

Constant-Round Asynchronous Multi-Party Computation Based on One-Way Functions^{*}

Sandro Coretti^{1**}, Juan Garay^{2***}, Martin Hirt³, and Vassilis
Zikas^{4***†}

¹ New York University
corettis@nyu.edu

² Yahoo Research

garay@yahoo-inc.com

³ Department of Computer Science, ETH Zurich
hirt@inf.ethz.ch

⁴ Department of Computer Science, RPI
vzikas@cs.rpi.edu

Abstract. Secure multi-party computation (MPC) allows several mutually distrustful parties to securely compute a joint function of their inputs and exists in two main variants: In *synchronous* MPC parties are connected by a synchronous network with a global clock, and protocols proceed in *rounds* with strong delivery guarantees, whereas *asynchronous* MPC protocols can be deployed even in networks that deliver messages in an arbitrary order and impose arbitrary delays on them.

The two models—synchronous and asynchronous—have to a large extent developed in parallel with results on both feasibility and asymptotic efficiency improvements in either track. The most notable gap in this parallel development is with respect to round complexity. In particular, although under standard assumptions on a synchronous communication network (availability of secure channels and broadcast), synchronous MPC protocols with (exact) constant rounds have been constructed, to the best of our knowledge, thus far no constant-round asynchronous MPC protocols based on standard assumptions are known, with the best protocols requiring a number of rounds that is linear in the multiplicative depth of the arithmetic circuit computing the desired function.

In this work we close this gap by providing the first constant-round asynchronous MPC protocol that is optimally resilient (i.e., it tolerates up to $t < n/3$ corrupted parties), adaptively secure, and makes black-box use of a pseudo-random function. It works under the standard network

^{*} The full version of this paper can be found on the Cryptology ePrint Archive [18].

^{**} Work supported by the Swiss NSF project no. 200020-132794.

^{***} Work done in part while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant #CNS-1523467.

[†] Work supported in part by the Swiss NSF Ambizione grant PZ00P2_142549.

assumptions for protocols in the asynchronous MPC setting, namely, a complete network of point-to-point (secure) asynchronous channels with eventual delivery and asynchronous Byzantine agreement (aka consensus). We provide formal definitions of these primitives and a proof of security in the Universal Composability framework.

1 Introduction

In secure multi-party computation (MPC), a set of n parties p_1, \dots, p_n , each holding some private input, wish to jointly compute a function on these inputs in a fashion such that even up to t colluding adversarial parties are unable to obtain any information beyond what they can extract from their inputs and outputs or to affect the computation in any way other than contributing their desired inputs. The problem of MPC has been studied in the two important settings of *synchronous* and *asynchronous* networks, respectively.

MPC protocols for the synchronous setting assume a network in which parties proceed in rounds, with the guarantee that messages sent by any party in any given round are delivered to all recipients in the same round. Consequently, in all such protocols the parties are assumed to be (at least partially) synchronized, i.e., to be in the same round at all times.

In real-world networks, such as the Internet, this synchrony assumption corresponds to assuming that the parties have (partially) synchronized clocks and communicate over channels with a known upper-bounded latency. The synchronous structure is then imposed by “timeouts,” i.e., in each round the parties wait for an amount T of time defined by their estimate of when other parties send their messages and the bound on the network latency. If their estimate is accurate and their clocks are indeed synchronized, this will ensure that parties receive all messages sent to them from honest senders before the end of the round (timeout). Thus, after time T has passed, they can safely assume that if a message was expected for the current round but has not been received, then the sender must be adversarial. The security of synchronous protocols heavily relies on this assumption. In fact, many of them would become completely insecure if there is even a single delayed message. As a result, the round length T must typically be set much higher than the average transmission time.

A natural question is therefore to study the security one can obtain if no synchrony assumption is made but merely under the assumption that

messages sent by honest parties are *eventually* delivered.⁵ In particular, messages sent by parties can be reordered arbitrarily and delayed by arbitrary (albeit finite) amounts of time in such an asynchronous network. Note that one could consider even more pessimistic networks where the adversary can block messages sent by honest parties; this is for example the case in the base network assumed in Canetti’s UC framework [13]. In such networks, however, protocols cannot be guaranteed to (eventually) terminate as the adversary can delay the computation indefinitely.

In asynchronous MPC protocols parties do not wait until a round times out. Rather, as soon as a party has received enough messages to compute its next message⁶, it computes that message, sends it, and moves on. In that sense, asynchronous MPC protocols are “opportunistic” and terminate essentially as quickly as the network allows. Hence, they can be much faster than their synchronous counterparts depending on the network latency.⁷

In this work, unless explicitly stated otherwise, whenever we refer to the *asynchronous (communication) model*, we mean the above asynchronous model with eventual delivery.

On the importance of round complexity. The inherent need for waiting until each round times out clearly makes round complexity an important consideration for the performance of *synchronous* MPC protocols. Indicatively, Schneider and Zohner [34] have shown that as the latency between machines increases, the cost of each round becomes more and more significant.

Despite their opportunistic nature, round complexity is just as important a consideration for *asynchronous* protocols, since a protocol’s round complexity can be a more relevant efficiency metric than, for example, its bit complexity. Indeed, at the conceptual/theoretical level, having constant-round protocols allows us to use them as sub-routines in a higher level protocol without blowing up (asymptotically) the round complexity of the calling protocol, while at the practical level, communication time is often dominated by the round-trip times in the network and not by the size of the messages. For example, it takes about the same amount of time to transmit a byte and a megabyte, while sending a message from

⁵ The eventual-delivery assumption is supported by the fact that whenever a message is dropped or delayed for too long, Internet protocols typically resend that message.

⁶ What “enough” means is concretely specified by the party’s protocol.

⁷ This speed up, however, does not come for free, as it inevitably allows the adversary to exclude some of the honest parties’ inputs from being considered in the computation.

A to B over many intermediate nodes, computing something at B , and sending an answer back to A may take a (comparatively) long time.

Our contributions. In this paper, we first formalize the asynchronous model with *eventual delivery* in the universal composability (UC) framework [13], introduce a suitable formal notion of *asynchronous* round complexity, and formulate the basic communication resources (such as asynchronous secure channel and asynchronous Byzantine agreement [A-BA]) as ideal functionalities in that model.⁸ (See Section 3.) We then present the—to the best of our knowledge—first *constant-round* MPC protocol for this asynchronous setting (i.e., a protocol whose round complexity is independent of the multiplicative depth of the evaluated circuit and the number n of parties) based on standard assumptions, namely, the existence of pseudo-random functions (PRFs).⁹ The protocol is UC-secure in the secure-channels model with A-BA, and makes *black-box* use of the underlying PRF, tolerating a computationally bounded, *adaptive* adversary actively corrupting up to $t < n/3$ parties, which is optimal for this setting.¹⁰

At a high level, here is how we construct our constant round protocol. First, we devise a constant-depth circuit for computing the keys, masked values, and (shares of the) garbled gates needed for a distributed evaluation of a Yao garbled circuit that encodes the function the parties wish to compute. This circuit is then evaluated by means of a linear-round (in the depth of the circuit and in n) asynchronous protocol. However, this circuit is Boolean whereas all existing asynchronous protocols evaluate arithmetic circuits. To deal with this mismatch we devise an asynchronous protocol for computing Boolean circuits by appropriately adapting the protocol by Ben-Or, Kelmer, and Rabin [8]. Any party who receives the output from the evaluation of the Boolean circuit uses it to encrypt shares of each garbled gate, which it sends to all other parties. Finally, each party *locally* evaluates the (distributed) garbled circuit by decrypting incoming encrypted shares of each gate and reconstructing the function table of the gate as soon as sufficiently many consistent shares

⁸ Note that while the UC framework already is asynchronous, asynchronous communication *with eventual delivery* has not been modeled in it so far; moreover, standard (asynchronous) UC protocols do not achieve eventual termination/delivery (cf. Section 3).

⁹ An approach based on threshold fully homomorphic encryption was recently proposed by Cohen [17]; see the discussion in the related work section below.

¹⁰ The necessity of this bound in the asynchronous setting is also discussed in related work below.

have arrived until all gates are evaluated. Once all gates are evaluated in this fashion, the party is in possession of the output. The protocol and its analysis are presented in Section 4.

Related work. Beaver, Micali, and Rogaway [2] were the first to provide a constant-round MPC protocol in the synchronous stand-alone model. (Refer to Appendix A for a more detailed and historical account of the development of MPC protocols in both the synchronous and asynchronous settings, together with the tools that are used in each setting.) Their protocol is secure in the computational setting and tolerates an adaptive adversary who actively corrupts up to $t < n/2$ parties. The complexity of [2] was improved by Damgård and Ishai [19], who provided the first constant-round protocol making black-box use of the underlying cryptographic primitive (a pseudo-random generator). Importantly, both [2] and [19] assume a broadcast channel, an assumption essential for obtaining constant-round MPC. Indeed, as proved in [22,20], it is impossible to implement such a broadcast channel from point-to-point communication in a constant number of rounds, and although expected constant-round broadcast protocols exist in the literature (e.g., [21,30]), using them to instantiate calls within the constructions of [2] or [19] would not yield an expected constant-round protocol [6]. The intuitive reason—formally argued by Ben-Or and El-Yaniv [6]—is that the process of running n such broadcast protocols (even in parallel) does not terminate in an expected constant number of rounds.

