties agree on str . Thus, we remain with the following simpler task: parties need to perform agreement given a shared sufficiently long high-entropy string str .

Semi-random string to a quorum. A “good” quorum is a collection of well-balanced “good” committees, one per party. A “good” committee is a collection of $c = \text{polylog}(n)$ parties where at least $2/3$ of them are honest. A good quorum consists of n such good committees, one per party, and where no single party participates in too many committees. Of course, a uniformly random collection of n small subsets of parties will be a good quorum, but we do not have a uniformly random shared string. Here, we use a probabilistic lemma saying that a semi-random string suffices to sample a good quorum with high probability. This lemma was proven via the probabilistic method in [KLST11].

Given a quorum, it is trivial to make sure that if one of the parties has a Fail input, then all parties will know it. The party that holds Fail will broadcast it to its committee and then the committees will propagate this in a tree-like fashion and inform all other committees of this. Finally, each committee will send the Fail symbol to its party.

Applications. Once we obtain a quorum, we can get all of the applications in a straight-forward way, similarly to how we disseminated the Fail symbol. For example, performing broadcast can be done by the following simple protocol: the broadcaster sends its message to its committee and then the committee broadcasts this value in a tree-like fashion among all other committees; the latter send this value to their respective parties. (By saying that a committee sends or received a value, we mean that more than $2/3$ of the parties in the committee send or receive it.) Similar protocols can be described for Byzantine Agreement and Committee Election; see Section 7 for details.

2.2 A Best-Possible Optimistic/Pessimistic Framework

We generalize the above idea and present a general-purpose framework for building protocols with an optimistic/pessimistic complexity. Specifically, consider a protocol Π_{light} with following properties:

- In an honest execution, it achieves some goal.
- In any other case, at least one party knows that something went wrong (but privacy, if considered, was not broken).

Additionally, we have a protocol Π_{heavy} that achieves the same goal, no matter what the adversary does.

We would like to combine these two protocol to a single protocol where we first run Π_{light} and then, if something went wrong, run Π_{heavy} . Unfortunately, this is impossible since the state of all parties at the end of Π_{light} is not consistent: not all parties know that something went wrong. Thus, we use the idea from above (the optimistic/pessimistic agreement) to design an agreement protocol for failures, which we call `BYZANTINEAGREEMENTFORFAILURES`.

This is a special case of Byzantine Agreement, where importantly, if at least one honest party has Fail as input, all parties agree on Fail. But, we do not have any requirement in any other case except that we require that in an honest execution, if all parties have Success, then all parties will agree on Success. How do we build this protocol? This is exactly the agreement protocol we described above, without the fall-back part: run almost-everywhere agreement to generate a string, poll poly-logarithmically many parties for their string, and finally perform Byzantine Agreement via constructed quorums. This sub-protocol costs $\tilde{O}(1)$ rounds and communication.

To summarize, our combined protocol has the following form:

1. Run Π_{light} . Each party has an output y_i and a symbol $o_i \in \{\text{Fail}, \text{Success}\}$.
2. Run `BYZANTINEAGREEMENTFORFAILURES` on the o_i 's to get o_i^* 's.

3. For each party P_i , if $o_i^* = \text{Fail}$, participate in Π_{heavy} ; Otherwise, output y_i .

Observe that in an honest execution, Π_{heavy} will never be invoked and so the complexity of the protocol is similar to the complexity of Π_{light} plus $\tilde{O}(1)$ (the complexity of `BYZANTINEAGREEMENTFORFAILURES`). If a failure is detected, either during Π_{light} or during `BYZANTINEAGREEMENTFORFAILURES`, Π_{heavy} is invoked and its complexity dominates. Refer to Section 6 for the precise statement and full details.

2.3 Secure MPC

The idea to obtain an MPC protocol is simple given the quorum structure we obtained above. Following [DKM⁺17], once we have a good quorum, we associate each wire to some committee, in a balanced way (i.e., if there are s wires, each committee is in charge of $\approx s/n$ wires). We compute the circuit gate by gate, basically by relying on some known perfectly secure MPC to compute each gate and provide input for the next gate. Notice that no committee member holds the wire, but instead they have it in a secret shared fashion, and so when we process a gate, we first reconstruct the input wires, and at the end share the outputs to the appropriate committee members. All of these details are taken care of by a generic MPC, such as the one of BGW [BGW88]. Since each committee has a strong majority of honest parties and since each committee consists of poly-logarithmically many parties, this MPC costs $\tilde{O}(1)$ communication and rounds.

3 The Model and Problem Definitions

3.1 The Model

Communication model. We consider n parties in a fully connected network (i.e., clique). Each party has a unique ID and the IDs are common knowledge. We assume the IDs are $1, \dots, n$. Parties can perform arbitrary computation and they have a source of private randomness. Communication is authenticated, that is, whenever a party sends a message directly to another, the identity of the sender is known to the recipient. We assume synchronous communication. That is, communication proceeds in rounds; messages are all sent out at the same time at the start of the round, and then received at the same time at the end of the same round. All parties have synchronized clocks.

Adversarial model (point-to-point full information). We assume that there is an adversary that controls up to t parties and whose goal is to cause the protocol to fail in some way, depending on the context. The adversary chooses which parties to control non-adaptively, i.e., it chooses the set of corrupted parties at the start of the protocol. The adversary is malicious: corrupted parties can engage in any kind of deviations from the protocol and send arbitrary messages in the name of the corrupted parties. We emphasize that corrupted parties can send arbitrarily long messages to arbitrary other parties. We assume that the adversary is *rushing*, that is, it can view all messages sent by the honest parties in a round before the corrupted parties send their messages in the same round. Moreover, the adversary sees all messages, even messages sent between two honest parties (but honest parties only see messages that are sent directly to them). This model is known as the *point-to-point full information model* [KSSV06].

3.2 Problem Definitions

Byzantine agreement. In this problem, there are n parties, P_1, \dots, P_n . At most t of the parties may be corrupted. Each party P_i begins with an input bit $x_i \in \{0, 1\}$ and outputs a bit y_i . The goal

of the protocol is (with high probability) for all honest parties to terminate and, upon termination, agree on a bit held by at least one honest party at the start. This should hold even when the t corrupted parties collude and actively try to prevent it. More precisely, the Byzantine agreement problem is defined as a protocol that satisfies with high probability the following *agreement*, *validity*, and *termination* properties:

- **Agreement:** For every pair of honest parties P_i and P_j it holds that $y_i = y_j$.
- **Validity:** If there exists a bit x such that for every honest party P_i it holds that $x_i = x$, then the common output is x .
- **Termination:** Every honest party eventually outputs a bit.

Other distributed tasks. There are few other intimately related basic distributed functionalities. The first is *committee election* and the second is *broadcast*. In *committee election*, the goal is to bring all honest parties to agree on a small subset of parties with a fraction of corrupted parties close to the fraction for the whole set. In the extreme case, the committee is of size 1 in which case a single leader is chosen (i.e., *leader election*) and the goal is to guarantee that the leader is an honest party with constant probability. In *broadcast*, the goal is to allow an honest party to communicate a message to all other parties so that the message that all other honest parties end up knowing is the same as the one sent.

4 Preliminaries

A function $\text{negl}: \mathbb{N} \rightarrow \mathbb{R}^+$ is *negligible* if for every constant $c > 0$ there exists an integer N_c such that $\text{negl}(n) < n^{-c}$ for all $n > N_c$. The *statistical distance* between two random variables X and Y over a finite domain Ω is defined by $\text{SD}(X, Y) \triangleq \frac{1}{2} \cdot \sum_{x \in \Omega} |\Pr[X = x] - \Pr[Y = x]|$. Two sequences of random variables $X = \{X_n\}_{n \in \mathbb{N}}$ and $Y = \{Y_n\}_{n \in \mathbb{N}}$ are *statistically close* (or statistically indistinguishable) if there exists a negligible function $\text{negl}(\cdot)$ for which $\text{SD}(X_n, Y_n) \leq \text{negl}(n)$ for all $n \in \mathbb{N}$. For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$.

4.1 Secure Computation

We consider execution of distributed protocols in adversarial environments. All of our protocols assume n parties denoted P_1, \dots, P_n .

Definition 4.1 (View). *Given a protocol Π , an input \vec{x} , and party P_i , we let $\text{view}_i^\Pi(\vec{x})$ be a random variable corresponding to the information seen by P_i throughout an execution of Π . This consists of $(x_i, r_i, m_{i_1}, \dots, m_{i_k})$, where x_i is P_i 's private input, r_i is its internal coin tosses, and m_{i_j} is the j th message that was received by P_i in the protocol execution. For a subset of parties B , we denote $\text{view}_B^\Pi(\vec{x}) = \cup_{i \in B} \text{view}_i^\Pi(\vec{x})$.*

Attacker classes. A class of adversaries that is often studied is called *semi-honest*. Such adversaries follow the protocol specification but try to learn more than they should by inspecting the protocol transcript. Call this class of adversaries \mathcal{A}_{SH} .

In the next definition we consider a larger class of adversaries that we call *semi-honest with drops*. Such adversaries can deviate from the protocol specification but only in a limited way, as follows. In every round, the adversary has to perform all internal computations correctly, following the protocol specification (given everything it saw until that point), but if a message needs to be sent, the adversary can decide whether to actually send it or not (i.e., drop it).

Definition 4.2 (Semi-honest with drops adversary). *An adversary is semi-honest with drops if every message that is sent by a corrupt party is computed by following the protocols specification. I.e., the adversary has to follow the protocol specification (like a semi-honest adversary) except it can decide to drop some of the outgoing messages. Let \mathcal{A}_{SHwD} denote this class of adversaries.*

Next, we consider malicious adversaries that can follow an arbitrary strategy in order to carry out their attack.

Definition 4.3 (Malicious adversary). *An adversary is malicious if it follows an arbitrary strategy, arbitrarily deviating from the protocol. Let \mathcal{A}_{ALL} denote this class of adversaries.*

It follows immediately from the definitions above that

$$\mathcal{A}_{SH} \subseteq \mathcal{A}_{SHwD} \subseteq \mathcal{A}_{ALL}.$$

We shall only consider static corruptions in this work. Namely, the set of corrupted parties will be chosen by the adversary after the protocol is specified but before an execution begins. Also, the adversary will be allowed to control some constant fraction of the parties, typically, $1/3 - \epsilon$ for any $\epsilon > 0$. Thus, we introduce the notation $\mathcal{A}_\alpha^{static}$ for the family of adversaries that statically corrupt up to $\alpha - \epsilon$ fraction of parties for any positive constant ϵ . Lastly, we emphasize that in all of the above definitions, we do not limit the computational power of the adversary in any way.

Security definition. Next, we provide a security definition that will be used throughout this work. We follow the real-ideal paradigm that is helpful in modularizing our proofs. Concretely, this method allows us to transition to a hybrid model where some of the building blocks are implemented using ideal functionalities. Since all of our simulators will be simple enough (so called “straight-line”) we will automatically be able to apply composition theorems [Can00, Can01, Gol04].

