

# A New Way to Achieve Round-Efficient Asynchronous Byzantine Agreement

Simon Holmgaard Kamp

CISPA Helmholtz Center for Information Security  
simon.kamp@cispa.de

**Abstract.** We translate the *expand-and-extract* framework by Fitzi, Liu-Zhang, and Loss (PODC 21) to the asynchronous setting. While they use it to obtain a synchronous BA with  $2^{-\lambda}$  error probability in  $\lambda + 1$  rounds, we make it work in asynchrony in  $\lambda + 3$  rounds. At the heart of their solution is a generalization of crusader agreement and graded agreement to any number of grades called *proxcensus*. They achieve graded consensus with  $2^r + 1$  grades in  $r$  rounds by reducing proxconsensus with  $2s - 1$  grades to proxconsensus with  $s$  grades in one round. The *expand-and-extract* paradigm uses proxconsensus to *expand* binary inputs to  $2^\lambda + 1$  grades in  $\lambda$  rounds before *extracting* a binary output by partitioning the grades using a  $\lambda$  bit common coin. However, the proxconsensus protocol by Fitzi et al. does not translate to the asynchronous setting without lowering the corruption threshold or using more rounds in each recursive step.

This is resolved by attaching *justifiers* to all messages: forcing the adversary to choose between being ignored by the honest parties, or sending messages with certain validity properties. Using these we define validated proxconsensus and show that it can be instantiated in asynchrony with the same recursive structure and round complexity as synchronous proxconsensus. In asynchrony the extraction phase incurs a security loss of one bit which is recovered by expanding to twice as many grades using an extra round of communication. This results in a  $\lambda + 2$  round VABA and a  $\lambda + 3$  round BA, both with  $2^{-\lambda}$  error probability and communication complexity matching Fitzi et al.

## 1 Introduction

Following the *expand-and-extract* paradigm presented by Fitzi, Liu-Zhang, and Loss [FLZL21] we present a concretely round efficient asynchronous Monte Carlo style BA which runs for a fixed number of rounds,  $\lambda + 3$ , to reach agreement on a binary decision with probability at least  $1 - 2^{-\lambda}$  using signatures and a common coin.

The expand-and-extract paradigm generalizes Feldman and Micali [FM97] (FM) in which the parties first run crusader agreement to obtain an output which is the input of an honest party or an inconclusive value “?”, with the guarantee that no honest parties get different output bits. Then a random coin is flipped and parties with output “?” use the coin as input to the next round. With probability  $> 1/2$  this makes the system enter a univalent configuration where honest parties input the same bit to all future iterations. While it is possible to detect this state and terminate in an expected constant number of rounds; in order to reach agreement with probability  $1 - 2^{-\lambda}$  the protocol must run  $\lambda$  FM iterations, i.e. for  $2\lambda$  rounds. Fitzi et al. cut this worst case round complexity almost in half to  $\lambda + 1$  by observing that you can remove all but one of the coin flips if you first *expand* to  $2^\lambda + 1$  grades in  $\lambda$  rounds and then flip a  $\lambda$ -bit coin to *extract* a bit decision that is consistent with probability  $1 - 2^{-\lambda}$ . Their expansion technique is inherently synchronous, but we adapt it to the asynchronous setting while matching the asymptotic communication complexity and

the concrete round complexity up to needing one extra round to implement validated BA (VABA) as defined in [CKPS01] or two extra rounds to implement a full-fledged BA.

To compare with the FM approach in the asynchronous setting, it is worth noting that crusader agreement followed by a coin flip does not solve BA in asynchrony. This is because the adversary is assumed to learn the output of the coin as soon as some honest party wishes to flip the coin. Assuming a bivalent configuration the adversary can make an honest party output “?” in order to learn the coin  $c$  and then let other honest parties receive output  $1 - c$  to maintain the bivalent configuration indefinitely. Traditionally this was solved by using graded agreement with at least 4 grades. But this requires an extra round of communication before flipping the coin, resulting in FM iterations with 3 rounds and a worst case round complexity of  $3\lambda$ . Thus, to the best of our knowledge our solution improves on the worst case round complexity of asynchronous by almost a factor 3 to  $\lambda + 3$ .

In the synchronous BA by Fitzi et al. the saved rounds come from taking the  $\lambda$  rounds spent on flipping a 1-bit coin in each FM iteration and replacing them with a single round to flip a  $\lambda$ -bit coin after expanding. The resulting round complexity is  $\lambda + 1$  instead of  $2\lambda$ . In the same vein we can view our VABA protocol as taking – from each of the  $\lambda$  FM iterations – both the round spent on coin flipping and the round spent on upgrading crusader agreement to graded consensus and replacing all of them by just two rounds that serve the same purpose after the expansion phase, resulting in a round complexity of  $\lambda + 2$  rather than  $3\lambda$ .