The model of *asynchronous* communication with eventual delivery was considered early on in seminal works on fault-tolerant distributed computing (e.g., [23]), although to our knowledge this paper is the first to formalize this capability in the UC framework. The study of optimally resilient MPC in this type of asynchronous networks was initiated by Ben-Or, Canetti, and Goldreich [5], who proved that any function can be computed by a perfectly secure asynchronous protocol if and only if at most $t < n/4$ parties are corrupted. Following that result, Ben-Or, Kelmer, and Rabin [8] showed that if a negligible error probability is allowed, the bound $t < n/3$ is necessary and sufficient for asynchronous MPC.¹¹ More recently, Hirt *et al.* [27,28] provided computationally secure

¹¹ The necessity of the $t < n/3$ bound follows from the result by Canetti *et al.* [5,12], who argue that this bound is necessary for fail-stop adversaries; it also applies to computational security and assuming A-BA. Moreover, note that in the asynchronous setting, all feasibility bounds are worse by an additive term of t compared to the synchronous setting. Intuitively, this stems from the fact that honest parties cannot distinguish between messages by other honest parties being delayed and

solutions (i.e., protocols tolerating a computationally bounded adversary) and Beerliová and Hirt [3] perfectly secure solutions with improved communication complexity.

The above asynchronous protocols are secure—according to simpler, stand-alone security definitions—if one assumes point-to-point communication and an A-BA protocol. Similarly to their synchronous counterparts, all the above protocols—even assuming an A-BA primitive—have round complexity linear in the multiplicative depth of the arithmetic circuit that computes the function, as they follow the standard gate-by-gate evaluation paradigm.

Concurrently and independently, Cohen [17] recently put forth an asynchronous MPC protocol that is secure against a computationally bounded attacker statically corrupting up to $t < n/3$ parties and in which all parties run in constant time. Notably, the protocol from [17] relies on fully homomorphic encryption, thus leaving the question of constant-round MPC from standard assumptions open, which is answered in this work.

We note in passing that although in the synchronous setting BA implies broadcast, this is not the case in the asynchronous setting. Indeed, Canetti and Rabin [15] provide an asynchronous BA protocol tolerating $t < n/3$ malicious parties, which if every honest party terminates at the latest after a poly-logarithmic number of rounds, securely implements asynchronous BA except with negligible probability. A broadcast protocol with similar guarantees is provably impossible [23], and existence of an asynchronous BA protocol which terminates in a strict constant number of rounds would contradict the impossibility from [22,20]. Similarly to the synchronous case, although solutions for asynchronous BA with expected constant number of rounds exist [11,15], using them in the above asynchronous protocol to replace invocations to asynchronous BA would not yield an expected constant-round MPC protocol [6].¹²

2 Model and Building Blocks

We denote the player set by $\mathcal{P} = \{p_1, \dots, p_n\}$ and consider a computationally bounded *adaptive t -adversary*, i.e., the adversary gets to corrupt

messages by corrupted parties not being sent. Thus, in particular, perfectly secure asynchronous MPC is possible only if $t < n/4$.

¹² Nonetheless, [6] does describe an alternative way of obtaining several asynchronous BA protocols that are guaranteed to all terminate in expected constant number of rounds.

up to t parties dynamically during the protocol execution and depending on its protocol view. The most common network model for the execution of asynchronous protocols is secure channels with eventual delivery, where the adversary is allowed to delay the delivery of any message by an arbitrary but finite amount of time, i.e., he is not able to block messages sent among honest parties. Moreover, as argued in the introduction, existing asynchronous protocols rely on an additional resource, namely, an asynchronous version of Byzantine agreement (A-BA) instead of a broadcast channel, and such a resource is even necessary to obtain an (exact) constant number of rounds. We formalize this model and formulate the ideal functionalities corresponding to these communication resources separately in Section 3.

We now present some basic tools we use in our protocol.

Secret sharing. Our construction makes use of Shamir’s secret sharing scheme [35], which allows to encode a secret into n shares such that any subset of t shares gives no information about the secret and such that from any subset of $t + 1$ shares the secret can be reconstructed.

For a sharing of a secret s , let $[s]_i$ denote the i^{th} share. A set of shares are called *t-consistent* if they lie on a polynomial of degree at most t . For a tuple of secrets $s = (s_1, \dots, s_\ell)$, denote—in slight abuse of notation—by $[s]_i := ([s_1]_i, \dots, [s_\ell]_i)$ the tuple of the i^{th} shares of the values and refer to it as the i^{th} share of s . A set of such tuples is called *t-consistent* if the property holds for all components.

A linear-round asynchronous MPC protocol. In [8], Ben-Or, Kelmer, and Rabin constructed a protocol, call it π_{BKR} , that computes an arbitrary n -party function f in an asynchronous environment assuming asynchronous point-to-point (secure) channels and asynchronous BA.¹³ The protocol follows the gate-by-gate evaluation paradigm [7,26,16], where the function to be evaluated is represented as an arithmetic circuit over a sufficiently large finite field, and the computation proceeds by evaluating sequentially the gates of depth one, then the gates of depth two, and so on. The evaluation of each gate requires a constant number of (asynchronous) rounds,¹⁴ thus making the round complexity of the overall protocol linear in the depth of the circuit.

¹³ [8] also assumes A-Cast to get a more efficient solution, but as argued in the introduction, A-Cast can be easily reduced to asynchronous BA in two rounds.

¹⁴ Note that in each such round the parties might invoke the asynchronous BA resource.

π_{BKR} was designed for a simpler, stand-alone security definition, which only ensures sequential composition. In the next section we show how it can be cast in our eventual-delivery model to give UC-security guarantees.

3 A UC Model for Asynchronous Computation with Eventual Message Delivery

In this section we formalize the asynchronous network model *with eventual message delivery* in the UC framework. We start with the basic task of point-to-point communication and proceed with asynchronous SFE and BA. Note that the asynchronous model with eventual delivery has previously been informally treated only in the stand-alone model without composition. Thus, although at first read one might consider our treatment pedantic, providing a UC proof of asynchronous MPC protocol with eventual termination/delivery requires the design of appropriate UC functionalities that can be used as hybrids. Indeed, while the plain UC framework is inherently asynchronous, the adversary has full control over message delivery and may even choose to delete messages sent between uncorrupted parties (by delaying them indefinitely). Hence, without the extensions in this section, the UC model does *not* capture eventual delivery.¹⁵

Asynchronous communication with eventual delivery. Our formulation of communication with eventual delivery within the UC framework builds on ideas from [31]. In particular, we capture such communication by allowing the parties access to multi-use bilateral secure channels, where a sender $p_i \in \mathcal{P}$ can input a messages to be delivered to some recipient $p_j \in \mathcal{P}$; messages are buffered and delivered in an order specified by the adversary.

To ensure that when p_s and p_r are honest, the adversary cannot delay the delivery of submitted messages arbitrarily, we make the following modifications: We first turn the UC secure channels functionality to work in a “fetch message” mode, where the channel delivers the message to its intended recipient p_j if and only if p_j asks to receive it by issuing a special “fetch” command. If the adversary wishes to delay the delivery of

¹⁵ Standard UC constant-round protocols in the plain (fully asynchronous) UC framework do not work in this setting as, in these protocols, a party waits for all his r -round messages before proceeding to round $r + 1$, which would allow the adversary to make honest parties wait indefinitely (for messages from corrupted parties), thereby preventing them from terminating. Instead, in asynchronous protocols with eventual delivery, parties need to proceed to the next round as soon as they have sufficiently many (but not necessarily all) their messages for the current round.

some message, he needs to submit to the channel functionality an integer value T —the *delay*. This will have the effect of the channel ignoring the first T fetch attempts following the reception of the sender’s message. Importantly, we require the adversary send the delay T in unary notation; this will ensure that the delay will be bounded by the adversary’s running time,¹⁶ and thus a polynomial environment will be able to observe the transmission through its completion. To allow the adversary freedom in scheduling delivery of messages, we allow him to input delays more than once, which are added to the current delay amount. (If the adversary wants to deliver the message in the next activation, all he needs to do is submit a negative delay.)

The detailed specification of secure channels with eventual delivery, denoted $\mathcal{F}_{\text{A-SMT}}$, is given in Figure 1. In the description, we denote by \mathbf{M} a vector of strings. We also use $\|$ to denote the operation which adds a string to \mathbf{M} ; concretely, if $\mathbf{M} = (m_1, \dots, m_\ell)$, then $\mathbf{M}\|m := (m_1, \dots, m_\ell, m)$ and $m\|\mathbf{M} = (m, m_1, \dots, m_\ell)$.