The real-world execution: In the real-world execution, the protocol Π is carried out among n parties, where some subset \mathcal{C} of parties is controlled by the adversary \mathcal{A} . We shall also assume that the execution is performed in the presence of some auxiliary input aux known to \mathcal{A} ; this is important for composition. Then, we let real-world execution consist of the view of the adversary and the output of the honest parties. That is, for an input $\vec{x} = (x_1, \dots, x_n)$ and auxiliary information aux , let $\text{real}_{\mathcal{A}}^{\Pi}(\vec{x}, \text{aux}) = (\text{view}_{\mathcal{C}}^{\Pi}(\vec{x}), \{y_i\}_{i \in [n] \setminus \mathcal{C}})$, where y_i is P_i 's output.

The ideal-world execution: The ideal-world execution is augmented with a “trusted party” constituting a general algorithmic entity called an ideal functionality and denoted \mathcal{F} . The ideal functionality, which is modeled as another machine, repeatedly receives inputs from the parties and provides them with appropriate output values, while maintaining local state in between. To this end, we define an experiment involving an imaginary adversary \mathcal{S} (often called the simulator) that knows \mathcal{A} and can interact with \mathcal{F} . The experiment's output, given input $\vec{x} = (x_1, \dots, x_n)$ and auxiliary information aux , is denoted $\text{ideal}_{\mathcal{S}, \mathcal{A}}^{\mathcal{F}}(\vec{x}, \text{aux})$ and is defined as follows. Note that we consider only static corruptions.

1. **Initialization:** \mathcal{S} chooses $\mathcal{C} \subset [n]$, the corrupted parties. \mathcal{S} receives $\{x_i\}_{i \in \mathcal{C}}$ and aux .
2. **Evaluation:** \mathcal{F} receives \mathcal{C} and $\{x_i\}_{i \in [n] \setminus \mathcal{C}}$ and then interacts with \mathcal{S} . At the end of the interaction, \mathcal{S} knows $\{y_i\}_{i \in \mathcal{C}}$ and \mathcal{F} outputs $\{y_i\}_{i \in [n]}$.
3. **Output:** \mathcal{S} computes $\{\text{view}_i\}_{i \in \mathcal{C}}$ (based on everything that it knows) and the output of the experiment is $(\{\text{view}_i\}_{i \in \mathcal{C}}, \{y_i\}_{i \in [n] \setminus \mathcal{C}})$.

The following definition says that a protocol is considered secure with respect to a functionality and a class of adversaries if it is possible to come up with a simulator that causes the above two worlds to be indistinguishable.

Definition 4.4. *We say that a protocol Π realizes functionality \mathcal{F} against adversaries family \mathcal{A} , if for every $\mathcal{A} \in \mathcal{A}$ there exists a simulator \mathcal{S} such that for every input \vec{x} and every $\text{aux} \in \{0, 1\}^*$, it holds that $\text{real}_{\mathcal{A}}^{\Pi}(\vec{x}, \text{aux})$ is statistically indistinguishable from $\text{ideal}_{\mathcal{S}, \mathcal{A}}^{\mathcal{F}}(\vec{x}, \text{aux})$.*

If the statistical distance between $\text{real}_{\mathcal{A}}^{\Pi}(\vec{x}, \text{aux})$ and $\text{ideal}_{\mathcal{S}, \mathcal{A}}^{\mathcal{F}}(\vec{x}, \text{aux})$ is zero, we say that the protocol Π is *perfectly* secure.

4.2 (Attacker-Dependent) Complexity Measures of Protocols

Given a protocol Π , consider some abstract complexity measure for Π , denoted $\text{complexity}(\cdot)$. For concreteness, assume complexity stands for communication complexity, denoted also cc , but it could be anything else like round complexity, denoted rc , or computational complexity. Typically, $\text{complexity}(\cdot)$ of a protocol is defined as the maximal total communication (of honest parties) in Π , over all possible executions, in the presence of an arbitrary adversary. However, we want to be more fine grained.

Definition 4.5 (Complexity). *Given a protocol Π , an adversary $\mathcal{A} \in \mathcal{A}_{ALL}$, and randomness r , we define $\text{complexity}(\Pi, \mathcal{A}, i, r)$ as the number of bits send by party i throughout the execution of Π in the presence of \mathcal{A} and randomness r .*

The standard definition of the complexity of the i th party in Π is then the worst case over all possible adversaries and randomnesses. Namely,

$$\text{complexity}(\Pi, i) = \max_{r, \mathcal{A}} \{\text{complexity}(\Pi, \mathcal{A}, i, r)\}.$$

For the complexity of the protocol Π , we sum up the complexities of individual parties (worse case over all possible adversaries and randomnesses).

$$\text{complexity}(\Pi) = \max_{r, \mathcal{A}} \left\{ \sum_{i=1}^n \text{complexity}(\Pi, \mathcal{A}, i, r) \right\}.$$

Definition 4.6. [Balanced protocol] *A Protocol is said to be balanced if*

$$\frac{\text{complexity}(\Pi)}{\max_{i \in [n]} \{\text{complexity}(\Pi, i)\}} \in \tilde{\Theta}(n).$$

Adversary-dependent complexity. It is useful to adapt the above definitions with respect to a particular sub-class of attackers. For $\mathcal{A} \subseteq \mathcal{A}_{ALL}$, define

$$\text{complexity}(\Pi, \mathcal{A}, i) = \max_{r, \mathcal{A} \in \mathcal{A}} \{\text{complexity}(\Pi, \mathcal{A}, i, r)\}$$

and analogously,

$$\text{complexity}(\Pi, \mathcal{A}) = \max_{r, \mathcal{A} \in \mathcal{A}} \left\{ \sum_{i=1}^n \text{complexity}(\Pi, \mathcal{A}, i, r) \right\}.$$

In words, we restrict the possible set of actions the adversary can do, and then measure how expensive the protocol is if such an adversary is present. Note that if $\mathcal{A} = \mathcal{A}_{ALL}$, then the above definition coincides with the general (worst-case over all adversaries) definition above. That is, $\text{complexity}(\Pi, \mathcal{A}_{ALL}, i) = \text{complexity}(\Pi, i)$ for every $i \in [n]$ and $\text{complexity}(\Pi, \mathcal{A}_{ALL}) = \text{complexity}(\Pi)$.

5 Propagating the Detection of Failures

In this section we show how to propagate the knowledge that a failure occurred during the execution of a protocol from a single party to everyone. Of course this task can be easily achieved by using a generic broadcast protocol or some variant of Byzantine Agreement, but this is expensive in terms of communication and we would like to do it more efficiently. First, we explain the task at hand by defining the functionalities we want to achieve formally. Then, we will show how to implement our transformation efficiently.

5.1 Failure Detection Functionalities

Generically, the behavior of a functionality in case some of the adversarial parties deviate from the protocol's specification in an evident way (say by completely aborting), depends on the exact specification of the functionality. For us, we need to be more precise and we do so next.

We consider attackers that cheat and some honest party notices that something went wrong, and how it affects the outcome of the protocol. Specifically, we define two variants (or wrappers) of a given functionality \mathcal{F} . The first variant, denoted \mathcal{F}^{gFD} (where gFD stands for *global failure detection*), says that if some party noticed that something went wrong during an execution, then every honest party should know about this. The second variant, denoted \mathcal{F}^{lFD} (where lFD stands for *local failure detection*), says that at least one honest party should end up being aware that the protocol did not execute honestly. In this case, we have no guarantee or requirement about the output of all other parties.

It is worth pointing out that a priori it seems much easier to satisfy local failure detection. This is true because going from local failure detection to global failure detection requires some form of broadcast or agreement which are very close to the problem we are to solve in the first place.

We formalize both variants (\mathcal{F}^{gFD} first and then \mathcal{F}^{lFD}) as wrappers of a given functionality \mathcal{F} . In the above, the simulator has to identify which parties noticed that something went wrong during the execution and inform the functionality. The latter will either set the output of all honest parties to Fail or just a single party that noticed, respectively.

Functionality 1: \mathcal{F}^{gFD}

1. \mathcal{S} and \mathcal{F}^{gFD} interact according to \mathcal{F} . At the end, \mathcal{F}^{gFD} knows $\vec{y} = (y_1, \dots, y_n)$.
2. \mathcal{S} submits Fail or Success.
3. If \mathcal{S} submitted Success, then output \vec{y} .
4. If \mathcal{S} submitted Fail, output (Fail, ..., Fail).

Functionality 2: \mathcal{F}^{lFD}

1. \mathcal{S} and \mathcal{F}^{lFD} interact according to \mathcal{F} . At the end, \mathcal{F}^{lFD} knows $\vec{y} = (y_1, \dots, y_n)$.
2. \mathcal{S} submits Fail or Success.
3. If \mathcal{S} submitted Success, then output \vec{y} .
4. If \mathcal{S} submitted Fail, then \mathcal{S} specifies a vector (y'_1, \dots, y'_n) , where at least one y'_i corresponding to an honest party is equal to Fail. Output (y'_1, \dots, y'_n) .

Relation between \mathcal{F} , \mathcal{F}^{gFD} , and \mathcal{F}^{lFD} . A simulator \mathcal{S} for \mathcal{F} directly implies a simulator \mathcal{S}^{gFD} for \mathcal{F}^{gFD} . Indeed, notice that \mathcal{S}^{gFD} can merely submit **Success** in step 2 and output the output of the interaction between \mathcal{S} and \mathcal{F} . By similar logic, a simulator \mathcal{S}^{gFD} for \mathcal{F}^{gFD} directly implies a simulator \mathcal{S}^{lFD} for \mathcal{F}^{lFD} . In other words, it is easiest to realize the functionality \mathcal{F}^{lFD} , harder to realize the functionality \mathcal{F}^{gFD} , and hardest to realize the functionality \mathcal{F} .

We summarize these relations below.

Observation 5.1. *Consider a class of adversaries \mathcal{A} , a functionality \mathcal{F} , and a protocol Π . The following holds:*

- If Π realizes \mathcal{F} against \mathcal{A} , then Π also realizes \mathcal{F}^{gFD} against \mathcal{A} .
- If Π realizes \mathcal{F}^{gFD} against \mathcal{A} , then Π also realizes \mathcal{F}^{lFD} against \mathcal{A} .

5.2 Propagating Knowledge of Failures

We show a method for propagating knowledge of failures from a single party to all of them. Specifically, we translate protocols that realize *local* failure detection to ones that realize global failure detection. Recall that it is only interesting to consider the above task if we insist on preserving some other non-trivial property of the protocol. Indeed, if we only care about realizing \mathcal{F}^{gFD} for some functionality \mathcal{F} , then it is trivial: all parties will just output **Fail**. So, naturally, we will require the propagation procedure to preserve the fact that the protocol realizes \mathcal{F} . Details follow.

We start off with a protocol Π that realizes a functionality \mathcal{F} against a class of attackers \mathcal{A} . Further Π realizes \mathcal{F}^{lFD} against \mathcal{A}^* , where $\mathcal{A} \subset \mathcal{A}^*$. Thus, by Observation 5.1, it is immediate that Π realizes \mathcal{F}^{gFD} against \mathcal{A} . Our next theorem shows that it is possible to transform Π into Π' that realizes \mathcal{F}^{gFD} against \mathcal{A}^* (which contains \mathcal{A} and therefore this task is non-trivial); at the same time, we preserve the fact that Π' realizes \mathcal{F} against \mathcal{A} . The transformation costs only an additive poly-logarithmic term in n overhead in rounds and communication complexity (per party).