*Techniques.* The expand-and-extract approach first *expands* to  $2^\lambda + 1$  grades and then uses a  $\lambda$ -bit coin to *extract* a decision. In the expansion step the parties input their bit to a *proxcensus* protocol [FLZL21] which outputs a bit and a grade which serves as an indicator of confidence in the bit. Crusader agreement and graded agreement are well-known special cases of proxcensus with respectively 3 and 5 grades. To simplify the presentation and proofs we depart from the syntax of proxcensus given by Fitzi et al., so proxcensus outputs a nonnegative integer smaller than the number of grades. I.e. the output of crusader agreement is in  $\{0, 1, 2\}$  instead of  $\{(0, 1), (?, 0), (1, 1)\}$  and if all honest parties input the same bit 0 (or 1), their output must be the minimal (or maximal) grade. All honest parties are guaranteed to have approximate agreement in the sense that they have adjacent integer outputs. After expanding to a large space of  $2^\lambda + 1$  grades, the output bit is *extracted* by splitting the space in two using a random coin. If the obtained grade is strictly greater than the coin, the output is 1, otherwise it is 0. This means that if a party has the minimal (or maximal) grade then they output 0 (or 1) regardless of the coin, so the validity of BA reduces to the validity of proxcensus.

When we enter the extraction phase, a more general version of the attack on the FM approach in asynchrony described above also applies to asynchronous expand-and-extract. As we cannot wait to flip the coin until all honest parties are done expanding: each party must initiate the coin flip when they are done expanding and this leaks the value of the common coin to adversary. The output grades of proxcensus are guaranteed to be adjacent, but when the first honest party gives output, there are still potentially two candidates for the adjacent grade. This gives the adversary two shots at guessing the random coin instead

of one, doubling the error probability. However, adding an extra round to expand to twice as many grades recovers this lost bit of security.

Finally, an additional round is used to reduce BA to VABA by establishing a threshold signature on the input bit as suggested in [CKPS01].

*Related work.* Expand-and-extract along with the notion of proxconsensus were introduced in [FLZL21]. Proxconsensus is in turn an adaption of the proxcast definition used in [CFF<sup>+</sup>05]. Justifiers were introduced in [DYMM<sup>+</sup>20], but we use the definitions from [KN23]. A very similar notion of transferable justifiers that also considers a version of adversarially crafted justified outputs exist in [LN24], however we need to apply it to properties of messages rather than just outputs, which matches the security game in [KN23]. The logic used to get validated proxconsensus draws on the partially synchronous BA in [KNTT22], although the network model and definitions of justifiers are quite different. The result of [FLZL21] has later been improved in [GGLZ22], but the techniques do not appear to be compatible with asynchrony. There has not been much recent progress on the concrete round efficiency of Monte Carlo style BA in the asynchronous setting. However, the recent work of Erbes and Wattenhofer [EW24] provides a  $2^\lambda$  graded consensus protocol with  $6(\lambda + 1)$  rounds and suggests that [AC10] and [BBB<sup>+</sup>24] can be combined into a  $2^\lambda$  graded consensus using  $3(\lambda + 1)$  rounds. We improve on this by a factor of 3 (see Corollary 1).

## 1.1 Preliminaries

We provide protocols for  $n$  parties  $P_1, \dots, P_n$  where  $t < n/3$  can be adaptively Byzantine corrupted. We assume an asynchronous network where message delivery is handled by the adversary with no upper bound on the delay. Liveness properties hold under eventual delivery of the messages.

*Threshold Signatures and Common Coin* These primitives are standard and we only give brief informal descriptions. For simplicity we treat them as unconditionally secure. We assume key shares of a threshold signature scheme for multiple different thresholds have been setup between the parties. We will be using thresholds  $n - t$  and  $t + 1$  and assume that (partial) signatures have a size of  $\lambda$  bits. We also assume the parties can run a common coin primitive  $\Pi_{CC}$  that allow the parties to flip a  $\lambda$  bit coin which is unpredictable to the adversary until the first honest party initiates the protocol. This can be instantiated from threshold signatures following [CKS05].

*Justifiers* We use justifiers as defined in [KN23]. These are ways to demonstrate some notion of validity of each protocol message. A simple example is justifying the input bit to VABA protocol using threshold signatures as suggested in [CKPS01].