We refer to the model in which every two parties p_i and p_j in \mathcal{P} have access to an independent instance of $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$ as the $\mathcal{F}_{\text{A-SMT}}$ -*hybrid model*. An asynchronous protocol in such a model proceeds as follows: Whenever a party p_j gets activated, if its current protocol instructions include sending some message m to some other party p_j , then the party inputs **(send, m)** to $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$; otherwise, p_i sends a **fetch** message to every channel $\mathcal{F}_{\text{A-SMT}}(p_j, p_i)$, $j \in [n]$ in a round-robin fashion, i.e., if in the previous activation it sent a **(fetch)** message to $\mathcal{F}_{\text{A-SMT}}(p_j, p_i)$, then it sends a **(fetch)** message to $\mathcal{F}_{\text{A-SMT}}(p_{(j \bmod n)+1}, p_i)$.

Remark 1 (On permuting messages). Our formulation of an asynchronous channel captures the somewhat pessimistic view of asynchronous communication, implicit in many works on asynchronous distributed protocols, where the adversary has full scheduling power and can, in particular, reorder the messages sent by any party as he wishes. One could attempt to emulate a network which does not allow for reordering of the messages—the so-called *first-in-first-out (FIFO)* channel—by adding appropriate (publicly known) message identifiers and instructing the parties to wait until a specific identifier is delivered before outputting messages with other identifiers. However, we note that such an emulating protocol would be distinguished from the original when, for example, we consider an adversary that introduces no delay and an environment that inputs two

¹⁶ We refer to [13] for a formal definition of running time in the UC framework.

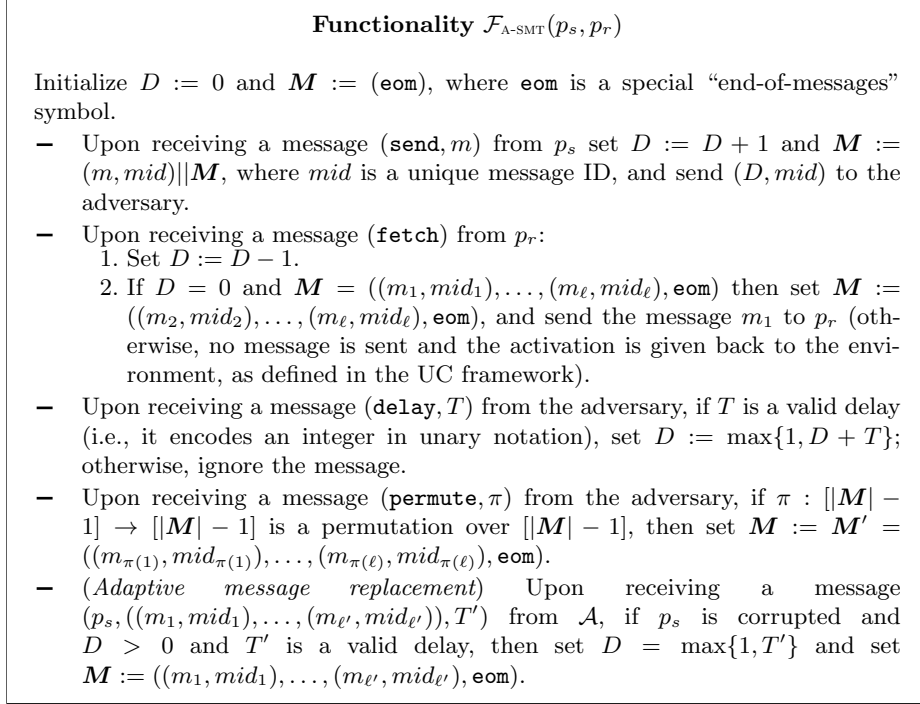


Fig. 1. Asynchronous secure channel with eventual delivery

messages in a row and corrupts the receiver as soon as the first message is supposed to have been delivered.

Asynchronous secure function evaluation (SFE). In an asynchronous environment, it is impossible to get guaranteed (eventual) termination and input completeness, i.e., take into account all inputs in the computation of the function (cf. [31] and early work on fault-tolerant distributed computing). The reason is that if honest parties wait until the inputs of all parties are delivered, then the adversary can make them wait indefinitely for corrupted parties to give their inputs (honest parties have no way of distinguishing between an honest sender whose message is delayed and a corrupt sender who did not send a message). Thus, to ensure eventual termination, the parties cannot afford to wait for input from more than $n - t$ parties, as the t remaining parties might be the corrupted ones. Therefore, protocols for asynchronous computation of a multi-party function f on inputs x_1, \dots, x_n from parties p_1, \dots, p_n end up computing the

Functionality $\mathcal{F}_{\text{A-SFE}}^f(\mathcal{P})$

$\mathcal{F}_{\text{A-SFE}}^f$ proceeds as follows, given a function $f : (\{0, 1\}^* \cup \{\perp\})^n \times R \rightarrow (\{0, 1\}^*)^n$ and a player set \mathcal{P} . For each $i \in \mathcal{P}$, initialize variables x_i and y_i to a default value \perp and a current delay $D_i := 0$. Additionally, set $\mathcal{I} := \mathcal{H}$. (Recall that \mathcal{H} denotes the set of honest parties)

- Upon receiving message (**no-input**, \mathcal{P}') from the adversary, if $|\mathcal{P}'| \leq |\mathcal{P} \setminus \mathcal{H}|$ and no party has received an output (**output**, y) yet, then set $\mathcal{I} = \mathcal{H} \setminus \mathcal{P}'$; otherwise ignore this message.
- Upon receiving input (**input**, v) from party $p_i \in \mathcal{P}$ (or from the adversary in case p_i is corrupted), do the following:
 - If some party (or the adversary) has received an output (**output**, y), then ignore this message; otherwise, set $x_i := v$.
 - If $x_i \neq \perp$ for every $p_i \in \mathcal{I}$, then compute $(y_1, \dots, y_n) = f((x'_1, \dots, x'_n), r)$ for a uniformly random r , where $x'_i = x_i$ for $p_i \in \mathcal{I} \cup (\mathcal{P} \setminus \mathcal{H})$ and $x'_i = \perp$ for all other parties.
 - Send (**input**, i) to the adversary.
- Upon receiving (**delay**, p_i, T), from the adversary, set $D_i := D_i + T$.
- Upon receiving message (**fetch**) from party $p_i \in \mathcal{P}$, if y_i has not yet been set (i.e., $y_i = \perp$) then ignore this message, else do:
 - Set $D_i := D_i - 1$
 - If $D_i = 0$, send (**output**, y_i) to p_i .

Fig. 2. Asynchronous SFE with eventual delivery

function $f|_{\mathcal{P}'}(x_1, \dots, x_n) = f(x'_1, \dots, x'_n)$ for some $\mathcal{P}' \subseteq \mathcal{P}$ with $|\mathcal{P}'| = t$, where $x'_i = x_i$ if $p_i \notin \mathcal{P}'$, and otherwise a default value (denoted \perp).

Moreover, by being able to schedule the delivery of messages from honest parties, the adversary can (in worst-case scenarios) choose exactly the set \mathcal{P}' . Therefore, the ideal functionality corresponding to asynchronous SFE with eventual termination needs to allow the simulator to choose this set depending on the adversary’s strategy. Moreover, the simulator should be allowed to schedule delivery of the outputs depending on the adversary’s strategy, but not allowed to delay them arbitrarily. This last requirement can be achieved, as in the case of $\mathcal{F}_{\text{A-SMT}}$, by turning the SFE functionality into a “fetch message”-mode functionality and allowing the simulator to specify a delay on the delivery to every party.

The SFE functionality with the above properties is described in Figure 2. In the description, we use $\mathcal{H} \subseteq \mathcal{P}$ to denote the set of honest parties; note that \mathcal{H} is dynamically updated as the adversary corrupts new parties. Moreover, we use \mathcal{I} to denote the set of honest parties whose input is allowed to be considered in the computation, and require that

$|\mathcal{I}| \geq n - 2t$. We provide a generic description of the functionality for an arbitrary number t of corruptions; however, and as implied by classical impossibility results, we are only able to realize it for $t < n/3$ [5].

