

Error-free Multi-valued Broadcast and Byzantine Agreement with Optimal Communication Complexity

Arpita Patra

Department of Computer Science, ETH Zurich, Switzerland
arpita.patra@inf.ethz.ch

Abstract. In this paper we present first ever *error-free, asynchronous* broadcast (called as A-cast) and Byzantine Agreement (called as ABA) protocols with optimal communication complexity and fault tolerance. Our protocols are multi-valued, meaning that they deal with ℓ bit input and achieve communication complexity of $\mathcal{O}(n\ell)$ bits for large enough ℓ for a set of $n \geq 3t + 1$ parties in which at most t can be Byzantine corrupted. Previously, Patra and Rangan (Latincrypt'10, ICITS'11) reported multi-valued, communication optimal A-cast and ABA protocols *that are only probabilistically correct*.

Following all the previous works on multi-valued protocols, we too follow reduction-based approach for our protocols, meaning that our protocols are designed given existing A-cast and ABA protocols for small message (possibly for single bit). Our reductions invoke less or equal number of instances of protocols for single bit in comparison to the reductions of Patra and Rangan. Furthermore, our reductions run in constant expected time, in contrast to $\mathcal{O}(n)$ of Patra and Rangan (ICITS'11). Also our reductions are much simpler and more elegant than their reductions.

By adapting our techniques from asynchronous settings, we present new *error-free*, communication optimal reduction-based protocols for broadcast (BC) and Byzantine Agreement (BA) in synchronous settings that are constant-round and call for only $\mathcal{O}(n^2)$ instances of protocols for single bit. Prior to this, communication optimality has been achieved by Fitzi and Hirt (PODC'06) who proposed *probabilistically correct* multi-valued BC and BA protocols with constant-round and $\mathcal{O}(n(n + \kappa))$ (κ is the error parameter) invocations to the single bit protocols. Recently, Liang and Vaidya (PODC'11) achieved the same *without* error probability. However, their reduction calls for round complexity and number of instances that are function of the message size, $\mathcal{O}(\sqrt{\ell} + n^2)$ and $\mathcal{O}(n^2\sqrt{\ell} + n^4)$, respectively where $\ell = \Omega(n^6)$.

Keywords: Multi-valued, Broadcast, Byzantine Agreement, A-cast, Asynchronous, Communication Complexity, Expected Running Time

1 Introduction

The problem of Broadcast (BC) and Byzantine Agreement (BA) (also popularly known as consensus) were introduced in [PSL80] and since then they have

been considered as the most fundamental problems in distributed computing. In brief, a BC protocol allows a special party among a set of parties, called sender, to send some message identically to all other parties. The challenge lies in achieving the above task despite the presence of some faulty parties (possibly including the sender), who may deviate from the protocol arbitrarily. The BA primitive is slightly different from BC. A BA protocol allows a set of parties, each holding some input bit, to agree on a common bit, even though some of the parties may act maliciously in order to make the honest parties disagree. The BC and BA primitives have been used as building blocks in several important secure distributed computing tasks such as Secure Multiparty Computation (MPC) [BOGW88,BKR94,RBO89], Verifiable Secret Sharing (VSS) [CGMA85,BOGW88,RBO89] etc.

An important, practically motivated variant of BC and BA problem are asynchronous broadcast (known as A-cast) and asynchronous BA (known as ABA) that study the conventional BC and BA problems in asynchronous network settings. It is well-known that asynchronous network setting is considered to be more realistic than synchronous network setting. The works of [BO83,Rab83,Bra84], [FM88,CR93,Can95,ADH08,PW92,PR11] have reported different A-cast and ABA protocols. In this paper, we focus on the communication complexity of *error-free* A-cast and ABA protocols and present first ever optimal protocols.

The Model. We follow the standard network model of [PSL80] for synchronous network and [CR93,Can95] for asynchronous network. Our A-cast, ABA, BC and BA protocols are carried out among a set of n parties, say $\mathcal{P} = \{P_1, \dots, P_n\}$, where every two parties are directly connected by an authenticated and secure channel and at most t out of the n parties can be under the influence of a *computationally unbounded Byzantine adversary*, denoted as \mathcal{A}_t . The adversary corrupts the parties *adaptively* at any point during the course of the protocol execution and the choice may base on the information gathered so far by the adversary. We assume that $n = 3t + 1$ which is the minimum number of parties required to design *error-free* A-cast, ABA, BC and BA protocols [Lyn96,PSL80]. The parties not under the influence of \mathcal{A}_t are called *honest or uncorrupted*.

We do not make any cryptographic assumptions such as public key infrastructure (PKI) etc in our protocols. All our protocols are *randomized*.

Definitions. We now define A-cast and ABA formally.

Definition 11 (A-cast [Can95]) *Let Π be an asynchronous protocol executed among the set of parties \mathcal{P} and initiated by a special party called sender $S \in \mathcal{P}$, having input m (the message to be sent). Π is an A-cast protocol tolerating \mathcal{A}_t if the following hold, for every behavior of \mathcal{A}_t and every input m :*

- **Termination:** *If S is honest, then all honest parties in \mathcal{P} will eventually terminate Π . If any honest party terminates Π , then all other honest parties will eventually terminate Π .*
- **Correctness:** *If the honest parties terminate Π , then they do so with a common output m^* . Furthermore, if the sender S is honest then $m^* = m$.*

Definition 12 (ABA [CR93]) Let Π be an asynchronous protocol executed among the set of parties \mathcal{P} , with each party having a private binary input. We say that Π is an ABA protocol tolerating \mathcal{A}_t if the following hold, for every possible behavior of \mathcal{A}_t and every possible input:

- **Termination:** All honest parties eventually terminate the protocol.
- **Correctness:** All honest parties who have terminated the protocol hold identical outputs. Furthermore, if all honest parties had same input, say ρ , then all honest parties output ρ .

The celebrated result of [FLP85] shows that any ABA protocol that never reaches disagreement must have some nonterminating executions. For a protocol that never reaches disagreement, the best we can hope for is that the set of nonterminating executions has probability zero. Such protocols are termed as *almost-surely terminating* by [ADH08]. In this work, we construct ABA protocol that is almost-surely terminating and has no error in correctness. The important complexity measures of A-cast and ABA protocol are: **Communication Complexity:** It is the total number of bits communicated by the *honest* parties in the protocol; **Expected Running Time:** Refer to [CR93,Can95] for a detailed definition of expected running time of a randomized asynchronous protocol.

While the basic definitions of A-cast and ABA consider message of single bit, *multi-valued* protocols allow message to be long string of bits and exploit the fact that the task is to be attained for the entire string and not bit by bit. This fact generally allows a *multi-valued* protocol to be considerably more efficient than many parallel executions of protocol for single bit.