**Definition 1 (Justifier [KN23]).** *For a message identifier  $\text{mid}$  we say that  $J^{\text{mid}}$  is a justifier if the following properties hold.  $J^{\text{mid}}$  is a predicate depending on the message  $m$  and the local state of a party. When we write pseudo-code then we write  $J^{\text{mid}}(m)$  to denote that the party  $P$  executing the code computes  $J^{\text{mid}}$  on  $m$  using its current state. In definitions and*

proofs we write  $J^{\text{mid}}(m, \mathsf{P}, \tau)$  to denote that we apply  $J^{\text{mid}}$  to  $m$  and the local state of  $\mathsf{P}$  at time  $\tau$ .

**Monotone:** *If for an honest  $\mathsf{P}$  and some time  $\tau$  it holds that  $J^{\text{mid}}(m, \mathsf{P}, \tau) = \top$  then at all  $\tau' \geq \tau$  it holds that  $J^{\text{mid}}(m, \mathsf{P}, \tau') = \top$ .*

**Propagating:** *If for honest  $\mathsf{P}$  and some point in time  $\tau$  it holds that  $J^{\text{mid}}(m, \mathsf{P}, \tau) = \top$ , then eventually the execution will reach a time  $\tau'$  such that  $J^{\text{mid}}(m, \mathsf{P}', \tau') = \top$  for all honest parties  $\mathsf{P}'$ .*

In our protocols all justifiers are explicit certificates based on threshold signatures and some auxiliary information send along with the message. We therefore omit the party and time from the justifiers. The following two definitions provide a framework to reason about properties of justified messages.

**Definition 2 (Possible Justified Messages [KN23]).** *Let  $\Pi$  be a protocol. When we say that an  $\ell$ -ary predicate  $P$  holds for all possible justified messages we mean: Run the protocol  $\Pi$  under attack by the adversary. At some point the adversary may output a sequence of triples  $(\mathsf{P}^1, \text{mid}^1, m^1), \dots, (\mathsf{P}^\ell, \text{mid}^\ell, m^\ell)$ . We say that the adversary wins if the message identifiers  $\text{mid}^1, \dots, \text{mid}^\ell$  identify messages of  $\Pi$ ,  $\mathsf{P}^1, \dots, \mathsf{P}^\ell$  are honest (but not necessarily distinct) parties, for  $j = 1, \dots, \ell$  it holds that  $J^{\text{mid}^j}(m^j) = \top$  at  $\mathsf{P}^j$ , and  $P(m^1, \dots, m^\ell) = \perp$ . Otherwise the adversary loses the game. Any PPT adversary should win with negligible probability.*

**Definition 3 (Possible Justified Outputs [KN23]).** *Let  $\Pi$  be a protocol with output justifier  $J$ . When we say that an  $\ell$ -ary predicate  $P$  holds for all possible justified outputs we mean: Let  $\Pi'$  be the protocol  $\Pi$  with only change being that each party on getting output, sends their output to all parties if this was not already done. Run the protocol  $\Pi'$  under attack by the adversary. At some point the adversary may output a sequence of triples  $(\mathsf{P}^1, \text{mid}^1, m^1), \dots, (\mathsf{P}^\ell, \text{mid}^\ell, m^\ell)$ . We say that the adversary wins if the  $\text{mid}^1, \dots, \text{mid}^\ell$  are identified with outputs of  $\Pi$ ,  $\mathsf{P}^1, \dots, \mathsf{P}^\ell$  are honest (but not necessarily distinct) parties, for  $j = 1, \dots, \ell$  it holds that  $J^{\text{mid}^j}(m^j) = \top$  at  $\mathsf{P}^j$ , and  $P(m^1, \dots, m^\ell) = \perp$ . Otherwise the adversary loses the game. Any PPT adversary should win with negligible probability.*

The protocol in Fig. 1 corresponds to the procedure for justifying inputs in [CKPS01]. The parties multicast their input bits and initially try to collect signature shares on some input bit from  $t + 1$  parties, implying that it is the input of an honest party. At the same time (assuming  $n > 3t$ ) it is guaranteed that parties eventually receive  $t + 1$  shares for one of the two inputs, since the bit input by the majority of the honest parties will account for at least  $t + 1$  shares. Using the definition of justifiers from [KN23] we can say that the  $\Pi_{\text{IVG}}$  in Fig. 1 has the liveness property that if all honest parties start running the protocol, then eventually they receive a justified bit as output. It also has the safety property that any possible justified output satisfying  $\Pi_{\text{IVG}} \cdot J_{\text{OUT}}$  was the input of an honest party.