Theorem 5.2. *Consider a protocol Π and a functionality \mathcal{F} . Further, consider two attacker classes $\mathcal{A}^* \subseteq \mathcal{A}_{1/3}^{static}$ and $\mathcal{A} \subseteq \mathcal{A}_{SHwD} \cap \mathcal{A}^*$. Assume that Π realizes \mathcal{F} against \mathcal{A} and realizes \mathcal{F}^{lFD} against \mathcal{A}^* .*

Then, there is a protocol Π' that realizes \mathcal{F} against \mathcal{A} and realizes \mathcal{F}^{gFD} against \mathcal{A}^ . Furthermore, $\text{rc}(\Pi') = \text{rc}(\Pi) + \text{polylog}(n)$, and for every party i , $\text{cc}(\Pi', i) = \text{cc}(\Pi, i) + \text{polylog}(n)$.*

The rest of the section is devoted to the proof of Theorem 5.2. At a very high level, we obtain Theorem 5.2 by designing a new type of Byzantine Agreement abstraction. This abstraction, referred to as `BYZANTINEAGREEMENTFORFAILURES` still requires agreement among all honest parties, like standard BA. The main difference is that we require that if some honest party has the input **Fail**, then all honest parties will agree on **Fail**. If all parties hold **Success** and the adversary behaves somewhat honestly, then everyone will agree on **Success**. Using the `BYZANTINEAGREEMENTFORFAILURES` protocol, we can first run a protocol with local failure detection (as guaranteed in Theorem 5.2), and then “boost” the knowledge of failure of any party. In the resulting protocol, either everyone gets the needed output or everyone agrees that something went wrong.

In Section 5.3 we give present our new `BYZANTINEAGREEMENTFORFAILURES` sub-protocol. In Section 5.4 we present the proof of Theorem 5.2.

5.3 Byzantine Agreement for Failures

In this section we implement a weak variant of Byzantine Agreement that we refer to as `BYZANTINEAGREEMENTFORFAILURES`. Each party P_i start with an input $x_i \in \{\text{Success}, \text{Fail}\}$, and output $y_i \in \{\text{Success}, \text{Fail}\}$. The properties of our protocol are:

- **Agreement:** For all honest parties the output is same.
- **Termination:** After fixed time (in n), each party must hold an output.
- **Failure Detection:** If an honest party P_i has $x_i = \text{Fail}$, then the agreed value is Fail .
- **Fairness:** For all adversaries in \mathcal{A}_{SHwD} (see Definition 4.2), if all parties' inputs are Success , then all parties output Success .

We do not require or provide any guarantee as to what happens if (say) the adversary is outside \mathcal{A}_{SHwD} and all of the inputs (of honest parties) are Success . This relaxation allows us to implement our protocol much more efficiently than full BA.

The protocol. Our protocol relies on the existence of a good quorum. A committee is said to be good if it contains $\text{polylog}(n)$ many parties and more than $2/3$ of them are good. A good quorum is a set of n good committees where additionally each party appears in roughly the same number of committees as any other party. In the next theorem we argue that there are efficient ways to come up with a quorum. Put another way, we use a pseudorandom object that takes a sufficiently random short string as input, and outputs many random-enough good committees. Such a construction is known to exist by the probabilistic method; see [KLST11, Lemma 3] for a proof.⁴

Theorem 5.3 (Quorum generation [KLST11]). *For every integer $n \in \mathbb{N}$, $c = c(n) \in \Omega(\log^2 n)$, there exist $d = O(c \cdot \log n)$ a mapping $H: [n] \times [n^c] \rightarrow [n]^d$ such that, viewing each $H(i, x)$ as outputting a multiset of $[n]$ of size d , the following holds. For every $B \subseteq [n]$, for at least $1 - \text{negl}(n)$ fraction of the x 's, it holds that*

1. For every $i \in [n]$, $H(i, x)$ contains at most $|B|/n + o(1)$ fraction of parties from B .
2. Define $H^{-1}(j, x) = \{i \mid j \in H(i, x)\}$. Then, for every $j \in [n]$, $|H^{-1}(j, x)| \in O(d)$.

The above theorem allows us, at a high level, to turn a sufficiently high (min-)entropy source to a good quorum (with high probability). We now proceed with the description of the protocol. Below, we associate Fail with 1 and Success with 0 so that Boolean operations (such as logical OR, denoted \vee) are well defined.

The BYZANTINEAGREEMENTFORFAILURES Protocol

- **Setup Phase:** Each party P_i set $y_i \leftarrow x_i$.
- **Almost-Everywhere Agreement Phase:** All parties run together the almost-everywhere Byzantine agreement protocol of [KLST11].⁵ At the end, each party P_i holds a string str_i of size $\Theta(\log^3 n)$.
- **Agreement Phase:**
 1. Each party P_i samples a set of $\log^2 n$ parties (denoted S_i), and sends them “ping”.
Filtering: Each party accepts at most $2 \cdot \log^2 n$ “ping”s.

⁴The statement in [KLST11, Lemma 3] slightly differs from ours. Specifically, they assume that c is constant and hide c -dependent terms in the size of $d = O(\cdot)$. Following their proof, it is easy to support also super constant c 's and explicitly write the multiplicative dependence on c in d .

⁵The original almost-everywhere protocol provided by King et al. [KSSV06] is described as a leader election protocol. However, a simple modification gives an almost-everywhere agreement protocol where the output of all but $o(1)$ -fraction of parties is a poly-logarithmically long string with poly-logarithmic min-entropy.

2. For each party P_j : if it received less than $2 \cdot \log^2 n$ “ping”s, P_j answers with str_j ; else, P_j sets $y_j \leftarrow \text{Fail}$.
Filtering: P_i accepts messages with size at most $\log^4 n$, only from parties in S_i .
3. For each party P_i : if there is a string str^* that was received more than $|S_i|/2$ times, P_i sets $\text{str}_i \leftarrow \text{str}^*$; else, P_i sets $y_i \leftarrow \text{Fail}$.

- **Distribution Phase:** This phase uses the quorum generation functions H, H^{-1} from Theorem 5.3.

1. For each party P_i , if $y_i = \text{Fail}$, P_i sends “ping” to $H(i, \text{str}_i)$.
Filtering: P_j accepts messages of $|\text{“ping”}|$ bit size, only from parties in $H^{-1}(j, \text{str}_j)$.
2. For each committee Q_i , if “ping” is received, then each party in the committee set $y_i \leftarrow \text{Fail}$, else $y_i \leftarrow \text{Success}$.
3. The committees compute the logical OR of the y_i values: all committees communicate their values upstream in a (depth $O(\log n)$) tree-like manner, calculating a logical OR along the way, and then broadcast the results in reverse order. (The exact pseudocode is given in Appendix A as the DISTRIBUTE protocol.)
4. Each party P_j sends x_i to each party in $H^{-1}(j, \text{str}_j)$.
Filtering: P_i accepts messages of 1 bit size, only from parties in $H(i, \text{str}_i)$.

- **Output Phase:** If there is x^* that was received more than $|H(i, \text{str}_i)|/2$ times, then P_i sets $x \leftarrow x^*$; else, P_i sets $x \leftarrow \text{Success}$. After that, P_i outputs $y_i \vee x$.

Lemma 5.4. *Consider an adversary $\mathcal{A} \in \mathcal{A}_{1/3}^{\text{static}}$ corrupting a set of parties \mathcal{C} . At the end of the BYZANTINEAGREEMENTFORFAILURES protocol, with probability $1 - \text{negl}(n)$, it holds that:*

1. **Agreement:** *All honest parties output the same value.*
2. **Failure Detection:** *If $\bigvee_{j \in [n] \setminus \mathcal{C}} x_j = \text{Fail}$, then the output is Fail.*
3. **Fairness:** *If $\mathcal{A} \in \mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{\text{static}}$ and $\bigwedge_{j \in [n]} x_j = \text{Success}$, then the output is Success.*

Efficiency. We analyze the round and communication complexity of each party in the above protocol and argue that both of them are bounded by some $\text{polylog}(n)$. First, the almost-everywhere agreement phase requires $\text{polylog}(n)$ rounds and each party sends and receives $\text{polylog}(n)$ bits. Now, we analyze the remaining phases separately. The agreement phase consists of two communication rounds. In the first, each party sends $O(\log^2 n)$ bits and receives at most $O(\log^2 n)$ bits. In the second, each party sends $O(\log^6 n)$ bits and receives at most $\log^6 n$ bits. The distribution phase consists of 3 communication phases. The first requires each party to send $O(\log^5 n)$ bits and receive at most $O(\log^5 n)$ bits. The “Distribute” sub-protocol requires $O(\log n)$ rounds, and each party sends and receives $O(\log^{10} n)$ bits within each round. In the last round, each party sends $O(\log^5 n)$ bits and receives at most $\log^5 n$ bits. Overall, the total number of rounds is $\text{polylog}(n)$ and the total communication is also $\text{polylog}(n)$, per party, as needed.

Correctness. Next, we prove that the protocol satisfies all the claimed properties (Lemma 5.4). First, we prove that after the agreement phase every honest party gained some knowledge in the following sense: either it knows the “correct” string that defines the quorum (denoted str), or it knows that something went wrong.

Claim 5.5. *Assume that for at least $1 - o(1)$ fraction of honest parties $\text{str}_i = \text{str}$. For every adversary $\mathcal{A} \in \mathcal{A}_{1/3}^{\text{static}}$, at the end of the Agreement Phase, with probability $1 - \text{negl}(n)$, for all honest parties $\{P_i\}_{i \in [n] \setminus \mathcal{C}}$, either $\text{str}_i = \text{str}$ or $y_i = \text{Fail}$.*

Proof. By Step 3 of the protocol, the conclusion could be wrong only if more than half of the parties in some S_i hold $\text{str}^* \neq \text{str}$. Recall that each S_i is chosen uniformly at random. Let $\mu = (2/3 + \epsilon) \cdot \log^2 n$, and let $\delta = 1/\sqrt{(2/3 + \epsilon) \cdot \log \log n}$. Thus, for every $i \in [n] \setminus \mathcal{C}$, by a Chernoff bound, with probability $1 - e^{-\delta^2 \mu/2} = 1 - e^{-\log^2 n/2 \log \log n} = 1 - \text{negl}(n)$, the fraction of honest parties in S_i is at least $2/3 + \epsilon - o(1)$. So, by the assumption, at least $(1 - o(1)) \cdot (2/3 + \epsilon - o(1)) > 1/2$ fraction of honest parties in S_i holds str . By a union bound, the above holds for all i 's simultaneously with probability $1 - \text{negl}(n)$. \square

Next, we prove that if there is an honest party for which $\text{str}_i = \text{str}$ but $y_i = \text{Fail}$, then all honest parties that know str will necessarily also output Fail .