Brief Literature. Error-free BC and BA protocol in synchronous network are possible if and only if $n \geq 3t + 1$ [PSL80,Lyn96]. The same bound holds for A-cast and ABA both with and without error probability [Lyn96]. The seminal result of [DR85] shows that any error-free BA or BC must communicate $\Omega(n^2)$ bits (which again carry over for the case of A-cast and ABA). Since the message must be at least single bit, the lower bound on the communication complexity for single bit is $\Omega(n^2)$ bits. However, communication complexity of $\mathcal{O}(n\ell)$ bits can be achieved for large enough value of ℓ (at least $\ell \geq n$ bits) as shown in [FH06]. Requiring large value for ℓ is practically motivated in many distributed computing applications, like reaching agreement on a large file in fault-tolerant distributed storage system, distributed voting where ballots containing gigabytes of data is to be handled, MPC where many broadcasts and agreements are invoked which can be combined into fewer executions of multi-valued protocols.

Following the approach of Turpin and Coan [TC84], all the subsequent multi-valued protocols apply reduction-based approach [FH06,LV11,PR11,PR10], meaning that they are constructed based on access to protocols for small message or single bit message. The reductions presented in [FH06,LV11] for synchronous settings and in [PR11,PR10] for asynchronous settings achieve optimal communication complexity of $\mathcal{O}(n\ell)$ bits. While the reduction presented in [FH06] involves error probability, the reduction of [LV11] is error-free. In the asynchronous settings, [PR11,PR10] reported multi-valued protocols with error probability.

Our Contribution. We achieve optimal complexity of $\mathcal{O}(n\ell)$ bits for *error-free* A-cast and ABA with optimal fault-tolerance of $t < n/3$. We too follow reduction-based approach of [TC84]. We now compare our reductions with that of [PR10,PR11] and show that our reductions are better in all the following aspects: (a) error-free, (b) running time and (c) number of invocations to protocols for single bit. All the protocols have optimal fault tolerance of $t < n/3$.

Ref.	Type	Running Time	# Invocations to single bit protocol
[PR10], A-cast	Probabilistic	constant	$\mathcal{O}(n^2 \log n)$ A-cast
[PR11], ABA	Probabilistic	$\mathcal{O}(n)$	$\mathcal{O}(n^3)$ ABA
This paper, A-cast	Error-free	constant	$\mathcal{O}(n^2 \log n)$ A-cast
This paper, ABA	Error-free	constant	$\mathcal{O}(n)$ ABA

We now compare our results with the current best *error-free* A-cast and ABA for single bit. The only error-free A-cast is due to [Bra84] that communicates $\mathcal{O}(n^2)$ bits and runs in constant time. Similarly, the only error-free ABA is due to [ADH08] that runs in $\mathcal{O}(n^2)$ time and requires communication of $\mathcal{O}(n^8 \log n)$ bits and A-cast of same number of bits. Our protocols in this paper show clear improvement over ℓ executions of these protocols for large enough ℓ .

Technically, our reductions are simple and are based on linear error correcting code (e.g. Reed-Solomon Code) and a graph theoretic algorithm for finding some special structure (called (n, t) -star; defined in Section 2) in undirected graph [CR93,Can95]. While the existing reductions for multi-valued protocols [PR11,LV11] are constructed in player-elimination [HMP00] or dispute control [BTH06] framework, our reductions do not require them and therefore they are more elegant. Finally, we note that the multi-valued A-cast protocol of [PR10] also employs the algorithm for finding (n, t) -star [CR93,Can95]. However, we mark an important and crucial observation about the outcome of the algorithm in our context that allows to construct our protocol in an error-free manner.

Finally, we discuss our results in synchronous settings. By adapting our techniques from asynchronous settings, we present new *error-free* reduction that is constant-round and calls for $\mathcal{O}(n^2)$ instances of protocols for single bit. We now compare our result with the communication optimal reductions of [FH06,LV11].

Ref.	Type	Fault Tolerance	Round Complexity	# Invocations to single bit protocol
[FH06]	Probabilistic	$t < n/2$	constant	$\mathcal{O}(n(n + \kappa))$
[LV11]	Error-free	$t < n/3$	$\mathcal{O}(\sqrt{\ell} + n^2)$	$\mathcal{O}(n^2 \sqrt{\ell} + n^4)$
This paper	Error-free	$t < n/3$	constant	$\mathcal{O}(n^2)$

Road-map. In section 2 and 3, we present our construction for A-cast and ABA respectively. We present our BA and BC protocols in Section 4 and then conclude in Section 5.

2 Error-free Communication Optimal A-cast

Here we present our A-cast protocol. We start with brief presentation of the tools that we use: (a) A-cast protocol of Bracha [Bra84]; (b) An algorithm for

finding a graphical structure called (n, t) -star in an undirected graph; (c) Linear Error Correcting Code. We discuss them one by one.

Bracha's A-cast. The first ever protocol for A-cast is due to Bracha [Bra84] (a good description is available in [Can95]). The protocol is error-free, runs with $n \geq 3t+1$ in constant time and communicates $\mathcal{O}(n^2)$ bits for a *single bit* message.

Notation 21 *By saying that ' P_i A-casts M ', we mean that P_i as a sender, initiates Bracha's A-cast protocol with M as the message. Similarly ' P_j receives M from the A-cast of P_i ' will mean that P_j terminates the A-cast protocol initiated by P_i and outputs M . By the property of A-cast, if some honest party P_j terminates the A-cast of some sender P_i with M as the output, then every other honest party will eventually do so, irrespective of the behavior of the sender P_i .*

Finding (n, t) -star in an Undirected Graph. We now describe an existing solution for a graph theoretic problem, called finding (n, t) -star in an undirected graph $G = (V, E)$. Let G be an undirected graph with the n parties in \mathcal{P} as its vertex set. A pair $(\mathcal{C}, \mathcal{D})$ of sets with $\mathcal{C} \subseteq \mathcal{D} \subseteq \mathcal{P}$ is an (n, t) -star [Can95,BOCG93] in G , if: (i) $|\mathcal{C}| \geq n - 2t$; (ii) $|\mathcal{D}| \geq n - t$; (iii) for every $P_j \in \mathcal{C}$ and every $P_k \in \mathcal{D}$ the edge (P_j, P_k) exists in G .