**Input validation gadget  $\Pi_{IVG}$ .**

- On input  $b_i$ ,  $P_i$  multicasts  $b_i$  with a signature and a partial signature with threshold  $t + 1$ .
- On receiving  $b$  for some bit  $b$  from  $t + 1$  distinct parties:  $P_i$  combines the shares of the  $t + 1$  threshold signature scheme and outputs  $b$  justified by the threshold signature.

**Fig. 1.** An input validation gadget.

## 2 Validated Asynchronous Proxcensus

In this section we define validated proxconsensus, which is to the definition of proxconsensus in [FLZL21] what VABA [CKPS01] is to BA. To ease the notation in the protocol and proofs we will define the output over the integers, such that  $VAProx-(G)$  has outputs in  $\{0, \dots, G-1\}$  rather than a bit (or ‘?’) with a grade up to  $\lfloor G/2 \rfloor$ .

We will follow the approach of reducing the problem of  $VAProx-(2s-1)$  to  $VAProx-(s)$  using one round of communication.

**Definition 4 (Validated Proxcensus).** *Let  $\Pi_{VAPROX-(G)}(J_{IN})$  be a protocol for  $n$  parties, parameterized by an input justifier  $J_{IN}$ , and outputting  $y \in \{0, \dots, G-1\}$  satisfying an output justifier  $J_{OUT}$ . We say that  $\Pi_{VAPROX-(G)}(J_{IN})$  solves  $VAProx-(G)$  if the following holds:*

**Liveness** *If every honest party  $P_i$  have justified input  $x \in \{0, 1\}$  where  $J_{IN}(x) = \top$ , then eventually every honest party  $P_j$  will have justified output  $y \in \{0, \dots, G-1\}$  where  $J_{OUT}(y) = \top$ .*

**Justified Approximate Agreement** *For all possible justified outputs  $y$  and  $y'$ :  $|y - y'| \leq 1$ .*

**Justified Validity** *If  $b$  is the only possible  $J_{IN}$  justified bit, then  $y = b \cdot (G-1)$  for all possible justified outputs  $y$ .*

*Remark 1.* It is easy to map from an output  $y \in [G]$  to the  $(b, g)$  representation used in [FLZL21]. In our case  $G$  is always odd, so we can define a middle grade  $G' = \lfloor G/2 \rfloor$  and

$$\text{map } y \text{ to } (b, g) \text{ where } g = |G' - y| \text{ and } b = \begin{cases} ? & \text{if } y = G' \\ 0 & \text{if } y < G' \\ 1 & \text{if } y > G' \end{cases}.$$

To motivate the definition and the use of justifiers, let us first consider the simplest (non-trivial) version of proxconsensus: Crusader Agreement. Parties have input in  $\{0, 1\}$  and outputs in  $\{0, 1, 2\}$ . If we use the logic from Fitzi et al. in the asynchronous setting we run into problems with validity, liveness or agreement (unless we assume fewer corruptions or use more than one round). First, a party should not output  $2b$  unless it has seen  $n - t$  votes for  $b$ . Otherwise approximate agreement is easily broken. The validity property says that if all honest parties have input  $b$ , then the only valid output is  $2b$ . In order to make sure that you pick a valid output you can only “trust” your own input, or the other input if it is seen from at least  $t + 1$  parties. But what do you do if you had input  $b$ , only see  $n - t$  votes, and the majority but not all of these votes are on  $b$ ? You do not have enough information to

know that  $1 - b$  is an honest input and thus that 1 is a valid output, but neither do you have the  $n - t$  votes for  $b$  that would allow outputting  $2b$ . As the protocol is asynchronous, there is no way to wait for more than  $n - t$  votes. While this is far from a formal lower bound, it illustrates the problem which occurs because we cannot (as opposed to in the synchronous setting) wait for all honest votes. It can be solved by assuming  $n > 4t$ , or by using two rounds in each step of the recursion.<sup>1</sup>

In order to solve this problem, let us define a *validated* flavour of Crusader Agreement which is to Crusader Agreement what VABA is to BA. Namely, where inputs are checked by a predicate and the validity property only holds with respect to the predicate. We define Justified Validity to say that if only the input  $b$  is justifiable, then  $2b$  should be the only justifiable output bit. With this definition: if you see a justifier for both possible input bits, then you can output bottom without violating Justified Validity. You can also justify it by forwarding the 2 justifiers. To ensure agreement you can now simply wait until you see  $n - t$  votes for the same bit  $b$  in order to output  $2b$ . Finally, for Liveness observe that you will eventually see  $n - t$  votes from honest parties. Either these  $n - t$  votes are on the same bit, or you saw two that were different votes. In either case you obtain an output.