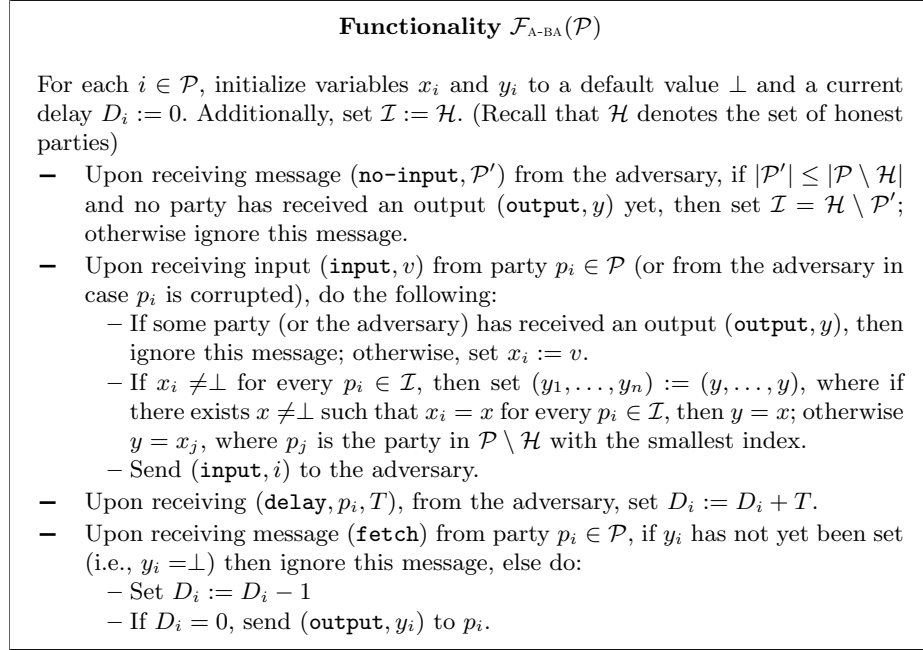


Fig. 3. Asynchronous BA with eventual delivery

Asynchronous BA with eventual delivery. The last primitive we describe is (UC) asynchronous BA with eventual message delivery. In such a BA primitive, every party has an input, and we want to ensure the following properties: All honest parties (eventually) output the same value y (consistency), and if all honest parties have the same input x , then this output is $y = x$. Intuitively, asynchronous BA can be cast as a version of asynchronous SFE for the function that looks at the set of received inputs and, if all inputs contributed by honest parties are identical¹⁷ and equal to some x , sets the output to x for every party; otherwise, it sets the output to the input of some corrupted party (for example, the first in

¹⁷ Similarly to the SFE case, the adversary might prevent some of the honest parties from providing an input.

any ordering, e.g., lexicographic). The formal definition of $\mathcal{F}_{\text{A-BA}}$ is given in Figure 3.

We will refer to the setting where every two parties p_i and p_j in \mathcal{P} have access to an independent instance of $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$ and, additionally, the parties have access to independent instances of $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$ as the $\{\mathcal{F}_{\text{A-SMT}}, \mathcal{F}_{\text{A-BA}}\}$ -*hybrid model*. The execution in the $\{\mathcal{F}_{\text{A-SMT}}, \mathcal{F}_{\text{A-BA}}\}$ -hybrid model is analogous to the execution in the $\mathcal{F}_{\text{A-SMT}}$ -hybrid model: Whenever a party p_i gets activated, if its current protocol instructions include sending some message m to some other party p_j or inputting a message m' to $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$, then the party inputs (send, m) to $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$ or m' to $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$, respectively; otherwise, p_i keeps sending (with each activation) a **fetch** to every channel $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$, $j \in [n]$ and then to $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$ in a round-robin fashion.

Asynchronous rounds. We now briefly elaborate on the notion of rounds in an asynchronous environment. Unlike the situation in the synchronous case, where rounds are well specified by the protocol, the definition of rounds in an asynchronous setting requires a bit more care. Intuitively, two messages m_i and m_j sent by some party p_i in an asynchronous protocol are considered to be sent in rounds i and j , $j > i$, if m_j is generated by computation which takes as input a message received after p_i sent m_i . Following the above intuition, we define for each p_i and for each point in the protocol execution, the current round in which p_i is to be the number of times p_i alternated between sending (send, m) to some channel $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$, $p_j \in \mathcal{P}$ (or to the asynchronous BA functionality $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$) and sending **fetch** to some $\mathcal{F}_{\text{A-SMT}}(p_k, p_i)$, $p_k \in \mathcal{P}$ or to $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$. That is, every round (except for the first one)¹⁸ starts by sending a (send, m) to some $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$ or to $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$ after some **fetch** was sent by p_i and finishes with the first **fetch** command that p_i sends. The round complexity of the protocol is the maximum (over all honest parties) number of rounds that an honest party uses in the protocol execution.

We note in passing that, similarly to [31], the above formulation allows for any party to send several messages in each round: the party buffers the messages and while the buffer is not empty, in each activation the party pops the next message and sends it to its intended recipient.

A UC-secure linear-round MPC protocol with eventual delivery. Finally, we argue the security of protocol π_{BKR} mentioned in Section 2 in our model. π_{BKR} is information-theoretic and proved simulation-based secure,

¹⁸ The first round starts as soon as the party receives its protocol input from the environment.

where the simulation is in fact black-box (i.e., the simulator uses the corresponding adversary in a black-box manner) and straight-line (the simulator does not rewind the adversary). Moreover, the protocol tolerates any adaptive t -adversary where $t < n/3$, a bound which is also tight [5]. Thus, by casting π_{BKR} in our UC $\{\mathcal{F}_{\text{A-SMT}}, \mathcal{F}_{\text{A-BA}}\}$ -hybrid model—where every bilateral message exchange is implemented by the sender p_i and the receiver p_j using (an instance of the) channel $\mathcal{F}_{\text{A-SMT}}(p_i, p_j)$ and every call to asynchronous BA done by invocation of $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$ —we obtain a protocol for UC securely evaluating $\mathcal{F}_{\text{A-SFE}}^f(\mathcal{P})$, which is linear in the depth of the circuit computing f . More formally:

Theorem 1 ([8]). *Let f be an n -ary function and C be an arithmetic circuit for computing f by parties in \mathcal{P} . Then there exists a protocol, π_{BKR} , which UC-securely realizes $\mathcal{F}_{\text{A-SFE}}^f$ in the $\{\mathcal{F}_{\text{A-SMT}}, \mathcal{F}_{\text{A-BA}}\}$ -hybrid model tolerating an adaptive t -adversary in a linear (in the depth of the circuit) number of rounds, provided $t < n/3$.*

4 Constant-Round Asynchronous SFE

In this section we present our asynchronous SFE protocol and prove its security and round complexity.

4.1 Description of the Protocol

Let Circ be a Boolean circuit that is to be evaluated in a multi-party computation. In our protocol for securely evaluating the function that Circ computes, denoted $\pi_{\text{A-SFE}}(\text{Circ}, \mathcal{P})$, the parties first jointly compute a garbled version of Circ (along the lines of [37,2,4]); every party then evaluates this garbled circuit locally to obtain the output of the computation. Computing the garbled circuit takes place in two phases: First, the parties evaluate a function $f_{\text{PREP}}^{\text{Circ}}$ (described below) which is represented by a *constant-depth* arithmetic circuit over a finite field using a (non-constant-round) asynchronous MPC protocol. Given the outputs of this function, the parties can then complete the computation of the garbled circuit within one additional asynchronous round.¹⁹ Since the evaluation of the garbled circuit takes place locally and $f_{\text{PREP}}^{\text{Circ}}$ is computed via a constant-depth circuit, the entire protocol is a constant-round protocol.

We define and analyze our protocol in the $\{\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}, \mathcal{F}_{\text{A-SMT}}\}$ -hybrid model. Furthermore we provide a protocol for UC securely realizing the

¹⁹ Refer to Section 2 for a definition of asynchronous round complexity.

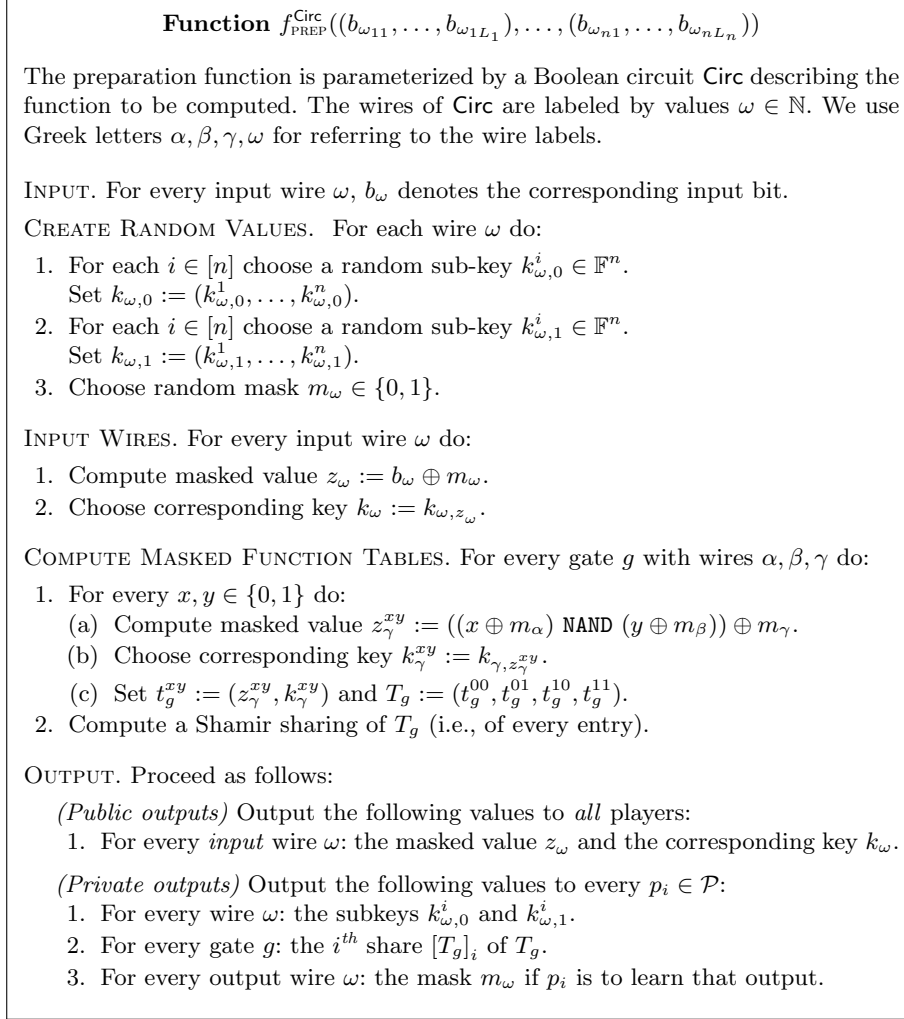


Fig. 4. The description of function $f_{\text{PREP}}^{\text{Circ}}$ corresponding to the distributed version of circuit garbling.

functionality $\mathcal{F}_{\text{A-SFE}}^{\text{Circ}}$ from asynchronous secure channels and BA with eventual delivery based on π_{BKR} [8] (cf. Lemma 1).