Following the idea of [GJ79], [BOCG93] presented an elegant and efficient algorithm for finding an (n, t) -star in a graph of n nodes, *provided that the graph contains a clique of size $n - t$* . Actually, the algorithm, called as Find-STAR takes the complementary graph \overline{G} of G as input and tries to find (n, t) - $\overline{\text{star}}$ in \overline{G} , where (n, t) - $\overline{\text{star}}$ is a pair $(\mathcal{C}, \mathcal{D})$ of sets with $\mathcal{C} \subseteq \mathcal{D} \subseteq \mathcal{P}$, satisfying the following conditions: (a) $|\mathcal{C}| \geq n - 2t$; (b) $|\mathcal{D}| \geq n - t$; (c) There are no edges between the nodes in \mathcal{C} and nodes in $\mathcal{C} \cup \mathcal{D}$ in \overline{G} . Clearly, a pair $(\mathcal{C}, \mathcal{D})$ representing an (n, t) - $\overline{\text{star}}$ in \overline{G} , is an (n, t) -star in G . Find-STAR outputs either an (n, t) - $\overline{\text{star}}$, or a message **star-Not-Found**. Whenever *the input graph \overline{G} contains an independent set of size $n - t$* , Find-STAR always outputs an (n, t) - $\overline{\text{star}}$. For simple notation, we denote \overline{G} by H . The algorithm Find-STAR is presented below:

Algorithm Find-STAR(H)

1. Find a maximum matching M in H . Let N be the set of matched nodes (namely, the endpoints of the edges in M), and let $\overline{N} = \mathcal{P} \setminus N$.
2. Compute output as follows:
 - (a) Let $T = \{P_i \in \overline{N} | \exists P_j, P_k \text{ s.t. } (P_j, P_k) \in M \text{ and } (P_i, P_j), (P_i, P_k) \in E\}$. T is called the set of triangle-heads. Let $\mathcal{C} = \overline{N} \setminus T$.
 - (b) Let B be the set of matched nodes that have neighbors in \mathcal{C} . So $B = \{P_j \in N | \exists P_i \in \mathcal{C} \text{ s.t. } (P_i, P_j) \in E\}$. Let $\mathcal{D} = \mathcal{P} \setminus B$.
 - (c) If $|\mathcal{C}| \geq n - 2t$ and $|\mathcal{D}| \geq n - t$, output $(\mathcal{C}, \mathcal{D})$. Otherwise, output **star-Not-Found**.

Linear Error Correcting Code. We use Reed-Solomon (RS) codes in our protocols. We consider an $(n, t + 1)$ RS code in Galois Field $\mathbb{F} = GF(2^c)$, where $n \leq 2^c$. Each element of \mathbb{F} is represented by c bits. An $(n, t + 1)$ RS code encodes

$t+1$ elements of \mathbb{F} into a codeword consisting of n elements from \mathbb{F} . We denote the encoding function as $\text{ENC}()$ and the corresponding decoding function as $\text{DEC}()$. Let m_0, m_1, \dots, m_t be the input to ENC , then ENC computes a codeword of length n , (s_1, \dots, s_n) , as follows: It constructs a polynomial of degree- t , $f(x) = m_0 + m_1x + \dots + m_tx^t$. It then computes $s_i = f(i)$. We use the following syntax for ENC : $(s_1, s_2, \dots, s_n) = \text{ENC}(m_0, m_1, \dots, m_t)$. Each element of the codeword is computed as a linear combination of the $t+1$ input data elements, such that every subset of $(t+1)$ elements from the codeword uniquely determine the input data elements. Similarly, knowledge of any $t+1$ elements from the codeword suffices to determine the remaining elements of the codeword.

The decoding function DEC can be applied as long as $t+1$ elements from a codeword are available. A RS code is capable of error correction and detection. The task of error correction is to find the error locations and error values in a received vector. On the other hand, error detection means an indication that errors have occurred, without attempting to correct them. We recall the following well known result from coding theory. DEC can correct up to c Byzantine error and simultaneously detect up to additional d Byzantine errors in a vector of length N (where $N \leq n$) if and only if $N - t - 1 \geq 2c + d$. In our protocols, we may invoke DEC on a vector of length $N \leq n$ with specific value of c and d . If c , d and N satisfy the above relation, then DEC returns back the correct data elements corresponding to the vector; otherwise DEC returns ‘failure’.

2.1 Multi-valued A-cast Protocol

With the above tools, we are now ready to present our multi-valued A-cast protocol, called **Multi-Valued-Acast**. We assume that the sender S has a message m containing ℓ bits that he would like to communicate to all the parties in \mathcal{P} identically. Our protocol is structured into two phases, (a) **S-dependent Phase** and (b) **S-independent Phase**. In the S-dependent phase, S proves that it has communicated the same message to at least a set of $2t+1$ parties, say CORE . The S-dependent phase, as the name suggests, demands S to perform some special roles. For an honest S , this phase will always be completed successfully. However, a corrupted S may choose not to perform his actions and therefore this phase may not be terminated for a corrupted S . The second phase, called S-independent phase is initiated upon completion of the first phase. If S successfully proves the existence of some CORE in the first phase, then the parties in CORE propagate their common message to the remaining parties without any help from S .

In the first phase, S communicates his message m to every party over private channel. Upon receiving a message from S , a party applies ENC on the message to get a codeword and communicates elements of the codeword to different party. Intuitively, the parties here check if they received the same message from S . They A-cast [Bra84] their responses. Based on the response of the parties, a consistency graph is constructed by the parties individually. S now finds a special structure in the graph, namely a quadruple $(\mathcal{C}, \mathcal{D}, \mathcal{F}, \mathcal{E})$ such that $(\mathcal{C}, \mathcal{D})$ is an (n, t) -star, $|\mathcal{F}| \geq 2t+1$ and every party in \mathcal{F} has at least $t+1$ neighbors in \mathcal{C} , $|\mathcal{E}| \geq 2t+1$ and every party in \mathcal{E} has at least $2t+1$ neighbors in \mathcal{F} . Such a quadruple essentially

proves that there is a set of at least $2t + 1$ parties, $CORE$ (same as \mathcal{E}), to whom S indeed communicated same message. On finding such a quadruple, S A-casts the same and all other parties can verify if indeed such quadruple exists in their individual graph. In this process, all the (honest) parties agree on $CORE$ and proceed to second phase. The algorithm for finding (n, t) -star and an important observation are combined intelligibly in order to find a quadruple in a graph. The observation is that *if S is honest then eventually, the set \mathcal{C} of an (n, t) -star will contain at least $t + 1$ honest parties and when it happens, \mathcal{F} and \mathcal{E} can be computed such that a valid quadruple can be formed.* In the second phase, the parties use error correction and detection of RS code to compute and agree on the common message of the parties in $CORE$. We present the protocol in Figure 1 and Figure 2 and subsequently prove the properties.

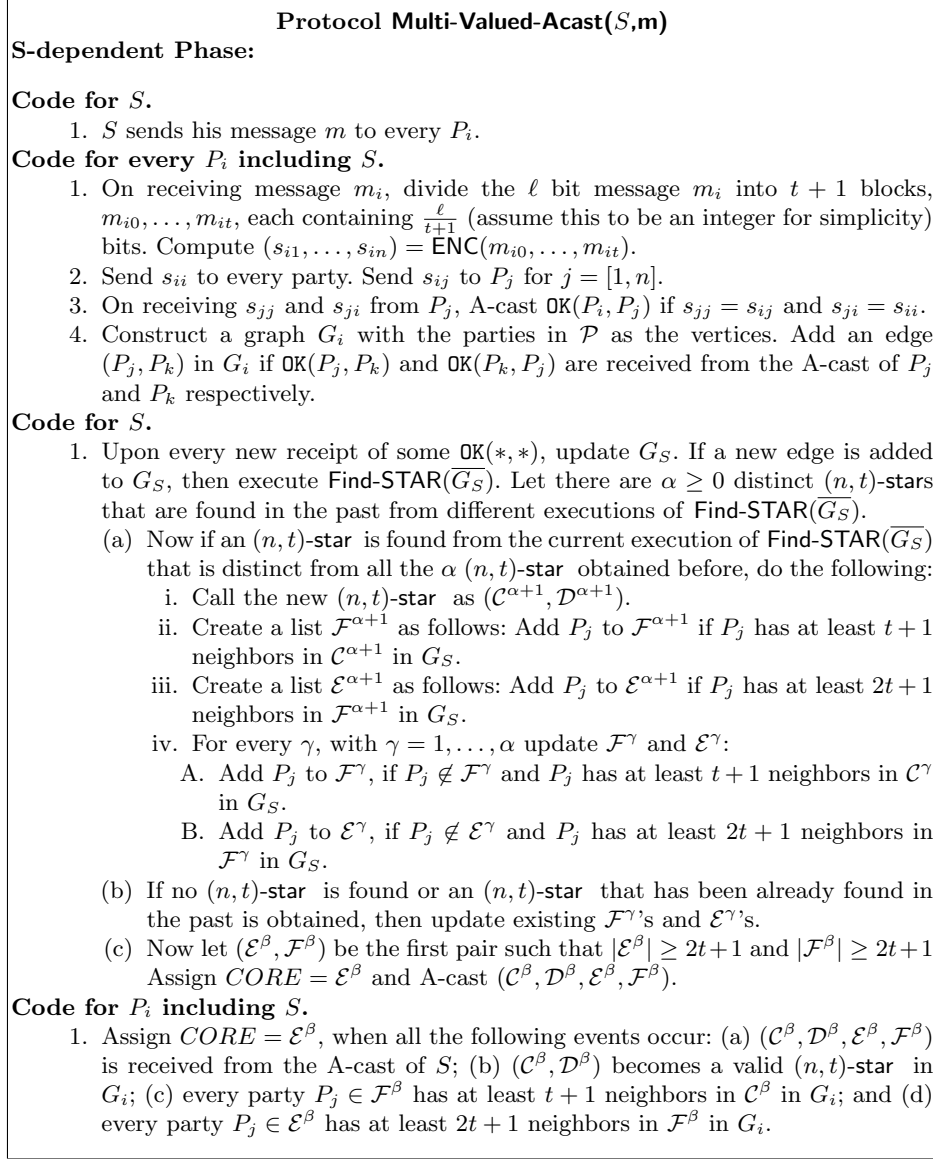
Lemma 22 *The honest parties in $CORE$ hold same message of length ℓ . If S is honest then the message is S 's message.*

Proof. The set $CORE$ is the \mathcal{E} component of a quadruple $(\mathcal{C}, \mathcal{D}, \mathcal{F}, \mathcal{E})$. We start with proving that the honest parties in \mathcal{C} hold the same message of length ℓ . We recall that \mathcal{D} contains at least $t + 1$ honest parties and every $P_i \in \mathcal{C}$ is neighbor of every party in \mathcal{D} . Let $\{P_{i_1}, \dots, P_{i_\alpha}\}$ be the set of α honest parties in \mathcal{D} , where $\alpha \geq t + 1$. Then for every P_i in \mathcal{C} , s_{ii_k} is same as $s_{i_k i_k}$ of all $k = [1, \alpha]$. Therefore the codewords corresponding to the messages of the honest parties in \mathcal{C} are same at least at $t + 1$ locations corresponding to the identities of the honest parties in \mathcal{D} . Since the codewords belong to $(n, t + 1)$ RS code, the messages of the honest parties in \mathcal{C} are same. Let the common message be m , $|m| = \ell$. Let $(s_1, \dots, s_n) = \text{ENC}(m_0, m_1, \dots, m_t)$, where $m = m_0 | m_1 | \dots | m_t$. Now we show that every honest party $P_i \in \mathcal{F}$ holds s_i . Recall that P_i has at least $t + 1$ neighbors in \mathcal{C} in which at least one is honest, say P_j . This implies that s_{ii} of P_i is same as s_{ji} of P_j . However, $s_{ji} = s_i$, since P_j holds m . Hence $s_{ii} = s_i$. Therefore every honest P_i in \mathcal{F} holds s_i which is same as s_{ii} . Finally, we show that every honest $P_i \in \mathcal{E}$ holds m . Recall that P_i has at least $2t + 1$ neighbors in \mathcal{F} in which at least $t + 1$ are honest. Let $\{P_{i_1}, \dots, P_{i_\alpha}\}$ be the set of α honest parties in \mathcal{F} , where $\alpha \geq t + 1$. Then s_{ii_k} of P_i is same as $s_{i_k i_k}$ of every honest P_{i_k} for $k = [1, \alpha]$. Now $s_{i_k i_k}$ of P_{i_k} is same as s_{i_k} . Therefore the codeword corresponding to the message of $P_i \in \mathcal{E}$ matches with (s_1, \dots, s_n) at least at $t + 1$ locations corresponding to the identities of the honest parties in \mathcal{F} . This implies the codeword of P_i is identical to (s_1, \dots, s_n) , since they belong to $(n, t + 1)$ RS code. Hence $P_i \in \mathcal{E}$ holds m . This completes the proof for the first part of the lemma. The second part of the lemma is easy to prove. \square

To prove the lemma below, we will show that when S is honest then eventually an (n, t) -star can be found such that the set \mathcal{C} will contain at least $t + 1$ honest parties. This observation is very crucial and is at the heart of our protocol.

Lemma 23 *If S is honest, then all the parties terminate **S-dependent Phase**, after agreeing on $CORE$.*

Fig. 1. Error-free Communication Optimal A-cast.



Proof. If S is honest, then he sends same message m to all the parties. Therefore, all honest parties generate same codeword, $(s_1, \dots, s_n) = \text{ENC}(m_0, \dots, m_t)$, such that $m = m_0|m_1| \dots |m_t$. Therefore eventually there will be an edge between every pair of honest parties. This implies that there will be a clique of size

Fig. 2. Error-free Communication Optimal A-cast.

S-independent Phase:

Code for P_i including S .

1. If $CORE$ is constructed and $P_i \in CORE$, then assign $s_i = s_{ii}$.
2. If $CORE$ is constructed and $P_i \notin CORE$, then assign s_i to be the value s_{ji} that is received from at least $t + 1$ P_j 's in $CORE$.
3. Send s_i to all the parties.
4. On receiving $2t + 1 + r$, $r \geq 0$, s_j 's apply DEC with $c = r$ and $d = t - r$, if DEC returns 'failure', then wait for more values. If DEC returns data elements m_0, \dots, m_t , then output $m = m_0|m_1|\dots|m_t$, where $|$ denotes concatenation.

at least $2t + 1$ eventually. This guarantees that S will eventually find at least one (n, t) -star in G_S . Now we show that S will eventually find a quadruple $(\mathcal{C}, \mathcal{D}, \mathcal{F}, \mathcal{E})$ such that $(\mathcal{C}, \mathcal{D})$ is an (n, t) -star and every party in \mathcal{F} has at least $t + 1$ neighbors in \mathcal{C} and every party in \mathcal{E} has at least $2t + 1$ neighbors in \mathcal{F} . To prove this we start with proving that an honest S will eventually find an (n, t) -star such that the set \mathcal{C} will contain at least $t + 1$ honest parties. For an honest S , eventually the edges between each pair of honest parties will vanish from the complementary graph $\overline{G_S}$. So the edges in $\overline{G_S}$ will be either (a) between an honest and a corrupted party OR (b) between two corrupted parties. Let β be the first index, such that (n, t) -star $(\mathcal{C}^\beta, \mathcal{D}^\beta)$ is generated in $\overline{G_S}$, when $\overline{G_S}$ contains edges of above two types only. Now, by construction of \mathcal{C}^β (see Algorithm Find-STAR), it excludes the parties in N (set of parties that are endpoints of the edges of maximum matching M) and T (set of parties that are triangle-heads). An honest P_i belonging to N implies that $(P_i, P_j) \in M$ for some P_j and hence P_j is corrupted (as the current $\overline{G_S}$ does not have edge between two honest parties). Similarly, an honest party P_i belonging to T implies that there is some $(P_j, P_k) \in M$ such that (P_i, P_j) and (P_j, P_k) are edges in $\overline{G_S}$. This clearly implies that both P_j and P_k are surely corrupted. So for every honest P_i not in \mathcal{C}^β , at least one (if P_i belongs to N , then one; if P_i belongs to T , then two) corrupted party also remains outside \mathcal{C}^β . As there are at most t corrupted parties, \mathcal{C}^β may exclude at most t honest parties. Still \mathcal{C}^β is bound to contain at least $t + 1$ honest parties.

Now all honest parties will be neighbors of the $t + 1$ honest parties in \mathcal{C}^β in G_S . Therefore \mathcal{F}^β will eventually contain all the honest parties. Finally since all honest parties are neighbors of each other, \mathcal{E}^β will contain all honest parties eventually and therefore it is guaranteed to contain at least $2t + 1$ parties. Hence we proved that S can find a quadruple $(\mathcal{C}, \mathcal{D}, \mathcal{F}, \mathcal{E})$ with the required properties. S now A-casts the quadruple.

We now argue that every honest party will find $(\mathcal{C}, \mathcal{D}, \mathcal{F}, \mathcal{E})$ in their graphs and agree on the same. Though the graphs are constructed and maintained by parties individually in their local memory, it is always guaranteed that if an edge

appears in the graph of an honest party, then the edge will eventually appear in the graphs of the other honest parties. This is ensured since the graphs are updated based on the responses of the parties that are A-casted. It now follows that if some honest party agree on *CORE*, then eventually all honest parties will also agree on the same. So we proved that all the honest parties will terminate **S-dependent Phase**, after agreeing on $CORE = \mathcal{E}$. \square

Lemma 24 *If the honest parties initiate S-independent Phase, then they terminate the phase with the common message of the parties in CORE as output.*

Proof. An honest party initiates **S-independent Phase**, if he agrees on *CORE*. By Lemma 22, all the honest parties in *CORE* hold common message, say m of length ℓ and therefore same codeword $(s_1, \dots, s_n) = \text{ENC}(m_0, m_1, \dots, m_t)$, where $m = m_0|m_1|\dots|m_t$. Then every honest P_i in *CORE* already holds s_i , the i th element in the codeword. Every party P_i not in *CORE* would receive s_i from the $t+1$ honest parties of *CORE*. Therefore every honest P_i will eventually hold the i th component of the codeword. Now every P_i send his s_i to every other party. Now on receiving at least $2t+1+r$, $0 \leq r \leq t$ s_j 's, party P_i applies DEC with $c=r$ and $d=t-r$. Note that $c+d=t$, where t is the maximum number of corruption. Therefore if there are more than r wrong values (sent by Byzantine corrupted parties), DEC will return 'failure'. However for at least one value of r , $0 \leq r \leq t$, there will be at most r errors in the received vector and then the message can be reconstructed back successfully. This technique has been previously used in [CR93, Can95]. They call it as *Online Error Correction*. \square

Theorem 21 *Multi-Valued-Acast is an A-cast protocol satisfying Definition 11.*

Proof. We first consider the case of an honest S . By Lemma 23, for an honest S all the parties terminate **S-dependent Phase**, after agreeing on *CORE*. By Lemma 22, the honest parties in *CORE* hold the message of S , i.e. m . By Lemma 24, all honest parties will terminate with the common message m .

For a corrupted S , all we need to show is that if some honest party terminates with message m^* , then every other honest party do the same. Let P_i be the first honest party to terminate the protocol with m^* as output. Then P_i must have agreed on *CORE* and the parties in *CORE* holds m^* . Then every other honest party will agree on the same *CORE* and eventually terminate with m^* as the output (by Lemma 24). \square

Theorem 22 *Multi-Valued-Acast communicates $\mathcal{O}(n\ell)$ bits and invokes $\mathcal{O}(n^2 \log n)$ A-cast protocol for single bit.*

Proof. S communicates his message m , $|m| = \ell$ to all the parties. This requires $n\ell$ bits of communication. Every party P_i sends two values s_{ii} and s_{ij} to every other party P_j . The values are $\frac{\ell}{t+1}$ bits long each. Therefore in total there are $\frac{\ell}{t+1} \mathcal{O}(n^2) = \mathcal{O}(n\ell)$ bits of communication.

S A-casts $(\mathcal{C}, \mathcal{D}, \mathcal{F}, \mathcal{E})$. Each set in the quadruple can be represented by an n length bit vector. Therefore $4n$ invocations to A-cast protocol for single bit are

required. Finally every party may A-cast OK signal for every other party. Each OK signal includes identities of two parties that can be represented by $2 \log n$ bits. Therefore $\mathcal{O}(n^2 \log n)$ invocations to A-cast for single bit are required. \square

We note that for an $(n, t + 1)$ RS code, the field $\mathbb{F} = GF(2^c)$ in which the code is defined should satisfy $n \leq 2^c$ or $\log n \leq c$. In our case $c = \frac{\ell}{t+1}$ (recall that m is divided into $t + 1$ parts each containing $\frac{\ell}{t+1}$ bits). Therefore $\frac{\ell}{t+1} \geq \log n \rightarrow \ell \geq (t + 1) \log n$.

3 Error-free Communication Optimal ABA

In this section, we present our ABA protocol. We use our multi-valued A-cast protocol Multi-Valued-Acast from the previous section as one of the sub-protocols. Similar to Multi-Valued-Acast that uses A-cast protocol for single bit, our new ABA uses existing error-free ABA for single bit as another sub-protocol. In fact we use a very well-known asynchronous primitive called Agreement on Common Subset (ACS) introduced by [BKR94] that uses ABA for single bit as black box. We recall that the only error-free ABA is due to [ADH08]. We will use the following notation for invoking Multi-Valued-Acast.

Notation 31 *By saying that ‘ P_i Multi-casts M ’, we mean that P_i as a sender, initiates Multi-Valued-Acast protocol with M as the message. Similarly ‘ P_j receives M from the Multi-cast of P_i ’ will mean that P_j terminates the execution of Multi-Valued-Acast protocol initiated by P_i and outputs M . By the property of Multi-Valued-Acast, if some honest party P_j terminates the Multi-Valued-Acast protocol of some sender P_i with M as the output, then every other honest party will eventually do so, irrespective of the behavior of the sender P_i .*

Agreement on a Common Subset (ACS). Consider the following scenario. The parties in \mathcal{P} are asked to A-cast (or Multi-cast) some value. While the honest parties in \mathcal{P} will eventually execute the A-cast (Multi-cast), the corrupted parties may or may not do the same. So the (honest) parties in \mathcal{P} want to agree on a common set $\mathcal{T} \subset \mathcal{P}$, with $|\mathcal{T}| = 2t + 1$, such that A-cast (Multi-cast) of each party in \mathcal{T} will be eventually terminated by the (honest) parties in \mathcal{P} . For this, the parties use ACS primitive presented in [BKR94]. The ACS protocol uses n instances of ABA for single bit.

3.1 Multi-valued ABA Protocol

Given the above sub-protocols, our ABA is very simple. Every party P_i on having a message m_i of length ℓ , computes an n length codeword $(s_{i1}, \dots, s_{in}) = \text{ENC}(m_{i0}, \dots, m_{it})$ where $m_{i0} | \dots | m_{it}$. P_i Multi-casts s_{ii} . Using ACS, the parties then agree on some subset of $2t + 1$ parties, say \mathcal{X} whose Multi-casts will be terminated eventually. Every party then verifies if the values Multi-casted by the parties in \mathcal{X} match with their corresponding elements of the codeword and

then A-cast their response. The parties again agree on some subset of $2t + 1$ parties using ACS, say \mathcal{Y} . Based on the responses of the parties in \mathcal{Y} and the values Multi-casted by the parties in \mathcal{X} , the agreement is reached. Note that we use Multi-cast for the elements of the codewords (i.e. s_{ii} 's) and A-cast for the responses. The reason is that s_{ii} 's are message dependent and therefore can be arbitrarily large. Therefore by appropriately setting the value of ℓ , we can implement Multi-casting of s_{ii} values in $\mathcal{O}(n\ell)$ overall complexity. However, we will see from the protocol given below, the response vector will be always n length bit vector. Therefore, using Multi-cast for this case will worsen the complexity, as compared to the case when A-cast of Bracha is used for the same purpose. The protocol is now presented in Figure 3 and its properties are proved subsequently.

Fig. 3. Error-free Communication Optimal ABA.

Protocol Multi-Valued-ABA()

Code for P_i .

1. On having message m_i , divide the ℓ bit message m_i into $t + 1$ blocks, m_{i0}, \dots, m_{it} , each containing $\frac{\ell}{t+1}$ (assume this to be an integer) bits. Compute $(s_{i1}, \dots, s_{in}) = \text{ENC}(m_{i0}, \dots, m_{it})$. Multi-cast s_{ii} .
2. Participate in an instance of ACS to agree on \mathcal{X} containing $2t+1$ parties whose Multi-casts will be eventually terminated by all honest parties.
3. Construct a binary vector V_i of length n . Assign $V_i[j] = 1$, if $P_j \in \mathcal{X}$ and $s_{ij} = s_{jj}$ where s_{jj} is received from the Multi-cast of P_j . Otherwise assign $V_i[j] = 0$. A-cast V_i .
4. Participate in an instance of ACS to agree on \mathcal{Y} containing $2t+1$ parties whose A-casts will be terminated eventually by all honest parties.
5. Check if there are at least $t + 1$ parties in \mathcal{Y} , whose vectors are identical and have at least $t + 1$ 1's. Let $\{i_1, \dots, i_{t+1}\} \subseteq \mathcal{X}$ be the $t + 1$ minimum indices where they all have 1's. If there is no such set of $t + 1$ parties in \mathcal{Y} , then $\{i_1, \dots, i_{t+1}\}$ be the $t + 1$ minimum indices in \mathcal{X} . Then apply DEC on $s_{i_1, i_1}, \dots, s_{i_{t+1}, i_{t+1}}$ and let m_0, m_1, \dots, m_t be the data returned by DEC. Then output $m = m_0 | \dots | m_t$.

Theorem 31 *Protocol Multi-Valued-ABA is an ABA protocol.*

Proof. The termination is guaranteed due to the termination properties of Multi-Valued-Acast, A-cast protocol of Bracha [Bra84] and ACS (the termination of ACS is guaranteed due to the termination of the underlying ABA for single bit). Since Multi-Valued-Acast initiated by the honest parties will eventually terminate, the set \mathcal{X} will be agreed among the parties by the termination of ACS. Similarly, since A-cast (of Bracha) initiated by the honest parties will eventually terminate, the set \mathcal{Y} will be agreed among the parties by the termination of ACS. Once \mathcal{X} and \mathcal{Y} are agreed, the rest is local computation. Therefore termination of Multi-Valued-ABA is guaranteed.