Let us generalize this idea to solve validated proxconsensus. We now reinterpret the consensus rules as doubling an input that is seen  $n-t$  times, or taking the double of the average of two different inputs. Note that after giving output: since all justified outputs are on at most two different adjacent integers, we are in an abstract sense in the same situation as when we started with 0 and 1 being those two integers. If we apply this new interpretation of the consensus rules to those integers and use the invariant that all justifiable outputs of an iteration are adjacent integers, we can keep threading the justified outputs of the each instance of validated crusader agreement into the next one. The maximum grade doubles every round. To see why justified validity is maintained: there is only something to show if the only justified input is  $\beta$ . And in that case each round just doubles the unique justified input and the only justified output after  $r$  rounds will be  $2^r \cdot b$ .

We describe the general protocol in detail in Fig. 2 as reducing VAProx- $(2s - 1)$  to VAProx- $(s)$ . As a base case for the recursion, define  $\Pi_{\text{VAPROX-}(2)}(J_{\text{IN}})$  to output its input with  $J_{\text{OUT}} = J_{\text{IN}}$ . Rephrasing the above: the main insight is that *all possible justified* outputs of VAProx- $(s)$  are on at most 2 adjacent grades, so we can double the grade if we only receive the same grade from  $n - t$  parties, or double the average of 2 justified grades.

We show that the protocol in Fig. 2 satisfies VAProx- $(2^i + 1)$  as defined in Definition 4 for any nonnegative integer  $i$ , using  $i$  rounds of communication. For the base case  $i = 0$  define  $\Pi_{\text{VAPROX-}(2)}(J_{\text{IN}})$  to be the zero round protocol that just returns the input with output justifier  $J_{\text{OUT}} = J_{\text{IN}}$  trivially satisfying Definition 4. We show the induction step:

**Lemma 1.** *If  $\Pi_{\text{VAPROX-}(s)}(J_{\text{IN}})$  satisfies Definition 4, then  $\Pi_{\text{VAPROX-}(2s-1)}(J_{\text{IN}})$  satisfies Definition 4*

---

<sup>1</sup> After we introduce the proxconsensus protocol it should be clear how to combine it with the protocol in Fig. 1 to obtain a recursive solution with  $t$  rounds in each step. But also that the “extra” round is really only needed in the very beginning of the protocol.

**Reduction from  $\Pi_{\text{VAPROX}}-(2s-1)$  to  $\Pi_{\text{VAPROX}}-(s)$  for  $s \geq 2$ .**

- On input  $x_i$  where  $J_{\text{IN}}(x_i) = \top$ ,  $P_i$  runs  $\Pi_{\text{VAPROX}}-(s)(J_{\text{IN}})$  with input  $x_i$ .
- On output  $z_i$  from  $\Pi_{\text{VAPROX}}-(s)(J_{\text{IN}})$ ,  $P_i$  multicasts  $(\text{PROPOSAL}, z_i)$  with a signature and a partial signature with threshold  $n - t$ .
- On receiving justified  $(\text{PROPOSAL}, z)$  and  $(\text{PROPOSAL}, z + 1)$   $P_i$  lets  $y_i = 2z + 1$  and terminates with output  $y_i$  justified by the justifiers for  $z$  and  $z + 1$ .
- On receiving justified  $(\text{PROPOSAL}, z)$  from  $n - t$  distinct parties,  $P_i$  lets  $y_i = 2z$  and terminates with output  $y_i$  justified by a  $n - t$  threshold signature on  $z$ .