Claim 5.6. *Assume that for at least $1 - o(1)$ fraction of honest parties $\text{str}_i = \text{str}$ and that the min-entropy of str is $\Omega(\log^2 n)$. For every adversary $\mathcal{A} \in \mathcal{A}_{1/3}^{\text{static}}$, if at the end of the Agreement Phase there is an honest party P_i with $\text{str}_i = \text{str}$ and $y_i = \text{Fail}$, then at the end of BYZANTINEAGREEMENTFORFAILURES, with probability $1 - \text{negl}(n)$, every honest party with $\text{str}_i = \text{str}$ will output $y_i = \text{Fail}$.*

Proof. By Theorem 5.3, with probability $1 - \text{negl}(n)$ over str , each committee in the resulting quorum has at least $2/3$ fraction of honest parties. This means that with probability $1 - \text{negl}(n)$, the fraction of honest parties in each committee with $\text{str}_i = \text{str}$ is $2/3 - o(1) > 1/2$. Therefore, the prerequisite for the DISTRIBUTE sub-protocol is satisfied and it results with every honest party setting $y_i = \text{Fail}$. \square

We remain with the case where after the Agreement Phase only parties without knowledge of str have $y = \text{Fail}$. However, as we show next, such a scenario is unlikely to occur.

Claim 5.7. *Assume that for at least $1 - o(1)$ fraction of honest parties $\text{str}_i = \text{str}$. For every adversary $\mathcal{A} \in \mathcal{A}_{1/3}^{\text{static}}$, at the end of the Agreement Phase, if there is an honest party P_i with $y_i = \text{Fail}$, then with probability $1 - \text{negl}(n)$, there is an honest party P_j with $\text{str}_j = \text{str}$ and $y_j = \text{Fail}$.*

Proof. If for P_i it holds that $\text{str}_i = \text{str}$, we are done (as we can use $j = i$). So, assume that $\text{str}_i \neq \text{str}$ at the end of the Agreement Phase. As we argued in the proof of Claim 5.5, with probability $1 - \text{negl}(n)$, this scenario can occur only if there is at least one honest party $P_j \in S_i$ with $\text{str}_j = \text{str}$ that does not send a message to P_i in Step 2. From the way the protocol works, in this case, P_j sets $y_j \leftarrow \text{Fail}$, as needed. \square

Now, we show that an adversary that is semi-honest with drops cannot force an honest party to not learn str .

Claim 5.8. *Assume that for at least $1 - o(1)$ fraction of honest parties $\text{str}_i = \text{str}$. For every adversary $\mathcal{A} \in \mathcal{A}_{\text{SHwD}} \cap \mathcal{A}_{1/3}^{\text{static}}$, at the end of the Agreement Phase, with probability $1 - \text{negl}(n)$, each honest party P_i holds $\text{str}_i = \text{str}$ and $y_i = x_i$.*

Proof. As in the proof of Claim 5.7, for an honest party P_i , $\text{str}_i \neq \text{str}$ only if there is at least one honest party $P_j \in S_i$ that did not answer to P_i at Step 2. From the protocol description, the latter is possible only if P_j receives $2 \cdot \log^2 n$ ‘‘pings’’. However, \mathcal{A} is semi-honest with drops meaning that it has to follow the protocol except it can decide to drop some of its outgoing messages. So, by a

Chernoff bound, too many “pings” arrive to P_j only with probability $1 - e^{-(2/3+\epsilon)\cdot\log^2 n} = 1 - \text{negl}(n)$, and from a union bound, too many “pings” arrive to *some* honest party in S_i with probability $1 - \text{negl}(n)$. Finally, in Step 2 and Step 3 each honest party will never set $y \leftarrow \text{Fail}$. \square

Finally, we are ready to combine all the above claims to prove Lemma 5.4.

Proof of Lemma 5.4. From the correctness of the almost-everywhere agreement protocol, after the Almost-Everywhere Agreement Phase at least $1 - o(1)$ fraction of honest parties holds $\text{str}_i = \text{str}$, where the min-entropy of str is $\Theta(\log^2 n)$. Let K be the set of honest parties that do not know str after the Agreement Phase, that is, $K \equiv \{P_i \mid P_i \in [n] \setminus \mathcal{C} \wedge \text{str}_i \neq \text{str}\}$ at the end of the Agreement Phase.

• If $K = \emptyset$:

1. **Agreement:** By Claim 5.6, w.h.p ($1 - \text{negl}(n)$) for each honest party P_i , $y_i \leftarrow \bigvee_{j \in [n] \setminus \mathcal{C}} y_j$.
2. **Failure Detection:** From the protocol description, if $x_i = \text{Fail}$, then at every step $y_i = \text{Fail}$. So, if $\bigvee_{j \in [n] \setminus \mathcal{C}} x_j = \text{Fail}$, then also $\bigvee_{j \in [n] \setminus \mathcal{C}} y_j = \text{Fail}$.
3. **Fairness:** If $\mathcal{A} \in \mathcal{A}_{SHwD}$ and $\bigwedge_{j \in [n]} x_j = \text{Success}$, then by Claim 5.8 w.h.p at the start of the Distribution Phase, for each party P_i , $y_i = \text{Success}$, and therefore nobody will send “ping” at Step 1. So, at the end, each honest party will output Success.

• If $K \neq \emptyset$:

1. **Agreement:** By Claim 5.5, w.h.p it follows that each party in K will output Fail. By Claim 5.7, w.h.p it follows that there is a party $P_j \in [n] \setminus (\mathcal{C} \cup K)$ with $y_j = \text{Fail}$, so from Claim 5.6 w.h.p it follows that each party in $[n] \setminus (\mathcal{C} \cup K)$ will output Fail.
2. **Failure Detection:** The output of all parties is Fail by correctness of the Distribution sub-protocol.
3. **Fairness:** By Claim 5.8, w.h.p for each honest party P_i , $\text{str}_i = \text{str}$ which mean that $P_i \notin K$, and so $K = \emptyset$. This is a contradiction to $K \neq \emptyset$, which can only happen if $\mathcal{A} \notin \mathcal{A}_{SHwD}$ which is again a contradiction. \square

5.4 Proof of Theorem 5.2

In this section we prove Theorem 5.2, boosting protocols with local failure detection to ones with global failure detection.

The protocol Π'

Assumption: Π is a protocol that realizes \mathcal{F} against \mathcal{A} and realizes \mathcal{F}^{lFD} against \mathcal{A}^* .

The protocol:

1. All parties participate in an execution of Π . At the end, each party P_i holds an output y_i .
2. For each party P_i , if $y_i = \text{Fail}$, then P_i sets $x_i^* \leftarrow \text{Fail}$; else $x_i^* \leftarrow \text{Success}$.
3. All parties run the `BYZANTINEAGREEMENTFORFAILURES` protocol from Section 5.3 with $\{x_i^*\}_{i \in [n]}$. The outputs are denoted $\{y_i^*\}_i$.
4. For each party P_i , if $y_i^* = \text{Fail}$, then P_i outputs Fail; else P_i outputs y_i .

In what follows we show that Π' realizes \mathcal{F} against \mathcal{A} and realizes \mathcal{F}^{gFD} against \mathcal{A}^* . For this we present two simulators.

5.4.1 Realizing \mathcal{F}^{gFD} Against \mathcal{A}^*

Since Π realizes \mathcal{F}^{lFD} against any $A \in \mathcal{A}^*$, then there is a simulator \mathcal{S}^{lFD} that communicate with \mathcal{F}^{lFD} in the ideal experiment and causes the ideal and real worlds to be indistinguishable; see Functionality 2 for the local failure detection functionality and Section 4.1 for the ideal world experiment. We will design a simulator for the global failure detection functionality as described in Functionality 1. The simulator is described next:

1. **Initialization:** \mathcal{S}^{gFD} is initialized the same as \mathcal{S}^{lFD} .
2. **Evaluation:** \mathcal{S}^{gFD} executes \mathcal{S}^{lFD} . In the case that \mathcal{S}^{lFD} sends Fail in Step 2 of Functionality 2, then in Step 4, \mathcal{S}^{lFD} provides a vector (y'_1, \dots, y'_n) . \mathcal{S}^{gFD} does not need to provide anything in Step 4 of Functionality 1, so it saves the vector for internal use.
3. **Output:**
 - (a) \mathcal{S}^{gFD} computes $\{\text{view}_i^{lFD}\}_{i \in \mathcal{C}}$ exactly like \mathcal{S}^{lFD} .
 - (b) For each honest party P_i , if $y'_i = \text{Fail}$, then \mathcal{S}^{gFD} set $x_i^* \leftarrow \text{Fail}$, else $x_i^* \leftarrow \text{Success}$.
 - (c) \mathcal{S}^{gFD} and \mathcal{A} run the `BYZANTINEAGREEMENTFORFAILURES` protocol with $\{x_i^*\}_{i \in [n]}$, where \mathcal{S}^{gFD} simulates the honest parties by running honestly. At the end, \mathcal{S}^{gFD} holds $\{\text{view}'_i\}_{i \in [n]}$, the view of the execution.
 - (d) \mathcal{S}^{gFD} outputs $\{\text{view}_i\}_{i \in \mathcal{C}} = \{\text{view}_i^{lFD} \parallel \text{view}'_i\}_{i \in \mathcal{C}}$

Claim 5.9. *Let Sim^{lFD} be a simulator for Π and the functionality \mathcal{F}^{lFD} . Then, \mathcal{S}^{gFD} , as described above, is a simulator for Π' and the functionality \mathcal{F}^{gFD} .*

Proof. Let $\{\text{view}_i^{\Pi'}\}_{i \in \mathcal{C}} = \{\text{view}_i^{\Pi'_1} \parallel \text{view}_i^{\Pi'_{2-4}}\}_{i \in \mathcal{C}}$ be the adversary view in the real world execution, split to the stages in Π' . Since Π realizes \mathcal{F}^{lFD} against every $\mathcal{A} \in \mathcal{A}^*$, then $(\{\text{view}_i^{\Pi'_1}\}_{i \in \mathcal{C}}, \{y'_i\}_{i \in [n] \setminus \mathcal{C}})$ is statistical indistinguishable from $(\{\text{view}_i^{lFD}\}_{i \in \mathcal{C}}, \{y_i\}_{i \in [n] \setminus \mathcal{C}})$. Since \mathcal{S}^{gFD} knows the inputs of the `BYZANTINEAGREEMENTFORFAILURES` protocol and act honestly, $\{\text{view}_i^{\Pi'_{2-4}}\}_{i \in \mathcal{C}}$ and $\{\text{view}'_i\}_{i \in \mathcal{C}}$ are identically distributed. Also, the outputs of the `BYZANTINEAGREEMENTFORFAILURES` protocol are statistical indistinguishable. To see this, we consider two cases:

- **There is an honest party P_i with $y_i^* = \text{Fail}$:**
 - **Real world:** By bullet 1 of Lemma 5.4, we know that every honest party will output Fail, and so $\{y_i\}_{i \in \mathcal{H}} = \{\text{Fail}, \dots, \text{Fail}\}$.
 - **Ideal world:** If there is an honest party P_i with $y_i^* = \text{Fail}$, then \mathcal{S}^{lFD} sends Fail at step 2 of Functionality 2. Thus, \mathcal{S}^{gFD} also sends Fail at step 2 of Functionality 1, and so $\{y_i\}_{i \in \mathcal{H}} = \{\text{Fail}, \dots, \text{Fail}\}$.
- **Otherwise (no honest party P_i has $y_i^* = \text{Fail}$):**
 - **Real world:** By bullets 1 and 2 of Lemma 5.4 this case can occur (w.h.p) only if for every honest party P_i after the first stage of Π' , $y_i \neq \text{Fail}$. But, since Π realizes \mathcal{F}^{lFD} , then $\vec{y} \in \mathcal{F}(\mathcal{C}, \{\tilde{x}_i\}_{i \in [n]})$.