We now argue about the correctness of Multi-Valued-ABA. First we show that all the honest parties will agree on the same message. This follows from

the fact that all the honest parties agree on $\{i_1, \dots, i_{t+1}\} \subseteq \mathcal{X}$ in the last step of the protocol. Furthermore, by the correctness property of **Multi-Valued-Acast**, all the honest parties also agree on the values Multi-casted by the parties in $\{P_{i_1}, \dots, P_{i_{t+1}}\}$. Our claim now follows trivially. We now consider the case when all the honest parties start with same input message m of length ℓ and argue that all honest parties will agree on m eventually. If all the honest parties start with m , then they generate $(s_1, \dots, s_n) = \text{ENC}(m_0, \dots, m_t)$ locally, where $m = m_0 | \dots | m_t$. Every honest party P_i then Multi-casts s_i . By the property of **Multi-Valued-Acast**, all the parties will receive the same value Multi-casted by the parties in \mathcal{X} . Therefore the honest parties in \mathcal{Y} will have identical V_i vectors. Furthermore the V_i vectors will have 1's at least at $t+1$ locations corresponding to the parties in \mathcal{X} who Multi-casted correct value from the codeword (s_1, \dots, s_n) . So $\{i_1, \dots, i_{t+1}\} \subseteq \mathcal{X}$ in the last step of the above protocol will be $t+1$ identities of the parties in \mathcal{X} (having $t+1$ minimum indices) who Multi-casted correct value from codeword (s_1, \dots, s_n) . So DEC when applied on the values Multi-casted by $\{P_{i_1}, \dots, P_{i_{t+1}}\}$ will return m_0, m_1, \dots, m_t where $m = m_0 | \dots | m_t$. \square

Theorem 32 *Multi-Valued-ABA communicates $\mathcal{O}(n\ell)$ bits, invokes $\mathcal{O}(n^3 \log n)$ instances of A-cast for single bit and invokes $2n$ instances of ABA for single bit.*

Proof. Every party Multi-casts $\frac{\ell}{t+1}$ bits. This requires communication of $\mathcal{O}(n\ell)$ bits and $\mathcal{O}(n^3 \log n)$ invocations to A-cast for single bit. Then every party A-casts an n length bit vector. Therefore n^2 invocations to A-cast is required. Finally two invocations to ACS calls for $2n$ instances of ABA for single bit. \square

To make the underlying protocol **Multi-Valued-Acast** work correctly, we require $\frac{\ell}{t+1} \geq (t+1) \log n$. Recall that when the input message size for **Multi-Valued-Acast** is ℓ , then we require that $\ell \geq (t+1) \log n$. In our ABA protocol, the input to **Multi-valued-Acast** is $\frac{\ell}{t+1}$. Therefore we have $\frac{\ell}{t+1} \geq (t+1) \log n \rightarrow \ell \geq (t+1)^2 \log n$.

4 Error-free Communication Optimal BA and BC

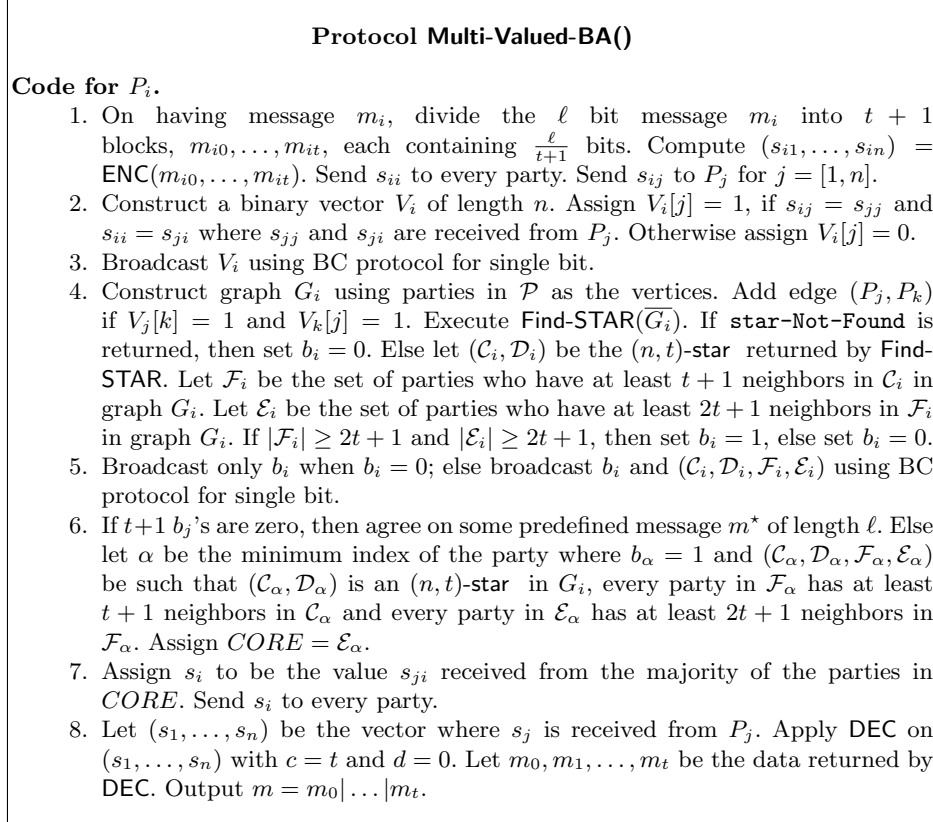
Here we present our new multi-valued BA and BC protocol. We first present a BA protocol. A BC protocol with same complexity of the BA protocol can be achieved by letting the sender send the message to all the parties and then running a BA to reach agreement. This is the standard reduction in synchronous settings from BA to BC [Lyn96]. Our BA protocol follows the idea of **Multi-Valued-Acast**. We use the BC protocol of [BGP09,CW92] for single bit that communicates $\mathcal{O}(n^2)$ bits. We now present the protocol in Figure 4.

Lemma 41 *The honest parties in CORE hold same message of length ℓ .*

The proof of Lemma 41 completely follows from the proof of Lemma 22.

Lemma 42 *If all honest parties start with same input m , then all the parties will agree on CORE, $|CORE| \geq 2t + 1$.*

Fig. 4. Error-free Multi-valued BA with Optimal Communication Complexity.



Proof. The proof here follows from the proof of Lemma 23. Briefly, when all honest parties start with same input, every pair of honest parties will have edge between them. In other words, the edges in the complementary graph will be either (a) between an honest and a corrupted party OR (b) between two corrupted parties. Therefore following the argument given in Lemma 23, \mathcal{C} component of an (n, t) -star will contain at least $t + 1$ honest parties, which subsequently will lead to the construction of \mathcal{F} and \mathcal{E} with size at least $2t + 1$. Although it is not guaranteed that all honest parties find same quadruple $(\mathcal{C}, \mathcal{D}, \mathcal{F}, \mathcal{E})$, but it is ensured that they will find some quadruple. So the honest parties never agree on predefined m^* in this case. Now since all the parties broadcast their quadruple, it is easy to reach agreement on a valid quadruple which the parties do by selecting the one broadcasted by the party with minimum index. Therefore all the parties will agree on CORE . \square

Lemma 43 *If CORE is agreed, all honest parties output the common message of the parties in CORE .*

Proof. By Lemma 41, all honest parties in *CORE* hold same message, say m . The proof now follows from the proof of Lemma 24. \square

Theorem 41 *Multi-Valued-BA is a BA protocol.*

Proof. If *CORE* is agreed, then all honest parties will output the common message of the parties in *CORE* (by Lemma 43). If *CORE* is not agreed, then there must be at least $t + 1$ parties who broadcasted $b_i = 0$. Since b_i 's are broadcasted, all honest parties will agree on predefined m^* of length. So agreement is always achieved at the end. Now if all the honest parties start with same m , then they will agree on *CORE* (by Lemma 42) and output m (by Lemma 43). \square

Theorem 42 *Multi-valued-BA communicates $\mathcal{O}(n\ell)$ bits and invokes $\mathcal{O}(n^2)$ broadcast protocol for single bit.*

Proof. Every party P_i sends two values s_{ii} and s_{ij} to every other party P_j . The values are $\frac{\ell}{t+1}$ bits long each. Therefore in total there are $\frac{\ell}{t+1}\mathcal{O}(n^2) = \mathcal{O}(n\ell)$ bits of communication. Every party P_i broadcasts n -length binary vector V_i , a bit b_i and quadruple $(\mathcal{C}_i, \mathcal{D}_i, \mathcal{F}_i, \mathcal{E}_i)$. Each set in the quadruple can be represented by n -length bit vector. Therefore every party invokes $5n + 1$ instances of broadcast for single bit. This leads to total $\mathcal{O}(n^2)$ instances of broadcast for single bit. \square

The value of ℓ should be at least $(t + 1) \log n$ to make the underlying $(n, t + 1)$ RS code work (following the same logic as explained for Multi-Valued-Acast).

5 Open Problems

An important open question is to investigate whether multi-valued communication optimal protocols can be achieved with less number of invocations to protocols for single bit in comparison to what we provide in this paper.

References

- [ADH08] I. Abraham, D. Dolev, and J. Y. Halpern. An almost-surely terminating polynomial protocol for asynchronous Byzantine Agreement with optimal resilience. In *PODC*, pages 405–414. ACM Press, 2008.
- [BGP09] P. Berman, G. A. Garay, and K. J. Perry. Bit optimal distributed consensus. In *Computer Science Research*, 2009.
- [BKR94] M. BenOr, B. Kelmer, and T. Rabin. Asynchronous secure computations with optimal resilience. In *PODC*, pages 183–192. ACM Press, 1994.
- [BO83] M. Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols. In *PODC*, pages 27–30. ACM Press, 1983.
- [BOCG93] M. Ben-Or, R. Canetti, and O. Goldreich. Asynchronous Secure Computation. In *STOC*, pages 52–61. ACM Press, 1993.
- [BOGW88] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *STOC*, pages 1–10. ACM Press, 1988.

- [Bra84] G. Bracha. An asynchronous $\lfloor (n-1)/3 \rfloor$ -resilient consensus protocol. In *PODC*, pages 154–162. ACM Press, 1984.
- [BTH06] Z. Beerliová-Trubíniová and M. Hirt. Efficient multi-party computation with dispute control. In *TCC*, LNCS 3876, pages 305–328, 2006.
- [Can95] R. Canetti. *Studies in Secure Multiparty Computation and Applications*. PhD thesis, Weizmann Institute, Israel, 1995.
- [CGMA85] B. Chor, S. Goldwasser, S. Micali, and B. Awerbuch. Verifiable secret sharing and achieving simultaneity in the presence of faults (extended abstract). In *STOC*, pages 383–395. ACM Press, 1985.
- [CR93] R. Canetti and T. Rabin. Fast asynchronous Byzantine Agreement with optimal resilience. In *STOC*, pages 42–51. ACM Press, 1993.
- [CW92] B. A. Coan and J. L. Welch. Modular construction of a Byzantine Agreement protocol with optimal message bit complexity. *Information and Computation*, 97(1):61–85, 1992.
- [DR85] D. Dolev and R. Reischuk. Bounds on information exchange for Byzantine Agreement. *JACM*, 32(1):191–204, 1985.
- [FH06] M. Fitzi and M. Hirt. Optimally Efficient Multi-valued Byzantine Agreement. In *PODC*, pages 163–168, 2006.
- [FLP85] M. J. Fischer, N. A. Lynch, and M. Paterson. Impossibility of distributed consensus with one faulty process. *JACM*, 32(2):374–382, 1985.
- [FM88] P. Feldman and S. Micali. An Optimal Algorithm for Synchronous Byzantine Agreement. In *STOC*, pages 639–648. ACM Press, 1988.
- [GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [HMP00] M. Hirt, U. Maurer, and B. Przydatek. Efficient Secure Multiparty Computation. In *ASIACRYPT*, LNCS 1976, pages 143–161, 2000.
- [LV11] G. Liang and N. H. Vaidya. Error-Free Multi-Valued Consensus with Byzantine Failures. In *PODC*, pages 11–20. ACM Press, 2011.
- [Lyn96] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [PR10] A. Patra and C. Pandu Rangan. Communication Optimal Multi-valued Asynchronous Broadcast Protocol. In *LATINCRYPT*, LNCS 6212, pages 162–177, 2010.
- [PR11] A. Patra and C. Pandu Rangan. Communication Optimal Multi-valued Asynchronous Byzantine Agreement with Optimal Resilience. In *ICITS*, LNCS 6673, pages 206–226, 2011.
- [PSL80] M. Pease, R. E. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *JACM*, 27(2):228–234, 1980.
- [PW92] B. Pfitzmann and M. Waidner. Unconditional Byzantine Agreement for any number of faulty processors. In *STACS*, LNCS 577, pages 339–350, 1992.
- [Rab83] M. O. Rabin. Randomized Byzantine generals. In *FOCS*, pages 403–409. IEEE Computer Society, 1983.
- [RBO89] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In *STOC*, pages 73–85. ACM Press, 1989.
- [TC84] R. Turpin and B. A. Coan. Extending binary Byzantine Agreement to multivalued Byzantine Agreement. *IPL*, 18(2):73–76, 1984.