**Fig. 2.** A recursive description of the validated proxensus protocol.

*Proof.* For liveness we observe that  $n - t$  honest parties  $\Pi_{\text{VAPROX}}-(s)(J_{\text{IN}}).J_{\text{OUT}}$  justified values  $z_i$  which by Justified Approximate Agreement are either all identical or split between two adjacent integers, so when these are propagated every party has a set of  $n - t$  values that allow defining  $y_i$  through one of the two cases. Justified Validity follows from Justified Validity of  $\Pi_{\text{VAPROX}}-(s)(J_{\text{IN}})$ , as the input justifier is shared. So since all justified  $z$  are identical, every party gets the same  $y$ . In particular, if the only justified input to the inner protocol is  $b$  then the only  $\Pi_{\text{VAPROX}}-(s)(J_{\text{IN}}).J_{\text{OUT}}$ -justified output is  $z = b \cdot (s - 1)$ , thus the only  $\Pi_{\text{VAPROX}}-(2s-1)(J_{\text{IN}}).J_{\text{OUT}}$ -justified output is  $2b \cdot (s - 1) = b \cdot (2s - 2)$ . For Justified Approximate Agreement we again rely on Justified Approximate Agreement of the  $\Pi_{\text{VAPROX}}-(s)(J_{\text{IN}}).J_{\text{OUT}}$ -justified  $z$  values to say that and for any justified  $z_i$  and  $z_j$  we have  $|z_i - z_j| \leq 1$ . We only need to argue that there cannot be  $n - t$  parties who send some  $z_i$  and  $n - t$  parties who sent  $z_j = z_i - 1$ . But such two sets would overlap on at least one honest party as  $n > 3t$ . Thus, by definition of step 3 all justified outputs  $y_i$  and  $y_j$  satisfy  $|z_i - z_j| \leq 1$ .  $\square$

*Constant sized justifiers of odd grades.* The proof of Lemma 1 completes the induction proof as far as correctness and round complexity goes. But notice that while each party only sends a single message in each round, the size of the odd justifiers keep growing, because they depend recursively on justifiers from the previous round. So, the size of the output justifier of an odd grade in round  $i$  is  $O(i\lambda + |J_{\text{IN}}|)$  bits, where  $|J_{\text{IN}}|$  is the size of the input justifier. Meanwhile the justifier of an even grade is compressing in the sense that it comes with a threshold signature which was produced by a set of parties including at least one honest party, and thus is of size  $O(\lambda)$  bits as it does not need to include justifiers from previous rounds. Let  $\Pi_i = \Pi_{\text{VAPROX}}-(2^i+1)$  for  $i \geq 0$ . For  $\Pi_0$  the output justifier is the input justifier of size  $O(|J_{\text{IN}}|)$ . To optimize the size of odd output justifiers to  $O(\lambda + |J_{\text{IN}}|)$  for  $i > 0$ , observe that the justifier of an odd grade is based on a justified even and odd grade from the previous round. Consider an odd output  $y$  of  $\Pi_i$ , which was calculated as  $2z - 1$  based on justified values  $z$  and  $z - 1$ . We can think of this as  $y = 2 \frac{|z' + z''|}{2}$  for some justified values  $z'$  and  $z'' = z' \pm 1$ . Assume W.L.O.G. that  $z'$  is the even of the two grades, then it was justified by an  $n - t$  threshold signature on  $z'/2$  while  $z''$  was justified by the same threshold signature justifier for  $z'/2$  in addition to the justifier for the odd grade  $z'/2 \pm 1$ . Since the threshold signature justifying  $z'$  has contributions from honest parties who saw a justifier for  $z'/2$ , it guarantees the existence of the justifier for  $z'/2$ . This means that – in the context of justifying  $y$  – the justification of  $z''$  only needs to justify its  $z'/2 \pm 1$  component. If  $z'/2 \pm 1$  is even, it

is justified by a threshold signature on  $(z'/2 \pm 1)/2$  (or is 0 and justified by  $J_{\text{IN}}$ ). Otherwise,  $z'/2 \pm 1$  is odd and we apply the above step until we hit another threshold signature or an input justified bit. In summary, the output for any grade at any level can be justified using  $O(\lambda + |J_{\text{IN}}|)$  bits.

This allows our asynchronous validated proxconsensus to match the complexity of the synchronous proxconsensus for  $n > 3t$  parties presented in [FLZL21].

**Corollary 1.** *For any  $r \geq 0$ ,  $\Pi_{\text{VAPROX-}(2^r+1)}(J_{\text{IN}})$  solves  $\text{VAProx-}(2^r + 1)$  with  $r$  rounds of communication and  $O(rn^2(\lambda + |J_{\text{IN}}|))$  bits of communication, where  $|J_{\text{IN}}|$  is the size of the input justifier. When using  $\Pi_{\text{IVG}}$  to establish  $J_{\text{IN}}$ , it solves  $\text{Prox-}(2^r + 1)$  as defined in [FLZL21] in  $r + 1$  rounds with  $O(rn^2\lambda)$  bits of communication.*

### 3 Validated Asynchronous BA

We give a brief description of the extraction phase, which largely follows [FLZL21]. The expansion step requires an extra round of communication to expand to  $2^{\lambda+1} + 1$  grades and make up for the error probability being doubled in asynchrony. (In contrast to the  $2^\lambda + 1$  that suffice in the synchronous case.) As the proxconsensus is validated, the combined expand-and-extract procedure only yields a validated BA protocol. A full BA protocol is given in Section 4. We first define validated asynchronous BA (VABA).