Circuit garbling. Before elaborating on the protocol, we describe what the garbled version of Circ looks like.²⁰ Boolean circuit Circ consists of

²⁰ Note that $f_{\text{PREP}}^{\text{Circ}}$ actually computes a “distributed” version of the garbled circuit (described below).

wires and NAND gates.²¹ In the garbled version, every wire ω of Circ has a corresponding (secret) random mask m_ω , which is used to hide the real value on that wire. Consequently, every gate g , with input wires α and β and output wire γ , has a special function table T_g that works on masked values. It contains four entries z_γ^{xy} , corresponding to the masked value on the outgoing wire γ under the four possible combinations of masked inputs $x, y \in \{0, 1\}$ on wires α and β . Each entry is obtained by unmasking the inputs, applying the gate function (NAND), and re-masking the result with the mask of the outgoing wire. That is,

$$z_\gamma^{xy} = ((x \oplus m_\alpha) \text{ NAND } (y \oplus m_\beta)) \oplus m_\gamma,$$

for $x, y \in \{0, 1\}$.

The entries of each function table need to be protected so that only the one entry necessary to evaluate the circuit can be accessed. To that end, for each wire ω there are two (secret) keys $k_{\omega,0}$ and $k_{\omega,1}$. In the function tables T_g , each entry z_γ^{xy} is now augmented by the corresponding key $k_{\gamma, z_\gamma^{xy}}$ of the outgoing wire γ . The pair $t_g^{xy} := (z_\gamma^{xy}, k_{\gamma, z_\gamma^{xy}})$ is encrypted with $k_{\alpha,x}$ and $k_{\beta,y}$ under the “tweak” (g, x, y) . The resulting ciphertexts

$$c_g^{xy} := \text{Enc}_{k_{\alpha,x}, k_{\beta,y}}^{g,x,y} (t_g^{xy}) = \text{Enc}_{k_{\alpha,x}, k_{\beta,y}}^{g,x,y} (z_\gamma^{xy}, k_{\gamma, z_\gamma^{xy}})$$

make up the garbled function table

$$C_g := (c_g^{00}, c_g^{01}, c_g^{10}, c_g^{11}),$$

where $(\text{Enc}_{k_1, k_2}^T(\cdot), \text{Dec}_{k_1, k_2}^T(\cdot))$ is a tweakable dual-key cipher. A suitable such cipher can be realized using a PRF [4].²²

In order to be compatible with the garbled function tables, inputs to the circuit must be garbled as well. That is, for the input bit b_ω on input wire ω , the garbled input is $z_\omega := b_\omega \oplus m_\omega$ and the corresponding key is $k_\omega := k_{\omega, z_\omega}$.

With the garbled inputs and function tables, any party can (locally) evaluate the circuit as follows: Given the masked values and the corresponding keys of the incoming wires of some gate, the party can decrypt the corresponding row, obtaining the masked value on the outgoing wire and the corresponding key. In the end, the values on the output wires can be unblinded if the corresponding masks are known.

²¹ Any (arithmetic or Boolean) circuit can be efficiently transformed into such a circuit.

²² The security required from such a cipher is roughly semantic security even if one of the keys is known (see [4] for more details). Moreover, we assume a canonical injective mapping of triples (g, x, y) to the tweak space of the cipher.

Distributed encryption. Given the input bits b_ω for all input wires ω , computing the garbled circuit could be described by a constant-depth circuit, since the garbled function tables can be computed in parallel after choosing the wire masks and keys. This circuit, however, would be rather large since it entails evaluating the cipher. Therefore, to avoid evaluating the cipher within the asynchronous MPC, the parties use the distributed-encryption technique by Damgård and Ishai [19]: Instead of computing $\text{Enc}_{k_1, k_2}^T(m)$ for a message m , two keys k_1 and k_2 , and a tweak T , the parties first jointly choose $2n$ subkeys k_1^1, \dots, k_1^n and k_2^1, \dots, k_2^n , compute a Shamir sharing of $m = ([m]_1, \dots, [m]_n)$, open $[m]_i$ as well as k_1^i and k_2^i to p_i for every i , and then each party encrypts its share $[m]_i$ of m using its two subkeys k_1^i and k_2^i and sends the resulting ciphertext $\text{Enc}_{k_1^i, k_2^i}^T([m]_i)$ to all parties.

In order to decrypt, a party in possession of all keys recovers the shares by decrypting the ciphertexts received from other players and waits until it has $2t + 1$ t -consistent shares, which it uses to reconstruct m .²³

Asynchronously evaluating Boolean circuits. Protocol π_{BKR} [8], which we wish to use to realize $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$, evaluates arithmetic circuits over fields with more than two elements; the circuit representing $f_{\text{PREP}}^{\text{Circ}}$, however, is Boolean. Thus, in order to evaluate it via π_{BKR} we transform it into an (arithmetic) circuit over an extension field F of $\text{GF}(2)$, denoted $C_{f_{\text{PREP}}^{\text{Circ}}}^F$, by having every NAND gate with inputs $x, y \in \{0, 1\}$ replaced by the computation $1 - xy$, which can be implemented using addition and multiplication over the extension field F . The above transformation, however, works only if all the inputs to the circuit corresponding to bits in the Boolean circuit are either 0 or 1 in the corresponding field F . For the honest parties this is easy to enforce: they encode a 0 (resp., 1) input bit into the 0 (resp., 1) element of F . The adversary, however, might try to cheat by giving “bad” inputs, namely, inputs in $F \setminus \{0, 1\}$. We now show an explicit construction to ensure that the adversary cannot give any value other than 0 or 1, resulting in a simple adaptation of protocol π_{BKR} .²⁴

²³ Our protocol ensures that each party eventually receives these many encrypted shares (see below).

²⁴ In principle, the arithmetic circuit “re-compiler” technique by Genkin *et al.* [24] could also be used for this purpose, although it is not shown to work for π_{BKR} nor be constant-depth. In addition, the functionality of the re-compiler is richer, as it allows to restrict possible malicious strategies during the evaluation of the circuit, which is not needed here.