- **Ideal world:** If no honest party P_i has $y_i^* = \text{Fail}$, then \mathcal{S}^{lFD} sends **Success** at step 2 of Functionality 2. This means that \mathcal{S}^{gFD} also sends **Success** at step 2 of Functionality 1, and so $\vec{y} \in \mathcal{F}(\mathcal{C}, \{\tilde{x}_i\}_{i \in [n]})$.

Overall, the tuple $(\{\text{view}_i\}_{i \in \mathcal{C}}, \{y_i\}_{i \in \mathcal{H}})$ is statistical indistinguishable between the real and ideal worlds. \square

5.4.2 Realizing \mathcal{F} Against \mathcal{A}

Since Π realizes \mathcal{F} against an adversary $\mathcal{A} \in \mathcal{A}$, then there is a simulator \mathcal{S}_Π that communicates with \mathcal{F} in the ideal experiment and is able to generate the right view. The behavior of \mathcal{S} in the ideal experiment (see Section 4.1) with \mathcal{F} is described next:

1. **Initialization:** \mathcal{S} simulates \mathcal{S}_Π .
2. **Evaluation:** \mathcal{S} simulates \mathcal{S}_Π .
3. **Output:**
 - (a) \mathcal{S} computes $\{\text{view}_i^\Pi\}_{i \in \mathcal{C}}$ by simulating \mathcal{S}_Π .
 - (b) For each honest party P_i , \mathcal{S} sets $x_i^* \leftarrow \text{Success}$.
 - (c) \mathcal{S} and \mathcal{A} run the `BYZANTINEAGREEMENTFORFAILURES` protocol with $\{x_i^*\}_{i \in [n]}$, where \mathcal{S} simulates the honest parties and acts honestly. At the end, \mathcal{S} holds $\{\text{view}'_i\}_{i \in [n]}$, the view of the execution.
 - (d) \mathcal{S} outputs $\{\text{view}_i\}_{i \in \mathcal{C}} = \{\text{view}_i^\Pi \parallel \text{view}'_i\}_{i \in \mathcal{C}}$.

Claim 5.10. *Let Sim_Π be a simulator for Π and the functionality \mathcal{F} . Then, \mathcal{S} , as described above, is a simulator for Π' and the functionality \mathcal{F} .*

Proof. Let $\{\text{view}_i^{\Pi'}\}_{i \in \mathcal{C}} = \{\text{view}_i^{\Pi'_1} \parallel \text{view}_i^{\Pi'_2-4}\}_{i \in \mathcal{C}}$ be the adversary view in the real world execution, split to the stages in Π' . Since Π realizes \mathcal{F} against \mathcal{A} , then $(\{\text{view}_i^{\Pi'_1}\}_{i \in \mathcal{C}}, \{y_i^{\text{real}}\}_{i \in [n] \setminus \mathcal{C}})$ is statistical indistinguishable from $(\{\text{view}_i^{lFD}\}_{i \in \mathcal{C}}, \{y_i^{\text{ideal}}\}_{i \in [n] \setminus \mathcal{C}})$, and for every honest party P_i , the value of y_i^{real} after the stage 1 is not **Fail**. Since \mathcal{S} knows the inputs of the `BYZANTINEAGREEMENTFORFAILURES` protocol and acts honestly, $\{\text{view}_i^{\Pi'_2-4}\}_{i \in \mathcal{C}}$ and $\{\text{view}'_i\}_{i \in \mathcal{C}}$ belong to the same distribution. This means that $\{\text{view}_i^{\Pi'}\}_{i \in \mathcal{C}}$ is statistically indistinguishable from $\{\text{view}_i\}_{i \in \mathcal{C}}$. From bullets 1 and 3 in Lemma 5.4 we obtain that $\{y_i^{\text{real}}\}_{i \in \mathcal{H}} \leftarrow \mathcal{F}(\mathcal{C}, \{\tilde{x}_i\}_{i \in [n]})$, which means that the tuple $(\{\text{view}_i\}_{i \in \mathcal{C}}, \{y_i\}_{i \in \mathcal{H}})$ is statistical indistinguishable between the real and ideal worlds. \square

6 Optimistic/Pessimistic Combiner Framework

In this section, we spell out the optimistic/pessimistic framework to obtain a full-fledged protocol implementing some given functionality. At a very high level, the starting point is an optimistic protocol for the functionality, where the protocol roughly assumes that parties behave honestly but has the property that it can recognize if some party deviated from the honest execution. We transform this protocol into a full-fledged protocol that is secure against all adversaries, even ones that do not follow the protocol specification.

In more detail, we first transform the given protocol into a protocol for the same functionality but with global failure detection (using Theorem 5.2). That is, the resulting protocol has the property

that if some party notices that another party deviated from the protocol than all honest parties will know of this. Now, we execute the protocol with global failure detection; by our guarantee, if a failure was noticed by any party, all honest parties know of this, and can then fallback to some less efficient but always secure protocol.

Theorem 6.1. *Consider a functionality \mathcal{F} and two attacker classes $\mathcal{A}^* \subseteq \mathcal{A}_{1/3}^{static}$ and $\mathcal{A} \subseteq \mathcal{A}_{SHwD} \cap \mathcal{A}^*$. Assume that*

- *protocol Π_{light} realizes \mathcal{F} against \mathcal{A} and \mathcal{F}^{lFD} against \mathcal{A}^* , and*
- *protocol Π_{heavy} realizes \mathcal{F} against \mathcal{A}^* .*

Then, there exists a protocol Π that realizes \mathcal{F} against \mathcal{A}^ with the following efficiency properties:*

1. *Worst-case communication complexity: For every party $i \in [n]$:*

$$\text{cc}(\Pi, i) = \text{cc}(\Pi_{light}, i) + \text{cc}(\Pi_{heavy}, i) + \text{polylog}(n).$$

2. *Adversary-dependent communication complexity: For every party $i \in [n]$*

$$\text{cc}(\Pi, \mathcal{A}, i) = \text{cc}(\Pi_{light}, i) + \text{polylog}(n).$$

3. *Worst-case round complexity: $\text{rc}(\Pi) = \text{rc}(\Pi_{light}) + \text{rc}(\Pi_{heavy}) + \text{polylog}(n)$.*

4. *Adversary-dependent round complexity: $\text{rc}(\Pi, \mathcal{A}) = \text{rc}(\Pi_{light}) + \text{polylog}(n)$.*

Proof. First, use Theorem 5.2 to compile Π_{light} (that only realizes \mathcal{F}^{lFD} against \mathcal{A}^*) into Π'_{light} (that realizes \mathcal{F}^{gFD} against \mathcal{A}^*). Now, we describe Π .

Protocol Π :

1. Run Π'_{light} .
2. For each (honest) party P_i , if $y_i = \text{Fail}$, then participant in Π_{heavy} ; Otherwise, output the output of Π'_{light} and terminate.

Efficiency. In the worst case, $\text{cc}(\Pi, i) = \text{cc}(\Pi'_{light}, i) + \text{cc}(\Pi_{heavy}, i)$ and $\text{rc}(\Pi, \mathcal{A}) = \text{rc}(\Pi'_{light}, \mathcal{A}) + \text{rc}(\Pi_{heavy}, \mathcal{A})$. Now, by Theorem 5.2, $\text{cc}(\Pi'_{light}, i) = \text{cc}(\Pi_{light}, i) + \text{polylog}(n)$ and $\text{rc}(\Pi'_{light}, \mathcal{A}) = \text{rc}(\Pi_{light}) + \text{polylog}(n)$. This proves bullets 1 and 3. If the adversary is in \mathcal{A} (meaning it does not deviate from the protocol except by possibly dropping messages), then by the guarantee of Π'_{light} , there will not be an honest party P_i with $y_i = \text{Fail}$, and so all honest parties will not participate in Π_{heavy} . Thus, in this case, the complexity of the protocol is: $\text{cc}(\Pi, \mathcal{A}, i) = \text{cc}(\Pi'_{light}, i) = \text{cc}(\Pi_{light}, i) + \text{polylog}(n)$ and $\text{rc}(\Pi, \mathcal{A}) = \text{rc}(\Pi'_{light}) = \text{rc}(\Pi_{light}) + \text{polylog}(n)$. This proves bullets 2 and 4.

Security. Since Π'_{light} realizes \mathcal{F}^{gFD} against \mathcal{A}^* , then by definition if an honest party P_i has $y_i = \text{Fail}$, then we are guaranteed that all honest parties have $y = \text{Fail}$. Thus, all honest parties will participate in Π_{heavy} , and since Π_{heavy} is secure against \mathcal{A}^* , then Π is also secure against \mathcal{A}^* . Formally, we define a simulator for Π and functionality \mathcal{F} as follows: First run the simulator of Π'_{light} and then if all honest parties have $y = \text{Fail}$, then run the simulator of Π_{heavy} . The proof that this simulator generated an indistinguishable view is straightforward given the guarantees of the underlying two simulators. \square

7 Applications

In this section we present several instantiations of our framework. First, in Section 7.1, we show how to build a quorum, i.e., a well-formed representative committee for each party in the protocol. This is a useful abstraction on its own because (as we show in Section 7.2) it can be used to directly obtain several other applications, essentially without added complexity. For example, we show how to use it to directly obtain a Byzantine Agreement protocol, a broadcast protocol, and a committee election protocol. All of these have super efficient “optimistic” complexity and the best known worst-case overhead as fallback. As another (less immediate) application, in Section 7.3, we build a scalable secure multiparty computation (SMPC) protocol with almost no overhead in an optimistic execution beyond the circuit size.

7.1 A Quorum

We start by defining the properties of a good quorum. We implement the construction of a quorum via a distributed protocol and so it is convenient to specify the properties of a quorum via a functionality. At a very high level, a quorum is a collection of n poly-logarithmic size subsets where (1) each subset contains roughly the same fraction of bad parties as in the total population, (2) Each party is a member in poly-logarithmically-many committees. We parametrize a quorum functionality with $\epsilon > 0$, where we the adversary controls constant $< \epsilon$ fraction of parties in the total population.

Functionality 3: $\mathcal{F}_{\epsilon\text{-quorum}}$ - Quorum Building

1. \mathcal{S} specifies a function $f: [n] \rightarrow [n]^d$, where $d \in \tilde{O}(1)$, such that:
 - (a) For all $i \in [n]$, $\frac{|f(i) \cap \mathcal{C}|}{|f(i)|} < \epsilon$.
 - (b) Let $f^{-1}(j) = \{i \mid j \in f(i)\}$. Then $|f^{-1}(j)| \in O(d)$ for all $j \in [n]$.
2. \mathcal{F} outputs f to every party P_i .

“Heavyweight” quorum building. A protocol for generating a $1/3$ -quorum in the presence of attackers that statically corrupts $1/3 - \epsilon$ fraction of parties for a constant $\epsilon > 0$ is known [KLST11]. The protocol consists of $\tilde{O}(1)$ rounds and $\tilde{\Theta}(\sqrt{n})$ communication per party.

Lemma 7.1 (Heavyweight quorum generation [KLST11]). *There is a protocol Π_{heavy} realizing $\mathcal{F}_{1/3\text{-quorum}}$ against $\mathcal{A}_{1/3}^{static}$. The protocol consists of $\tilde{O}(1)$ rounds and $\tilde{\Theta}(\sqrt{n})$ communication per party.*

“Lightweight” quorum building. For the lightweight protocol, we observe that it was actually obtained in `BYZANTINEAGREEMENTFORFAILURES` (Section 5.3). Specifically, we run this protocol up until the end of the Agreement Phase, where the inputs of all parties are `Success`, and the output of P_i is `Fail` if $y_i = \text{Fail}$. If not, the output is $f_{\text{str}_i}(\star) \triangleq H(\star, \text{str}_i)$, where H is the function from Theorem 5.3. See formal description in Appendix B. The proof that we obtain a realization of the $1/3$ -quorum functionality readily follows from Claims 5.5 and 5.8. We thus have the following lemma whose proof can be found in Appendix B.

Lemma 7.2 (Lightweight quorum generation). *There is a protocol Π_{light} (described above) that realizes $\mathcal{F}_{1/3\text{-quorum}}$ against $\mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{static}$ and realizes $\mathcal{F}_{1/3\text{-quorum}}^{lFD}$ against $\mathcal{A}_{1/3}^{static}$.*

We apply our optimistic-pessimistic framework from Theorem 6.1 using Lemmas 7.1 and 7.2. Thus, We get the following result.

Theorem 7.3 (Quorum building). *There is a protocol Π that realizes $\mathcal{F}_{1/3\text{-quorum}}$ against adversaries family $\mathcal{A}_{1/3}^{\text{static}}$. Furthermore, Π has the following efficiency properties:*

1. $\text{rc}(\Pi) = \tilde{O}(1)$. (Round complexity)
2. For each party P_i , $\text{cc}(\Pi, \mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{\text{static}}, i) \in \tilde{O}(1)$. (Optimistic complexity)
3. For each party P_i , $\text{cc}(\Pi, i) \in \tilde{O}(\sqrt{n})$. (Pessimistic complexity)

7.2 Byzantine Agreement, Broadcast, and Committee Election

Once we have a quorum (as in Section 7.1), it is relatively easy to obtain other useful abstractions with essentially no extra complexity. Specifically, we obtain Byzantine Agreement, Broadcast, and Committee Election (formal definitions in Section 3.2), and more with essentially optimal optimistic complexity and otherwise the best-known worst-case complexity.

Theorem 7.4. *There are $\tilde{O}(1)$ -round Byzantine agreement, broadcast, quorum/committee/leader election protocols, each with negligible (in n) probability of failure, and with the following complexity:*

- For each party, if the adversary is in $\mathcal{A}_{1/3}^{\text{static}}$, the communication complexity is $\tilde{O}(\sqrt{n})$.
- For each party, if the adversary is in $\mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{\text{static}}$, the communication complexity is $\tilde{O}(1)$.

It is rather straightforward to obtain the above theorem given the quorum functionality realization in Section 7.1. We sketch the details next. Consider the Byzantine Agreement problem in which every party has an input bit and the goal is to reach an agreement on an input which is held by one of the honest party. We first generate a quorum using Theorem 7.3. Then, each party sends its bit to its associated committee members. The committees now perform agreement in a tree-like fashion by say computing the majority value of their input bits (recall that the committees can be viewed as honestly-behaving parties so we need not worry about Byzantine behavior). Finally, all committees distribute the agreed upon bit to their respective associated parties. Broadcast is implemented in a similar fashion by performing broadcast in a tree-like fashion between the committees defined via the quorum (again, since these committees are behaving honestly, this is easy). Committee election is a special case of quorum (the first committee in the quorum is a good committee). Leader election can be implemented by running a naive (heavy-weight) protocol among a chosen committee.

7.3 Secure Multiparty Computation

In MPC there are n parties, each having an input x_i . The goal is to jointly evaluate a known functionality described by a circuit C without revealing unnecessary information about the private inputs. This can be captured within our general real-ideal simulation given in Section 4.1 using the following functionality.

Functionality 4: \mathcal{F}^C - MPC for a circuit C

1. Every party P_i submits \tilde{x}_i , where for honest parties $\tilde{x}_i = x_i$.
2. Compute $(y_1, \dots, y_n) \leftarrow C(\tilde{x}_1, \dots, \tilde{x}_n)$.
3. For every party P_i , output y_i .

The size of a circuit C is defined as the total number of wires in C , and is denoted $|C|$. We further denote the depth of C by $Depth(C)$ and define it as the length of the longest path from an input wire to an output wire. We shall assume that each gate in C has a constant number of input and output wires; say two inputs and two outputs.

Recall that our communication model assumes only point-to-point synchronous communication. The communication is authenticated and private. (This is in contrast to the rest of the paper where communication is only authenticated but not private.) We obtain an MPC protocol that has worst case complexity $\tilde{O}(|C|/n + \sqrt{n})$ per party when computing C . If the adversary is honest, then the complexity per party is $\tilde{O}(|C|/n)$. In either case, the round complexity is $\tilde{O}(Depth(C))$. Notice that if we are considering a linear size circuit, then the complexity of an honest execution is essentially the best possible. The protocol description follows [DKM⁺17] except that we use our scalable optimistic/pessimistic protocol for quorum generation.

Theorem 7.5 (Optimistic/Pessimistic Scalable MPC). *There is an MPC protocol Π such that for every circuit C , it has the following complexity:*

- $rc(\Pi) \in \tilde{O}(Depth(C))$. (Round complexity)
- For each party P_i , $cc(\Pi_{MPC}^C, \mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{static}, i) \in \tilde{O}(|C|/n)$. (Optimistic complexity)
- For each party P_i , $cc(\Pi_{MPC}^C, i) \in \tilde{O}(|C|/n + \sqrt{n})$. (Pessimistic complexity)

The BGW protocol and secret sharing. Our protocol will rely on the generic MPC protocol of Ben-Or et al. [BGW88]. This protocol gives a *perfectly* secure general purpose protocol. Security holds for a static malicious adversary corrupting $t < n/3$ of the parties, where n is the number of parties, and even with concurrent composition. The protocol supports all computations and assumes they are given as circuits. Given a circuit of depth d and size s (where size is measured as the number of wires), the protocol requires $O(d)$ rounds and each party sends $s \cdot \text{poly}(n)$ bits overall.

Looking forward, we will use this MPC protocol only among (poly-)logarithmic size groups of parties and only for computing (poly-)logarithmic size circuits. Thus, the communication and round overhead per invocation of these protocols will be $\tilde{O}(1)$.

Theorem 7.6 (Perfectly secure heavyweight MPC e.g., [AL17]). *Let f be an n -ary functionality and C its arithmetic circuit representation. There exists an explicit protocol Π_{Heavy}^{MPC} realizing \mathcal{F}^C with perfect universally composable security against static adversaries that corrupt $< n/3$ parties. Furthermore, $rc(\Pi_{Heavy}^{MPC}) \in O(Depth(C))$ and $cc(\Pi_{Heavy}^{MPC}) \in \text{poly}(|C|, n)$.*

Furthermore, we will rely on a robust secret sharing scheme, a technique for splitting a secret s among n parties in such a way that every coalition of t parties can reconstruct the secret and every coalition of $t - 1$ parties learn nothing about the secret. A well-known method achieving the above was suggested by Shamir [Sha79]. We will need a strong notion of correctness called *robustness*: large enough sets of parties that try to reconstruct the secret will succeed even if some of the parties try to cheat in the reconstruction phase. Concretely, we will need that the n parties always succeed in reconstructing the secret, even if there are $t - 1$ parties that try to cheat. This property is known to be achievable for Shamir's scheme if $t < n/3$. Note that the sharing and reconstruction procedures are efficient, namely, there is a polynomial time algorithm (in n) that implements them.

The MPC Protocol. We denote the circuit corresponding to $n/3$ -out-of- n Shamir's secret sharing scheme [Sha79] by $\text{share}_n(x, r)$. For robust reconstruction, we use the (efficient, polynomial time) procedure of [WB86]. When we say that committee is associated with outputs of the share circuit,

we mean that the i th share is associated with the i th member of the committee in lexicographic order.

Secure Multiparty Computation Protocol

- **Setup Phase (Quorum Building):**

1. All parties run the quorum building protocol from Theorem 7.3. At the end, each party holds a function f as described in Functionality 3. Denote the i th committee ($f(i)$) by Q_i , and let q_i be the size of each committee.
2. Associate each committee Q_i with $|C|/n$ wires.

- **Share Phase:** Each party P_i and committee Q_i perform a “heavy” MPC protocol (Theorem 7.6) on $\text{share}_{q_i}(x_i, r_i)$, where r_i is fresh randomness.

- **Evaluation Phase:**

1. For layer $\ell \in 1 \dots \text{Depth}(C)$:
 - For each gate G in layer ℓ , do (in parallel):
 - * All the committees that are associated with wires that are connected to G perform an MPC protocol (Theorem 7.6) on $C_G^{q_i}$, where $C_G^{q_i}$ is a circuit that wraps the functionality of the gate by first reconstructing the input wire values, then evaluating the gates, and finally resharing the output with the committee members of associated with the output wires. This circuit is formally defined in Section 7.3.3.

- **Output Phase (Reconstruct):** Each party P_i reconstructs the output value by pulling shares from its committee members Q_i .

7.3.1 Security

We specify the simulator \mathcal{S} for the ideal experiment with \mathcal{F}^C (Functionality 4). Note that we will get statistical security due to the statistical error in the quorum building protocol from Theorem 7.3.

1. **Initialization:** \mathcal{S} chooses the same set of corrupted parties \mathcal{C} as \mathcal{A} .
2. **Evaluation:** \mathcal{S} run the setup and share phases with \mathcal{A} , where \mathcal{S} simulates every honest party P_i with $x_i = 0$. At the end, \mathcal{S} extracts $\{\tilde{x}_i, r_i\}_{i \in \mathcal{C}}$ using $q_i/3$ shares of honest parties. Then, it sends $\{\tilde{x}_i\}_{i \in \mathcal{C}}$ to \mathcal{F}^C , and receives back $\{y_i\}_{i \in \mathcal{C}}$.
3. **Output:** \mathcal{S} run the protocol with \mathcal{A} with the same randomness as before, with the following changes:
 - (a) **Setup Phase:** \mathcal{S} invokes the simulator of the quorum building protocol from Theorem 7.3; at the end \mathcal{S} holds $f_{\mathcal{S}}$.
 - (b) **Share Phase:** \mathcal{S} invokes the simulator of n concurrent executions of the “heavy” MPC protocol, where \mathcal{S} simulates \mathcal{F}^C as follows:
 - For $P_i \in \mathcal{C}$: \mathcal{S} outputs $\text{share}_{q_i}(\tilde{x}_i, r_i)$ using the values that were extracted previously.
 - For $P_i \notin \mathcal{C}$: \mathcal{S} outputs $\text{share}_{q_i}(0, r)$ using fresh randomness.