**Definition 5 (Validated Asynchronous BA).** *Let  $\Pi_{\text{VABA}}(J_{\text{IN}})$  be a protocol for  $n$  parties parameterized by an input justifier  $J_{\text{IN}}$  outputting  $y \in \{0, 1\}$  satisfying an output justifier  $J_{\text{OUT}}$ . We say that  $\Pi_{\text{VABA}}(J_{\text{IN}})$  is a secure VABA protocol if the following properties hold:*

**Liveness** *If every honest party  $P_i$  has justified input  $x_i \in \{0, 1\}$  where  $J_{\text{IN}}(x_i) = \top$ , then eventually every honest party  $P_j$  will have justified output  $y \in \{0, 1\}$  where  $J_{\text{OUT}}(y) = \top$ .*

**Justified Agreement** *For all possible justified outputs  $y$  and  $y'$ :  $y = y'$ .*

**Justified Validity** *If  $J_{\text{OUT}}(y) = \top$ , then  $J_{\text{IN}}(y) = \top$ .*

The protocol  $\Pi_{\text{VABA}}$  in Fig. 3 satisfies Definition 5 except with probability  $2^{-\lambda}$ . As mentioned the parties first expand to  $2^{\lambda+1} + 1$  grades. Then  $\Pi_{\text{CC}}$  is run to obtain output  $c$  and the grades are compared with  $2c$ . This mitigates the security loss caused by the adversary being free to choose between more than two grades at the time the coin is flipped. The adversary can still learn  $c$  when the first honest party  $P_i$  gets their grade  $g_i$  and then decide to give some  $P_j$  grade  $g_j \in \{g_i - 1, g_i, g_i + 1\}$  based on the value of  $c$ . But note that since  $2c$  is even, agreement can only be broken if there are justified grades  $g$  and  $g' = g + 1$  where  $g$  is even. So even if the adversary has the ability to choose between  $g_i - 1$  and  $g_i + 1$  becoming justified grades after learning the value of  $c$ . At most one of the two grades can result in conflicting bit decisions, so the adversary has no chance of breaking agreement beyond guessing the exact value of  $c$  before any honest party initiates  $\Pi_{\text{CC}}$ .

**Theorem 1.**  *$\Pi_{\text{VABA}}$  (Fig. 3) is a secure VABA as defined in Definition 5 except with probability  $2^{-\lambda}$ . It uses  $\lambda + 2$  rounds of communication and  $O(\lambda n^2(\lambda + |J_{\text{IN}}|))$  bits of communication, where  $|J_{\text{IN}}|$  is the size of the input justifier.*



**Validated Asynchronous BA protocol  $\Pi_{\text{VABA}}$**

- On input  $b_i$ ,  $P_i$  initiates  $\Pi_{\text{VAPROX}-(2^{\lambda+1}+1)}$  with input  $b_i$ .
- On output  $g_i$  from  $\Pi_{\text{VAPROX}-(2^{\lambda+1}+1)}$ ,  $P_i$  initiates  $\Pi_{\text{CC}}$ .
- On output  $c_i$  from  $\Pi_{\text{CC}}$ ,  $P_i$  lets  $d_i = 1$  if  $g_i > 2c_i$ .

**Fig. 3.** A validated BA using expand-and-extract.

*Proof.* Since we expand to  $2^{\lambda+1} + 1$  grades and double the value of the coin: even if the adversary can choose between 3 different grades when the coin is leaked, the adversary needs to guess the value of the coin to split agreement. Justified validity reduces to justified validity of validated proxconsensus: If  $b$  is the only justified input, then the only justified possible justified grade is  $b \cdot (G - 1)$  and the only possible justified output is  $b$ .

## 4 Binary Agreement

We finally solve BA using the recipe from [CKPS01] where a justifier is formed using a threshold signature before running VABA. This resulting BA has a slightly stronger security definition than usual: the output is justified and all possible justified outputs are identical.