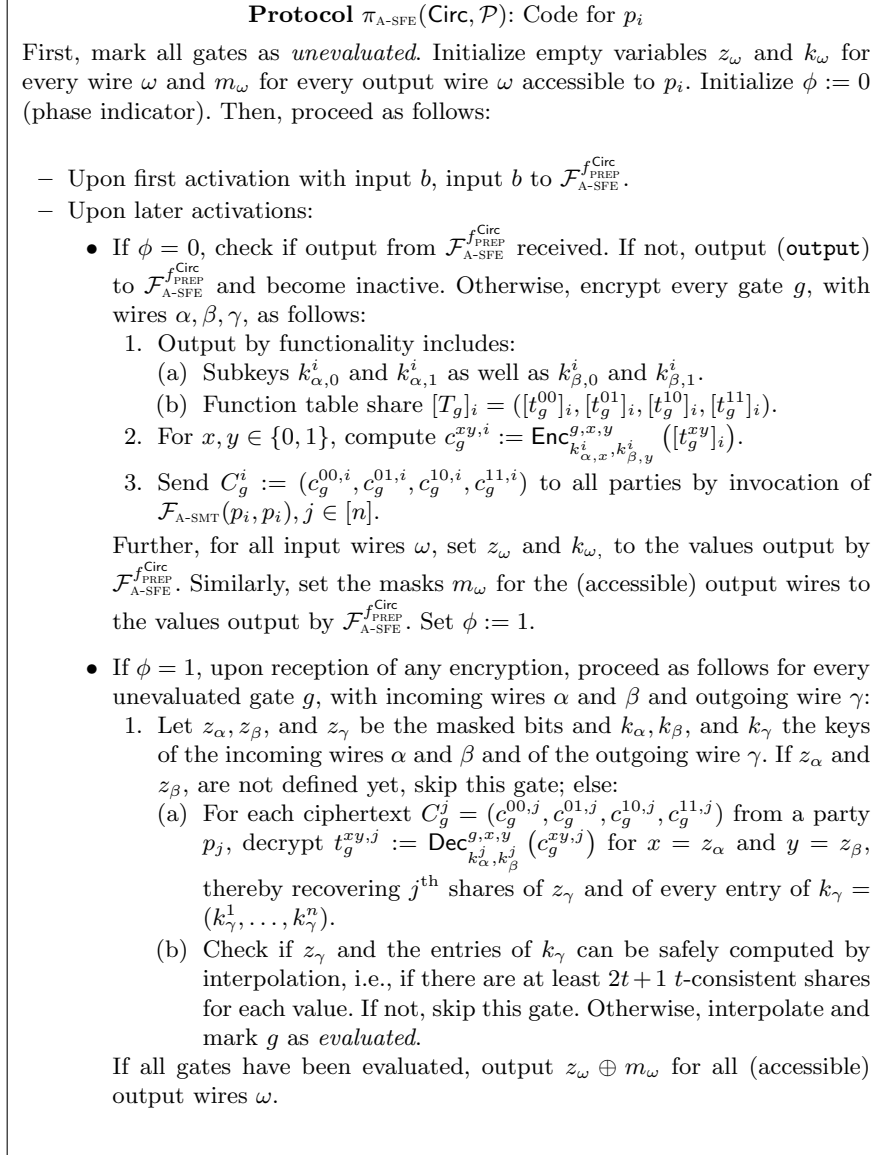


Fig. 5. Our constant-round asynchronous SFE protocol in the $\{\mathcal{F}_{\text{A-SFE}}^{\text{Circ}}^{\text{PREP}}, \mathcal{F}_{\text{A-SMT}}\}$ -hybrid model

Before describing the solution we recall the reader how π_{BKR} evaluates a given circuit. We omit several low-level details and keep the description at the level which is appropriate for a formal description of our adaptation; the interested reader is referred to [8] for further details. π_{BKR} follows the

gate-by-gate evaluation paradigm [26,7]: The circuit is evaluated in a gate-by-gate fashion, where the invariant is that after the evaluation of each gate (in fact, of each bulk of gates that are at the same multiplicative depth), the output of the gates is verifiably shared among all the parties. In fact, [8] defines the notion of *Ultimate Secret Sharing* (USS) which is a version of VSS that is appropriate for asynchronous computation with $t < n/3$;²⁵ More concretely, the first step is to process all input gates in parallel (i.e., receive inputs from all parties); this step finishes with every party holding a share of the input of each party p_i . As already mentioned, due to asynchrony, the inputs of some, up to t , (honest) parties might not be considered. The set **Core** of these parties whose inputs are considered (the so-called *core set* [8,5]) is decided by π_{BKR} (and agreed upon by all parties) during the evaluation of the input gates, while the input of the parties not in the core set is set to a default value, in our case to 0 (i.e., a default USS of 0 is adopted as a sharing of these parties' inputs [8]). Once any party has agreed on the core set parties giving input, it goes on to the evaluation of the next gate (in fact, of all gates which are one level deeper in the circuit in parallel).

Our modification to π_{BKR} is as follows. For any party p_j , as soon as p_j has processed all input gates (i.e., holds shares of inputs of all parties in the core set and default shares for the remaining parties), and before any other gate of the arithmetic circuit is computed, p_j does the following: For each party we denote by x'_i the value that is (eventually) shared as p_i 's input when all parties have evaluated the corresponding input gate, and denote by $[x'_i]_j$ p_j 's share of this value.²⁶

Now, instead of continuing to process the original circuit $C_{f_{\text{PREP}}}^F$, we use the following trick from [9], which will allow us to enforce zero/one inputs. Each party uses the shared values x'_i to compute the circuit evaluating the following function: output $\mathbf{c} = (c_1, \dots, c_n)$, where $c_i = x'_i - x_i'^2$ for each p_i . Each party that received the output \mathbf{c} does the following:²⁷ For each p_i , if $c_i \neq 0$, then the parties replace the sharing of x'_i with a default sharing of 0. (That is, as soon as p_j receives the vector \mathbf{c} , for each i with $c_i \neq 0$ p_i replaces his share $[x_i]_j$ of x'_i with a default sharing of 0.) Once a party has completed this step (and replaced his local shares), he continues

²⁵ USS is an adaptation of the bivariate-polynomial sharing technique [7,33] to the asynchronous setting.

²⁶ By the USS property, at this point p_i is committed to x'_i but the adversary has no information on it, i.e., the adversary holds random shares of a USS of x_i .

²⁷ Observe that the eventual delivery property ensures that every party will eventually receive the output.

the execution of π_{BKR} with the modified shares to compute the remainder of the circuit $C_{f_{\text{PREP}}^{\text{Circ}}}^F$.

We denote the above modification of protocol π_{BKR} (in the $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$ -hybrid world where calls to A-BA are replaced by invocations of $\mathcal{F}_{\text{A-BA}}(\mathcal{P})$) with the above mechanism for enforcing inputs in $\{0, 1\}$ by $\pi_{\text{BKR}}^{0/1}$. In order to evaluate the (Boolean) circuit for $f_{\text{PREP}}^{\text{Circ}}$, the parties execute $\pi_{\text{BKR}}^{0/1}$ encoding their inputs and outputs with the following trivial encoding: An input-bit 0 (resp., 1) is encoded as the 0 (resp., 1) element in F , and output 0 (resp., 1) in F is decoded back to the bit 0 (resp., 1). The following lemma states the achieved security.

Lemma 1. *Protocol $\pi_{\text{BKR}}^{0/1}$ for evaluating the circuit $C_{f_{\text{PREP}}^{\text{Circ}}}^F$ with the above trivial encoding UC-securely realizes $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$.*

Proof (sketch). First note that if the inputs of all (honest and corrupted) parties are 0 or 1 (in the arithmetic field F), then the (decoded) output of $C_{f_{\text{PREP}}^{\text{Circ}}}^F$ is the same as the output of the (Boolean) circuit for $f_{\text{PREP}}^{\text{Circ}}$ since all NAND gates with inputs $x, y \in \{0, 1\}$ are computed by $1 - xy$. Next, we argue that $\pi_{\text{BKR}}^{0/1}$ forces the inputs of the adversary to be 0 or 1 and does not modify the inputs of honest parties. Indeed, an honest party p_i in the core set will share inputs $x'_i \in \{0, 1\}$ and therefore $c_i = 0$, which means that his input sharing is not modified by $\pi_{\text{BKR}}^{0/1}$. The same holds for any corrupted party that shares $x'_i = 0$ or $x'_i = 1$. On the other hand, any corrupted party sharing a value other than 0 or 1 will result into an output $c_i \neq 0$ (since the non-zero elements in F form a multiplicative group of order $|F| - 1$) and its input will be set to 0.

Note that the eventual termination of π_{BKR} ensures that all parties will eventually receive the output vector c and will therefore resume the computation of the original circuit $C_{f_{\text{PREP}}^{\text{Circ}}}^F$, which (also due to the eventual termination of π_{BKR}) will terminate. The simulation of $\pi_{\text{BKR}}^{0/1}$ is easily reduced to the simulation of π_{BKR} : The evaluation of the extra component that computes the c_i 's can be easily simulated as they are random sharings of 0 for all honest parties in the core set, and for corrupted parties they are functions of the sharing of x'_i that the adversary creates in the input-processing phase, which for corrupted parties is fully simulatable. For the rest of the simulation, the simulator simply uses the π_{BKR} simulator. Thus the indistinguishability of the simulation follows from the security of π_{BKR} . \square

Putting things together. The detailed description of protocol $\pi_{\text{A-SFE}}(\text{Circ}, \mathcal{P})$ is presented in Figure 5. As already said, we describe the protocol in the $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$ -hybrid model, where $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$ can be replaced with $\pi_{\text{BKR}}^{0/1}$ using Lemma 1 and the universal composition theorem. At a high-level, the protocol proceeds as follows: In the first phase, the parties send their inputs to the functionality $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$. The (randomized) function $f_{\text{PREP}}^{\text{Circ}}$ chooses the random masks, the subkeys, and computes Shamir sharings of the masked function tables (which are the values that need to be encrypted). Moreover, based on the inputs, it computes the masked value and the corresponding key of every input wire. The formal specification of $f_{\text{PREP}}^{\text{Circ}}$ can be found in Figure 4. The fact that $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$ can be evaluated by a constant-depth circuit is illustrated in Figure 6, which provides a diagram describing the structure of such a circuit. Each of the rectangles corresponds to a collection of independent constant-depth circuits that are evaluated in parallel.

In the second phase of the protocol, as soon as a party receives output from $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$, it encrypts all the shares obtained using the appropriate subkeys and sends the resulting ciphertexts to all parties, as shown in Figure 5. Then, it proceeds to locally evaluate the gates. For each gate, the party waits for ciphertexts from other parties and decrypts them. For a specific entry in the function table, the party has to wait until it has $2t + 1$ t -consistent shares of that entry (see again Figure 5).²⁸ Note that since all of the at least $2t + 1$ honest parties are guaranteed to obtain an output from $\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$, they will all properly encrypt their function tables and send out the resulting ciphertexts. Therefore, the wait for $2t + 1$ t -consistent shares is finite.

4.2 Analysis of the Protocol

Theorem 2. *Let Circ be a given boolean circuit and f_{Circ} be the n -party function computed by Circ . Protocol $\pi_{\text{A-SFE}}(\text{Circ}, \mathcal{P})$ securely realizes $\mathcal{F}_{\text{A-SFE}}^{f_{\text{Circ}}}$ in the $\{\mathcal{F}_{\text{A-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}, \mathcal{F}_{\text{A-SMT}}\}$ -hybrid model tolerating an adaptive adversary who corrupts up to $t < n/3$ of the parties and making black-box use of a PRF.*

A full proof of Theorem 2 can be found in the full version of this paper [18]. Here we only provide a high-level sketch.

²⁸ Note that using the Berlekamp-Welch algorithm, this can be achieved efficiently.




Fig. 6. Bird’s-eye view of the arithmetic circuit **Prep** computing function $f_{\text{PREP}}^{\text{Circ}}$. Each box represents a constant-depth circuit.

Proof (sketch). The output of each party from the evaluation of $\mathcal{F}_{\mathcal{A}\text{-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$ contains (among other things) a t -out-of- n sharing of the garbled circuit for computing function f_{Circ} . After receiving the output from $\mathcal{F}_{\mathcal{A}\text{-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$ and encrypting as described in Figure 5, the only time the parties have to wait is for the encryptions of $2t+1$ t -consistent shares of garbled function-table entries from other parties. Since all of the at least $2t+1$ honest parties are guaranteed to obtain an output from $\mathcal{F}_{\mathcal{A}\text{-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$, they will all properly encrypt their function tables and send out the resulting ciphertexts at some point. Therefore, the wait for $2t+1$ t -consistent shares is finite.

Moreover, the adversary cannot make an honest party accept a wrong value for any entry of the garbled gate: Observe that in any set of $2t+1$ shares that a party receives, at least $t+1$ are from honest parties. These $t+1$ shares uniquely define the degree- t sharing polynomial F and, therefore, they can only be combined with correct shares (as output by $\mathcal{F}_{\mathcal{A}\text{-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$). This implies that wrong shares sent by the adversary cannot make any honest party choose any other polynomial than F .

The simulator \mathcal{S} for an adversary \mathcal{A} roughly proceeds as follows: It emulates towards \mathcal{A} the behavior of $\mathcal{F}_{\mathcal{A}\text{-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$ and the channels $\mathcal{F}_{\mathcal{A}\text{-SMT}}$. The security of the circuit-garbling technique and that of Shamir sharings allows \mathcal{S} to perfectly simulate the entries of the garbled function tables that one would decrypt during a local evaluation of the garbled circuit, even *without* knowing the actual inputs. Moreover, the security of the double-key cipher ensures that the remaining entries are hidden, and can thus be replaced by dummy values (which can again be done without knowing the inputs) without causing a noticeable difference in the view of \mathcal{A} . \square

We now turn to the analysis of the protocol’s round complexity. It is straightforward to verify that our protocol (assuming hybrid $\mathcal{F}_{A\text{-SFE}}^{f_{\text{PREP}}^{\text{Circ}}}$) needs only two communication rounds for each party p_i : one round in which p_i sends its inputs to the functionality and receives its outputs, and one round in which p_i sends all its corresponding encryptions and receives the encryptions of other parties. Moreover, the function $f_{\text{PREP}}^{\text{Circ}}$ can be represented by an arithmetic circuit Prep over a finite field \mathbb{F} with constant (multiplicative) depth: The players first jointly generate the subkeys and the masks. A straightforward method for generating a random field element (such as the subkeys) is to take a random input from every party and computing the sum. Generating a random bit $b \in \{0, 1\} \subseteq \mathbb{F}$ can be reduced to generating random field elements as shown by Genkin *et al.* [24]. Given the masks and the subkeys, computing the function table and a Shamir sharing thereof can clearly be done in constant depth and, most importantly, in parallel for every gate.

Combining the above and Theorem 2 with Theorem 1 yields the following corollary:

Corollary 1. *Let Circ be a given boolean circuit and f_{Circ} be the n -party function computed by Circ . There exists a constant-round protocol which securely realizes $\mathcal{F}_{A\text{-SFE}}^{f_{\text{Circ}}}$ in the $\{\mathcal{F}_{A\text{-BA}}, \mathcal{F}_{A\text{-SMT}}\}$ -hybrid model tolerating an adaptive adversary who corrupts up to $t < n/3$ of the parties and making black-box use of a PRF.*

References

1. Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing, 2-4 May 1988, Chicago, Illinois, USA. ACM (1988)
2. Beaver, D., Micali, S., Rogaway, P.: The round complexity of secure protocols (extended abstract). In: STOC. pp. 503–513. ACM (1990)
3. Beerliová-Trubíniová, Z., Hirt, M.: Simple and efficient perfectly-secure asynchronous mpc. In: Kurosawa, K. (ed.) ASIACRYPT. Lecture Notes in Computer Science, vol. 4833, pp. 376–392. Springer (2007)
4. Bellare, M., Hoang, V.T., Rogaway, P.: Foundations of garbled circuits. In: the ACM Conference on Computer and Communications Security, CCS’12, Raleigh, NC, USA, October 16-18, 2012. pp. 784–796 (2012), <http://doi.acm.org/10.1145/2382196.2382279>
5. Ben-Or, M., Canetti, R., Goldreich, O.: Asynchronous secure computation. In: STOC. pp. 52–61 (1993)
6. Ben-Or, M., El-Yaniv, R.: Resilient-optimal interactive consistency in constant time. Distributed Computing 16(4), 249–262 (2003), <http://dx.doi.org/10.1007/s00446-002-0083-3>
7. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: STOC [1], pp. 1–10

8. Ben-Or, M., Kelmer, B., Rabin, T.: Asynchronous secure computations with optimal resilience (extended abstract). In: PODC. pp. 183–192 (1994)
9. Boneh, D., Goh, E.J., Nissim, K.: Evaluating 2-DNF formulas on ciphertexts. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 325–341. Springer, Berlin, Germany, Cambridge, MA, USA (Feb 10–12, 2005)
10. Bracha, G.: An asynchronous $[(n-1)/3]$ -resilient consensus protocol. In: Probert, R.L., Lynch, N.A., Santoro, N. (eds.) 3rd ACM PODC. pp. 154–162. ACM Press, Vancouver, British Columbia, Canada (Aug 27–29, 1984)
11. Cachin, C., Kursawe, K., Shoup, V.: Random oracles in Constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology* 18(3), 219–246 (Jul 2005)
12. Canetti, R.: Studies in Secure Multiparty Computation and Applications. Ph.D. thesis, Weizmann Institute of Technology, (<http://www.wisdom.weizmann.ac.il/oded/PSX/ran-phd.pdf>) (6 1995)
13. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: 42nd FOCS. pp. 136–145. IEEE Computer Society Press, Las Vegas, Nevada, USA (Oct 14–17, 2001)
14. Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: 28th ACM STOC. pp. 639–648. ACM Press, Philadelphia, Pennsylvania, USA (May 22–24, 1996)
15. Canetti, R., Rabin, T.: Fast asynchronous byzantine agreement with optimal resilience. In: 25th ACM STOC. pp. 42–51. ACM Press, San Diego, California, USA (May 16–18, 1993)
16. Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols (extended abstract). In: STOC [1], pp. 11–19
17. Cohen, R.: Asynchronous secure multiparty computation in constant time. In: Public-Key Cryptography - PKC 2016 - 19th IACR International Conference on Practice and Theory in Public-Key Cryptography, Taipei, Taiwan, March 6–9, 2016, Proceedings, Part II. pp. 183–207 (2016), http://dx.doi.org/10.1007/978-3-662-49387-8_8
18. Coretti, S., Garay, J., Hirt, M., Zikas, V.: Constant-round asynchronous multiparty computation. *Cryptology ePrint Archive, Report 2016/208* (2016), <http://eprint.iacr.org/2016/208>
19. Damgård, I., Ishai, Y.: Constant-round multiparty computation using a black-box pseudorandom generator. In: Shoup, V. (ed.) CRYPTO. Lecture Notes in Computer Science, vol. 3621, pp. 378–394. Springer (2005)
20. Dolev, D., Strong, H.R.: Authenticated algorithms for byzantine agreement. *SIAM J. Comput.* 12(4), 656–666 (1983), <http://dx.doi.org/10.1137/0212045>
21. Feldman, P., Micali, S.: Optimal algorithms for byzantine agreement. In: 20th ACM STOC. pp. 148–161. ACM Press, Chicago, Illinois, USA (May 2–4, 1988)
22. Fischer, M.J., Lynch, N.A.: A lower bound for the time to assure interactive consistency. *Inf. Process. Lett.* 14(4), 183–186 (1982), [http://dx.doi.org/10.1016/0020-0190\(82\)90033-3](http://dx.doi.org/10.1016/0020-0190(82)90033-3)
23. Fischer, M.J., Lynch, N.A., Paterson, M.: Impossibility of distributed consensus with one faulty process. In: Fagin, R., Bernstein, P.A. (eds.) Proceedings of the Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, March 21–23, 1983, Colony Square Hotel, Atlanta, Georgia, USA. pp. 1–7. ACM (1983), <http://doi.acm.org/10.1145/588058.588060>
24. Genkin, D., Ishai, Y., Prabhakaran, M., Sahai, A., Tromer, E.: Circuits resilient to additive attacks with applications to secure computation. In: Symposium on

- Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014. pp. 495–504 (2014), <http://doi.acm.org/10.1145/2591796.2591861>
25. Goldreich, O.: Foundations of Cryptography: Basic Applications, vol. 2. Cambridge University Press, Cambridge, UK (2004)
 26. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or a completeness theorem for protocols with honest majority. In: STOC. pp. 218–229. ACM (1987)
 27. Hirt, M., Nielsen, J.B., Przydatek, B.: Cryptographic asynchronous multi-party computation with optimal resilience (extended abstract). In: Cramer, R. (ed.) EUROCRYPT. Lecture Notes in Computer Science, vol. 3494, pp. 322–340. Springer (2005)
 28. Hirt, M., Nielsen, J.B., Przydatek, B.: Asynchronous multi-party computation with quadratic communication. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP (2). Lecture Notes in Computer Science, vol. 5126, pp. 473–485. Springer (2008)
 29. Hirt, M., Zikas, V.: Adaptively secure broadcast. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 466–485. Springer, Berlin, Germany, French Riviera (May 30 – Jun 3, 2010)
 30. Katz, J., Koo, C.Y.: On expected constant-round protocols for byzantine agreement. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 445–462. Springer, Berlin, Germany, Santa Barbara, CA, USA (Aug 20–24, 2006)
 31. Katz, J., Maurer, U., Tackmann, B., Zikas, V.: Universally composable synchronous computation. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 477–498. Springer, Berlin, Germany, Tokyo, Japan (Mar 3–6, 2013)
 32. Pfitzmann, B., Waidner, M.: Unconditional byzantine agreement for any number of faulty processors. In: Finkel, A., Jantzen, M. (eds.) STACS 92, 9th Annual Symposium on Theoretical Aspects of Computer Science, Cachan, France, February 13–15, 1992, Proceedings. Lecture Notes in Computer Science, vol. 577, pp. 339–350. Springer (1992), http://dx.doi.org/10.1007/3-540-55210-3_195
 33. Rabin, T., Ben-Or, M.: Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In: 21st ACM STOC. pp. 73–85. ACM Press, Seattle, Washington, USA (May 15–17, 1989)
 34. Schneider, T., Zohner, M.: GMW vs. Yao? efficient secure two-party computation with low depth circuits. In: Sadeghi, A.R. (ed.) FC 2013. LNCS, vol. 7859, pp. 275–292. Springer, Berlin, Germany, Okinawa, Japan (Apr 1–5, 2013)
 35. Shamir, A.: How to share a secret. Commun. ACM 22(11), 612–613 (1979)
 36. Yao, A.C.C.: Protocols for secure computations (extended abstract). In: FOCS. pp. 160–164. IEEE (1982)
 37. Yao, A.C.C.: How to generate and exchange secrets (extended abstract). In: FOCS. pp. 162–167. IEEE (1986)

A History and Related Work (cont’d)

Here we provide a fuller account of related work and put our results in perspective. To give a more complete picture, we start by discussing the development of MPC protocols in the synchronous setting, and then contrast it with the development in the asynchronous setting. Along the

way we also discuss the tools (e.g., setup assumptions and communication resources) that are used in each setting.

Starting with Yao’s seminal paper [36], which introduced the problem of MPC and provided the first solution, a long line of interesting results proved feasibility bounds for *synchronous* networks in various adversarial settings. Goldreich, Micali, and Wigderson [26,25] proved that under computational assumptions (the existence of enhanced trapdoor permutations), any n -party function can be securely computed if and only if up to $t < n$ parties are corrupted passively or up to $t < n/2$ actively. Corresponding bounds for information-theoretic security were shown by Ben-Or, Goldwasser, and Wigderson [7], who proved that perfect security is possible if and only if the adversary corrupts up to $t < n/2$ parties passively or up to $t < n/3$ actively. Similar bounds were concurrently proved by Chaum, Crépeau, and Damgård [16] for the case where a negligible error-probability is allowed and were later improved by Rabin and Ben-Or [33] to achieve optimal resiliency $t < n/3$.

The above works assume point-to-point secure communication and a broadcast channel and, under these assumptions, are secure even against an adaptive adversary [14]. However, in [26,25] both these resources can be implemented with adaptive security assuming a public-key infrastructure and non-committing encryption [14,29]. Similarly, the broadcast channel in [7,16,33] can be emulated by an adaptively secure broadcast protocol [29].²⁹ The round complexity of all the above protocols in the malicious multi-party setting—even with the assumption of a broadcast channel—is linear in the multiplicative depth of the arithmetic circuit corresponding to the function the parties aim to compute.

Beaver, Micali, and Rogaway [2] were the first to provide a constant-round MPC protocol in the synchronous stand-alone model. Their protocol is secure in the computational setting and tolerates an adaptive adversary who actively corrupts up to $t < n/2$ parties. The complexity of [2] was improved by Damgård and Ishai [19], who provided the first constant-round protocol making black-box use of the underlying cryptographic primitive (a pseudo-random generator). Importantly, both [2] and [19] assume a broadcast channel, an assumption essential for obtaining constant-round MPC. Indeed, as proved in [22,20], it is impossible to implement such a broadcast channel from point-to-point communication in a constant number of rounds, and although expected constant-round

²⁹ Because [33] tolerates even $n/3 \leq t < n/2$ corrupted parties, the emulation of broadcast would require an additional setup of information-theoretic pseudo-signatures [32].

broadcast protocols exist in the literature (e.g., [21,30]), using them to instantiate calls within the constructions of [2] or [19] would not yield an expected constant-round protocol [6]. The intuitive reason—formally argued by Ben-Or and El-Yaniv [6]—is that the process of running n such broadcast protocols (even in parallel) does not terminate in an expected constant number of rounds.

The model of *asynchronous* communication with eventual delivery was considered early on in seminal works on fault-tolerant distributed computing (e.g., [23]). The study of optimally resilient MPC in such an asynchronous network was initiated by Ben-Or, Canetti, and Goldreich [5], who proved that any function can be computed by a perfectly secure asynchronous protocol if and only if at most $t < n/4$ parties are corrupted. Following that result, Ben-Or, Kelmer, and Rabin [8] showed that if a negligible error probability is allowed, the bound $t < n/3$ is necessary and sufficient for asynchronous MPC.³⁰ More recently, Hirt *et al.* [27,28] provided computationally secure solutions (i.e., protocols tolerating a computationally bounded adversary) and Beerliová and Hirt [3] perfectly secure solutions with improved communication complexity.

The above asynchronous protocols are secure if one assumes point-to-point communication and an A-BA protocol. Similarly to their synchronous counterparts, all the above protocols—even assuming an A-BA primitive—have round complexity linear in the multiplicative depth of the arithmetic circuit that computes the function, as they follow the standard gate-by-gate evaluation paradigm.

We note in passing that although in the synchronous setting BA implies broadcast, this is not the case in the asynchronous setting. Indeed, Canetti and Rabin [15] provide an asynchronous BA protocol tolerating $t < n/3$ malicious parties, which if every honest party terminates at the latest after a poly-logarithmic number of rounds, securely implements asynchronous BA except with negligible probability. A broadcast protocol with similar guarantees is provably impossible [23], and existence of an asynchronous BA protocol which terminates in a strict constant number of rounds would contradict the impossibility from [22,20]. Similarly to the

³⁰ The necessity of the $t < n/3$ bound follows from the result by Canetti *et al.* [5,12], who argue that this bound is necessary for fail-stop adversaries; it also applies to computational security and assuming A-BA. Moreover, note that in the asynchronous setting, all feasibility bounds are worse by an additive term of t compared to the synchronous setting. Intuitively, this stems from the fact that honest parties cannot distinguish between messages by other honest parties being delayed and messages by corrupted parties not being sent. Thus, in particular, perfectly secure asynchronous MPC is possible only if $t < n/4$.

synchronous case, although solutions for asynchronous BA with expected constant number of rounds exist [11,15], using them in the above asynchronous protocol to replace invocations to asynchronous BA would not yield an expected constant-round MPC protocol [6].³¹

Finally, if one gives up the requirement that the broadcast protocol (eventually) terminates when the sender is corrupted (this results in a primitive known as *A-Cast* [10]), then one can implement it even in a constant number of rounds. (In fact, A-Cast can be easily reduced to asynchronous BA by having the sender send his input to all parties, who then forward this input as soon as it is received to the asynchronous BA primitive).

³¹ Nonetheless, [6] does describe an alternative way of obtaining several asynchronous BA protocols that are guaranteed to all terminate in expected constant number of rounds.