- (c) **Evaluation Phase:** For layer $\ell \in 1 \dots \text{Depth}(C)$, let k be the number of gates in layer ℓ . \mathcal{S} invokes the simulator of k concurrent executions of the “heavy” MPC protocols, where for each such execution \mathcal{S} simulates the functionality $\mathcal{F}_G^{C_{q_i}}$ as follows:
- If the output of G corresponds to an output y_i : \mathcal{S} outputs $(\text{share}_{q_i}(y_i, r_1), \text{share}_{q_i}(0, r_2))$ using fresh randomness.
 - Else, \mathcal{S} outputs $(\text{share}_{q_i}(0, r_1), \text{share}_{q_i}(0, r_2))$ using fresh randomness.

To show that the above simulator succeeds in simulating the view of the real world, we proceed by considering five intermediate hybrids. The first hybrid corresponds to the ideal experiment, the last hybrid corresponds to the real experiment, and in each intermediate hybrid, we specify one component that we change.

- **Hybrid 0:** This world corresponds to the ideal experiment with \mathcal{S} as above.
- **Hybrid 1:** In this world \mathcal{S} behave as in Hybrid 0, except that in the Setup Phase, \mathcal{S} behaves like in the real world execution by using f instead of the simulated ones $f_{\mathcal{S}}$.
- **Hybrid 2:** In this world \mathcal{S} behave as in Hybrid 1, except that \mathcal{S} knows the honest parties inputs, and in the Share Phase, when \mathcal{S} simulates \mathcal{F}^C , for $P_i \notin \mathcal{C}$, \mathcal{S} outputs $\text{share}_{q_i}(x_i, r)$ using fresh randomness.
- **Hybrid 3:** In this world, \mathcal{S} behaves as in Hybrid 2, except that in the Share Phase, \mathcal{S} behaves as in the real world execution.
- **Hybrid 4:** In this world, \mathcal{S} behaves as in Hybrid 3, except that in the Evaluation Phase, \mathcal{S} behaves as in the real world execution. This hybrid is identical to the real world.

Let $(\text{view}_i^{(k)}, y_i^{(k)})$ be the view and output of P_i in Hybrid k . Denote $\mathcal{H} = [n] \setminus \mathcal{C}$, the set of honest parties.

Claim 7.7. *The distributions $(\{\text{view}_i^{(0)}\}_{i \in \mathcal{C}}, \{y_i^{(0)}\}_{i \in \mathcal{H}})$ and $(\{\text{view}_i^{(1)}\}_{i \in \mathcal{C}}, \{y_i^{(1)}\}_{i \in \mathcal{H}})$ are statistically close.*

Proof. Follows directly from the security of the protocol from Theorem 7.3. \square

Claim 7.8. *The distributions $(\{\text{view}_i^{(1)}\}_{i \in \mathcal{C}}, \{y_i^{(1)}\}_{i \in \mathcal{H}})$ and $(\{\text{view}_i^{(2)}\}_{i \in \mathcal{C}}, \{y_i^{(2)}\}_{i \in \mathcal{H}})$ are identical.*

Proof. By Theorem 7.6 it follows that the adversary learns nothing about \mathcal{S} 's inputs from the execution of the heavy MPC protocol. From the security of the secret sharing scheme, it follows that the adversary's outputs are independent of the inputs. Overall, the distribution of the view that is generated by the heavy MPC simulator is identical for all possible inputs. \square

Claim 7.9. *The distributions $(\{\text{view}_i^{(3)}\}_{i \in \mathcal{C}}, \{y_i^{(2)}\}_{i \in \mathcal{H}})$ and $(\{\text{view}_i^{(3)}\}_{i \in \mathcal{C}}, \{y_i^{(3)}\}_{i \in \mathcal{H}})$ are identical.*

Proof. For each party $P_i \notin \mathcal{C}$ the claim is true by Theorem 7.6. For each party $P_i \in \mathcal{C}$, we use the fact that \mathcal{S} extracts the right \tilde{x}_i and r_i . Indeed, by Theorem 7.6, \mathcal{S} receives at least $\frac{2q_i}{3}$ shares of \tilde{x}_i , and by the parameters of the secret sharing scheme, \mathcal{S} correctly extracts both \tilde{x}_i and r_i , as needed. \square

To prove that Hybrid 3 and Hybrid 4 are identical we define additional sub-hybrids as follows:

- **Hybrid $(3, \ell)$:** In this world, \mathcal{S} behaves as in Hybrid 3, except that in the first ℓ layers of the Evaluation Phase, \mathcal{S} behave as in the real world execution.

Clearly, Hybrid $(3, 0)$ is identical to Hybrid 3, and Hybrid $(3, \text{Depth}(C))$ is identical to Hybrid 4. Also, for convenience and without loss of generality, we assume that the gates connected to output wires are in layer $\text{Depth}(C)$.

Claim 7.10. *For every $\ell \in [\text{Depth}(C) - 1]$, the distributions $(\{\text{view}_i^{(3,\ell)}\}_{i \in \mathcal{C}}, \{y_i^{(3,\ell)}\}_{i \in \mathcal{H}})$ and $(\{\text{view}_i^{(3,\ell-1)}\}_{i \in \mathcal{C}}, \{y_i^{(3,\ell-1)}\}_{i \in \mathcal{H}})$ are identical.*

Proof. By the construction of the $C_G^{q_i}$ circuit, it follows that each output is computed by $\text{share}_{q_i}(\star, r^*)$, where r^* is XOR of random values submitted by committee members. From Theorem 7.6, it follows that the adversary learns nothing about the inputs from the execution of the heavy MPC protocol, and so since the XOR contains at least one unknown uniformly random value, r^* is also uniformly random and unknown to the adversary. From the security of the secret sharing scheme, it follows that the outputs of the adversary are independent of the private inputs. Overall, the distribution of the view that is generated by the heavy MPC simulator is the same for all inputs. \square

Note that, up to this point the outputs of the adversary are generated by the ideal functionality. In the following claim, we prove the “correctness” of the protocol.

Claim 7.11. *In the real execution, assuming the quorum building protocol succeeds, for every input vector \vec{x} and gate $G \in C$, the result of applying the reconstruction scheme on the outputs of share in the wrapped circuit is equal to the value of output of G in C (without wrapping).*

Proof. We prove this by induction on the layers of the circuit.

- Base case ($\ell = 0$): For each party P_i , by the security of the heavy MPC protocol, the output of each honest party $P_j \in Q_i$ at the end of the Share Phase is the output of $\text{share}_{q_i}(x_i, r_i)$. Thus, more than $2q_i/3$ honest parties hold the right shares of x_i , and so from the robustness property, the reconstruction scheme will output x_i .
- Inductive step: Assume that the claim holds for all $\ell \in [l]$. For layer $l + 1$, Let $C_G^{q_i}$ be a circuit on layer $l + 1$. In the construction of $C_G^{q_i}$, at first the reconstruction scheme is applied for each G input to obtain x_1 and x_2 , then $G(x_1, x_2)$ calculated, and lastly, the share scheme is applied for each output. From the security of the heavy MPC it follows that the result of applying the reconstruction scheme on each share circuit output will be same as G 's output.

\square

Claim 7.12. *The distributions $(\{\text{view}_i^{(3,\text{Depth}(C))}\}_{i \in \mathcal{C}}, \{y_i^{(3,\text{Depth}(C))}\}_{i \in \mathcal{H}})$ and $(\{\text{view}_i^{(3,\text{Depth}(C)-1)}\}_{i \in \mathcal{C}}, \{y_i^{(3,\text{Depth}(C)-1)}\}_{i \in \mathcal{H}})$ are identical.*

Proof. The proof of this claim is analogous to the proof of Claim 7.10. By Claim 7.11 with $\ell = \text{Depth}(C)$, it follows that the outputs in the real execution are equal to $C(\tilde{x}_1, \dots, \tilde{x}_n)$, i.e., the output of the functionality \mathcal{F}^C . \square

7.3.2 Efficiency

By Claim 7.13, each $C_G^{q_i}$ is of size polynomial in the committee size and the latter is poly-logarithmic in n (Functionality 3). That is, $|C_G^{q_i}| = \text{poly}(q_i) = \text{polylog}(n)$. Together with Theorem 7.6, each heavy MPC on $C_G^{q_i}$ costs $\text{polylog}(n)$ rounds and communication; the same goes for the share and output phases. Since for each layer in the circuit C we compute all the gates in parallel for the share, evaluation, and output phases, we need $\tilde{O}(\text{Depth}(C))$ rounds, and $\tilde{O}(|C|/n)$ communication per party. Together with Theorem 7.3, we got the following complexities:

- $\text{rc}(\Pi_{MPC}^C) \in \tilde{O}(\text{Depth}(C))$.
- For each party P_i , $\text{cc}(\Pi_{MPC}^C, i) \in \tilde{O}(|C|/n + \sqrt{n})$.
- For each party P_i , $\text{cc}(\Pi_{MPC}^C, \mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{\text{static}}, i) \in \tilde{O}(|C|/n)$.

7.3.3 The Gate Wrapper

Given a gate G with two inputs and two outputs, and a pair of circuits share_n and rec_n that represent the sharing and reconstruction mechanism of the robust variant of Shamir’s scheme [Sha79] (i.e., using the reconstruction of [WB86]), we build a circuit C_G^n with $4n$ inputs and $2n$ outputs as follows, and as depicted in Figure 1:

- Connect each input of G to the output of a new rec_n circuit.
- Connect each output of G to the first (value) input of a new share_n circuit.
- Connect the second (randomness) input of each share_n circuit to a new XOR_n circuit, where XOR_n is implemented in a tree-like manner.

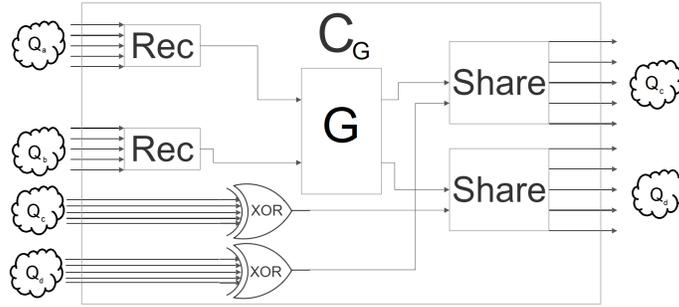


Figure 1: The wrapper of a gate G .

Since the sharing and reconstruction procedures can be implemented in polynomial time and since all other computations are efficient, we conclude that the wrapper can be written as a polynomial size circuit.

Claim 7.13. For each gate G , $|C_G^n| \in \text{poly}(n)$.

Acknowledgements

Ilan Komargodski is the incumbent of the Harry & Abe Sherman Senior Lectureship at the School of Computer Science and Engineering at the Hebrew University. Research supported in part by an Alon Young Faculty Fellowship, by a JPM Faculty Research Award, by a grant from the Israel Science Foundation (ISF Grant No. 1774/20), and by a grant from the US-Israel Binational Science Foundation and the US National Science Foundation (BSF-NSF Grant No. 2020643).

References

- [ACD⁺19] Ittai Abraham, T.-H. Hubert Chan, Danny Dolev, Kartik Nayak, Rafael Pass, Ling Ren, and Elaine Shi. Communication complexity of byzantine agreement, revisited. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC*, pages 317–326, 2019.

- [AL17] Gilad Asharov and Yehuda Lindell. A full proof of the bgw protocol for perfectly secure multiparty computation. *Journal of Cryptology*, 30(1):58–151, 2017.
- [BCG21] Elette Boyle, Ran Cohen, and Aarushi Goel. Breaking the $O(\sqrt{n})$ -bit barrier: Byzantine agreement with polylog bits per party. In *ACM Symposium on Principles of Distributed Computing, PODC*, pages 319–330, 2021.
- [BGH13] Nicolas Braud-Santoni, Rachid Guerraoui, and Florian Huc. Fast byzantine agreement. In *ACM Symposium on Principles of Distributed Computing, PODC*, pages 57–64, 2013.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, STOC*, pages 1–10, 1988.
- [BH08] Zuzana Beerliová-Trubíniová and Martin Hirt. Perfectly-secure MPC with linear communication complexity. In *Theory of Cryptography, TCC*, pages 213–230, 2008.
- [Bor96] Malte Borderding. Levels of authentication in distributed agreement. In *Distributed Algorithms, 10th International Workshop, WDAG*, pages 40–55, 1996.
- [Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science, FOCS*, pages 136–145. IEEE Computer Society, 2001.
- [CCD87] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (abstract). In *Advances in Cryptology - CRYPTO*, volume 293, page 462, 1987.
- [CCXY18] Ignacio Cascudo, Ronald Cramer, Chaoping Xing, and Chen Yuan. Amortized complexity of information-theoretically secure MPC revisited. In *Advances in Cryptology - CRYPTO*, pages 395–426, 2018.
- [CGH⁺18] Koji Chida, Daniel Genkin, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, Yehuda Lindell, and Ariel Nof. Fast large-scale honest-majority MPC for malicious adversaries. In *Advances in Cryptology - CRYPTO*, pages 34–64, 2018.
- [CL99] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the Third USENIX Symposium on Operating Systems Design and Implementation, (OSDI)*, pages 173–186, 1999.
- [DI06] Ivan Damgård and Yuval Ishai. Scalable secure multiparty computation. In *Advances in Cryptology - CRYPTO*, pages 501–520. Springer, 2006.
- [DIK⁺08] Ivan Damgård, Yuval Ishai, Mikkel Krøigaard, Jesper Buus Nielsen, and Adam D. Smith. Scalable multiparty computation with nearly optimal work and resilience. In *Advances in Cryptology - CRYPTO*, pages 241–261, 2008.
- [DIK10] Ivan Damgård, Yuval Ishai, and Mikkel Krøigaard. Perfectly secure multiparty computation and the computational overhead of cryptography. In *Advances in Cryptology - EUROCRYPT*, pages 445–465, 2010.

- [DKM⁺17] Varsha Dani, Valerie King, Mahnush Movahedi, Jared Saia, and Mahdi Zamani. Secure multi-party computation in large networks. *Distributed Computing*, 30:193–229, 2017.
- [DKMS12] Varsha Dani, Valerie King, Mahnush Movahedi, and Jared Saia. Breaking the $O(mn)$ bit barrier: Secure multiparty computation with a static adversary. In *8th Student Conference*, page 64, 2012.
- [DKMS14] Varsha Dani, Valerie King, Mahnush Movahedi, and Jared Saia. Quorums quicken queries: Efficient asynchronous secure multiparty computation. In *Distributed Computing and Networking - ICDCN*, pages 242–256, 2014.
- [DLS88] Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988.
- [DN07] Ivan Damgård and Jesper Buus Nielsen. Scalable and unconditionally secure multiparty computation. In *Advances in Cryptology - CRYPTO*, pages 572–590, 2007.
- [DS83] Danny Dolev and H. Raymond Strong. Authenticated algorithms for byzantine agreement. *SIAM J. Comput.*, 12(4):656–666, 1983.
- [FLM86] Michael J. Fischer, Nancy A. Lynch, and Michael Merritt. Easy impossibility proofs for distributed consensus problems. *Distributed Comput.*, 1(1):26–39, 1986.
- [GIP⁺14] Daniel Genkin, Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and Eran Tromer. Circuits resilient to additive attacks with applications to secure computation. In *Symposium on Theory of Computing, STOC*, pages 495–504, 2014.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, STOC*, pages 218–229, 1987.
- [Gol04] Oded Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.
- [HKK08] Dan Holtby, Bruce M. Kapron, and Valerie King. Lower bound for scalable byzantine agreement. *Distributed Comput.*, 21(4):239–248, 2008.
- [HM01] Martin Hirt and Ueli M. Maurer. Robustness for free in unconditional multi-party computation. In *Advances in Cryptology - CRYPTO*, pages 101–118, 2001.
- [IKP⁺16] Yuval Ishai, Eyal Kushilevitz, Manoj Prabhakaran, Amit Sahai, and Ching-Hua Yu. Secure protocol transformations. In *Advances in Cryptology - CRYPTO*, pages 430–458, 2016.
- [IPS09] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Secure arithmetic computation with no honest majority. In *Theory of Cryptography, TCC*, pages 294–314, 2009.
- [KLST11] Valerie King, Steven Lonargan, Jared Saia, and Amitabh Trehan. Load balanced scalable byzantine agreement through quorum building, with full information. In *Distributed Computing and Networking - ICDCN*, pages 203–214, 2011.
- [KS09] Valerie King and Jared Saia. From almost everywhere to everywhere: Byzantine agreement with $\tilde{O}(n^{3/2})$ bits. In *Distributed Computing, 23rd International Symposium, DISC*, pages 464–478, 2009.

- [KS10] Valerie King and Jared Saia. Breaking the $O(n^2)$ bit barrier: scalable byzantine agreement with an adaptive adversary. In *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 420–429, 2010.
- [KSSV06] Valerie King, Jared Saia, Vishal Sanwalani, and Erik Vee. Scalable leader election. In *17th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 990–999, 2006.
- [LSP82] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [PSL80] Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
- [Sha79] Adi Shamir. How to share a secret. *Communications of the ACM*, pages 612–613, 1979.
- [WB86] Lloyd R. Welch and Elwyn R. Berlekamp. Error correction for algebraic block codes, U.S. Patent US4633470A, 1986.

A Distribute Sub-Protocol

For convenience, we assume that party and committee IDs go from 0 to $n - 1$. In this protocol we assume a quorum, it mean that there are n committees ($Q_0 \dots Q_{n-1}$), where in each committee there is majority of honest parties that have knowledge about the quorum. When we say that committee sends a message it mean that all the parties in the committee sends this message, and when we say that a message was received from a committee, it mean that it was received from a majority of parties in the committee.

The protocol is very simple, all the committees communicate their values upstream in tree-like communication pattern, calculating a logical OR along the way, and then broadcast the results in reverse order, it takes $O(\log n)$ rounds. Since the communication pattern and messages sizes are fixed (given n), we avoid writing the filtering rules.

The DISTRIBUTE Protocol

Input: Each committee Q_i holds a bit x_i , indicating that P_i sent a message to Q_i .

Output: Each committee Q_i holds a bit x , which is the logical OR of all the indicators.

The Protocol:

- (a) For r in $1 \dots \lceil \log n \rceil + 1$:
 - i. For each committee Q_i and if $i \bmod 2^{r-1} = 0$:
 - A. If $r \neq 1 \wedge i + 2^{r-2} < n$, Q_i set y_i to be the value that was received most times from $Q_{i+2^{r-2}}$ in previous round (0 if there is a tie), else, Q_i set $y_i \leftarrow 0$.
 - B. Q_i set $x_i \leftarrow x_i \vee y_i$.
 - C. If $i \bmod 2^r \neq 0$, Q_i sends (i, x_i) to each party in $Q_{i-2^{r-1}}$.
- (b) Q_0 set $x \leftarrow x_0$, and all other committees set $x \leftarrow \perp$.
- (c) For r in $\lceil \log n \rceil \dots 0$:
 - i. For each committee Q_i and if $i \bmod 2^r = 0$:
 - A. If $i \neq 0 \wedge i \bmod 2^{r+1} \neq 0$, Q_i set x to be the value that was received most times from parties in Q_{i-2^r} in previous round (0 if there is a tie).
 - B. If $r \neq 0$, Q_i sends x_0 to each party in $Q_{i+2^{r-1}}$.

B Quorum Building Protocol

In this protocol, we associate Fail with 1 and Success with 0, so logical Boolean operations are well defined.

Lightweight Quorum Building Protocol (Π_{light})

- **Setup Phase:** Each party P_i sets $y_i \leftarrow \text{Success}$.
- **Almost-Everywhere Phase:** All the parties run together the almost-every agreement protocol from [KLST11]. At the end, each party P_i holds a string str_i of size $\Theta(\log^3 n)$.
- **Agreement Phase:**

1. Each party P_i samples a set of $\log^2 n$ parties (denoted S_i), and sends them “ping”.
Filtering: Each party accepts messages of size $|\text{“ping”}|$, and at most $2 \cdot \log^2 n$ of them.
2. For each party P_j : if it received less than $2 \cdot \log^2 n$ “ping”s, P_j answers with str_j ; else, P_j sets $y_j \leftarrow \text{Fail}$.
Filtering: P_i accepts messages of size at most $\log^4 n$ and only from parties in S_i .
3. For each party P_i , if there is a string str^* that was received more than $|S_i|/2$ times, P_i sets $\text{str}_i \leftarrow \text{str}^*$; else, P_i sets $y_i \leftarrow \text{Fail}$.

- **Output Phase:** For each party P_i if $y_i = \text{Fail}$ then output Fail ; else, output $y_i = f_{\text{str}_i}(\star) \triangleq H(\star, \text{str}_i)$, where H is the $\frac{1}{3}$ -quorum building function from Theorem 5.3.

Corollary B.1 (Restatement of Lemma 7.2). *There is a protocol Π_{light} (described above) that realizes $\mathcal{F}_{1/3\text{-quorum}}$ against $\mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{\text{static}}$ and realizes $\mathcal{F}_{1/3\text{-quorum}}^{\text{LFD}}$ against $\mathcal{A}_{1/3}^{\text{static}}$.*

Proof. Since the inputs are fixed in the protocol description, the simulator for an adversary \mathcal{A} simulates the honest parties to obtain the view of \mathcal{A} . By Claim 5.5, it follows that the output of each honest party is either f_{str_i} (the same function for all honest parties) or Fail . By Theorem 5.3 (see also [KLST11]) with all but negligible probability of error, for str with $\Omega(\log^2 n)$ (min-)entropy, $H(\star, \text{str})$ defines a good quorum. Thus, $f_{\text{str}} \in \mathcal{F}_{\frac{1}{3}\text{-quorum}}(\mathcal{C})$, and so Π_{light} realizes $\mathcal{F}_{\frac{1}{3}\text{-quorum}}^{\text{LFD}}$ against \mathcal{A} .

By Claim 5.8 it follows that if $\mathcal{A} \in \mathcal{A}_{SHwD} \cap \mathcal{A}_{1/3}^{\text{static}}$, then the output of all honest parties is f_{str} (nobody will output Fail). This means that Π_{light} realizes $\mathcal{F}_{\frac{1}{3}\text{-quorum}}$ against \mathcal{A} . \square

Finally, we apply our optimistic-pessimistic framework from Theorem 6.1 using Lemmas 7.1 and 7.2, to complete the proof of Theorem 7.3.