**Binary Agreement protocol  $\Pi_{\text{BA}}$ .**

- On input  $b_i$ ,  $P_i$  initiates  $\Pi_{\text{IVG}}$  with input  $b_i$ .
- On output  $x_i$  from  $\Pi_{\text{IVG}}$ ,  $P_i$  initiates  $\Pi_{\text{VABA}}(\Pi_{\text{IVG}}.J_{\text{OUT}})$  with input  $x_i$ .
- On output  $y_i$  from  $d_i \Pi_{\text{VABA}}(\Pi_{\text{IVG}}.J_{\text{OUT}})$ ,  $P_i$  outputs  $d_i$ .

**Fig. 4.** A Binary Agreement protocol.

**Definition 6 (Asynchronous BA).** Let  $\Pi_{\text{BA}}$  be a protocol for  $n$  parties outputting  $y \in \{0, 1\}$  satisfying an output justifier  $J_{\text{OUT}}$ . We say that  $\Pi_{\text{BA}}$  is a secure BA protocol if the following properties hold:

**Liveness** If every honest party  $P_i$  has input  $x_i \in \{0, 1\}$ , then eventually every honest party  $P_j$  will have justified output  $y_i \in \{0, 1\}$  where  $J_{\text{OUT}}(y_i) = \top$ .

**Justified Agreement** For all possible justified outputs  $y$  and  $y'$ :  $y = y'$ .

**Validity** If  $J_{\text{OUT}}(y) = \top$ , then some honest party gave input  $y$ .

We give a protocol  $\Pi_{\text{BA}}$  in Fig. 4 which implements BA except with probability  $2^{-\lambda}$ . The security follows as a corollary from Theorem 1.

**Corollary 2.**  $\Pi_{\text{BA}}$  given in Fig. 4 implements a secure BA as defined in Definition 6 except with probability  $2^{-\lambda}$ . It uses  $\lambda + 3$  rounds of communication and  $O((n\lambda)^2)$  bits of communication.

**Acknowledgements** I would like to thank Jesper Buus Nielsen for many helpful discussions.

## References

- AC10. Hagit Attiya and Keren Censor-Hillel. Lower bounds for randomized consensus under a weak adversary. *SIAM J. Comput.*, 39(8):3885–3904, 2010.
- BBB<sup>+</sup>24. Akhil Bandarupalli, Adithya Bhat, Saurabh Bagchi, Aniket Kate, Chen-Da Liu-Zhang, and Michael K. Reiter. Delphi: Efficient asynchronous approximate agreement for distributed oracles. In *DSN*, pages 456–469. IEEE, 2024.
- CFF<sup>+</sup>05. Jeffrey Considine, Matthias Fitzi, Matthew K. Franklin, Leonid A. Levin, Ueli M. Maurer, and David Metcalf. Byzantine agreement given partial broadcast. *J. Cryptol.*, 18(3):191–217, 2005.
- CKPS01. Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 524–541, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Berlin, Heidelberg, Germany.
- CKS05. Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in Constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, July 2005.
- DYMM<sup>+</sup>20. Thomas Dinsdale-Young, Bernardo Magri, Christian Matt, Jesper Buus Nielsen, and Daniel Tschudi. Afgjort: A partially synchronous finality layer for blockchains. In Clemente Galdi and Vladimir Kolesnikov, editors, *SCN 20*, volume 12238 of *LNCS*, pages 24–44, Amalfi, Italy, September 14–16, 2020. Springer, Cham, Switzerland.
- EW24. Mose Mizrahi Erbes and Roger Wattenhofer. Asynchronous approximate agreement with quadratic communication, 2024.
- FLZL21. Matthias Fitzi, Chen-Da Liu-Zhang, and Julian Loss. A new way to achieve round-efficient byzantine agreement. In Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen, editors, *40th ACM PODC*, pages 355–362, Virtual Event, Italy, July 26–30, 2021. ACM.
- FM97. Pesech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM J. Comput.*, 26(4):873–933, 1997.
- GGLZ22. Diana Ghinea, Vipul Goyal, and Chen-Da Liu-Zhang. Round-optimal byzantine agreement. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 96–119, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland.
- KN23. Simon Holmgaard Kamp and Jesper Buus Nielsen. Byzantine agreement decomposed: Honest majority asynchronous total-order broadcast from reliable broadcast. *IACR Cryptol. ePrint Arch.*, page 1738, 2023.
- KNTT22. Simon Holmgaard Kamp, Jesper Buus Nielsen, Søren Eller Thomsen, and Daniel Tschudi. Enig: Player replaceable finality layers with optimal validity. *Cryptology ePrint Archive*, Report 2022/201, 2022.
- LN24. Julian Loss and Jesper Buus Nielsen. Early stopping for any number of corruptions. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part III*, volume 14653 of *LNCS*, pages 457–488, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland.