

Distributed Randomness from Approximate Agreement

Luciano Freitas  

LTCI, Télécom Paris, Institut Polytechnique de Paris, France

Petr Kuznetsov  

LTCI, Télécom Paris, Institut Polytechnique de Paris, France

Andrei Tonkikh  

LTCI, Télécom Paris, Institut Polytechnique de Paris, France

Abstract

Randomisation is a critical tool in designing distributed systems. The *common coin* primitive, enabling the system members to agree on an unpredictable random number, has proven to be particularly useful. We observe, however, that it is impossible to implement a truly random common coin protocol in a fault-prone asynchronous system.

To circumvent this impossibility, we introduce two relaxations of the perfect common coin: (1) *approximate common coin* generating random numbers that are *close* to each other; and (2) *Monte Carlo common coin* generating a common random number with an arbitrarily small, but non-zero, probability of failure. Building atop the *approximate agreement* primitive, we obtain efficient asynchronous implementations of the two abstractions, tolerating up to one third of Byzantine processes. Our protocols do not assume trusted setup or public key infrastructure and converge to the perfect coin exponentially fast in the protocol running time.

By plugging one of our protocols for *Monte Carlo common coin* in a well-known consensus algorithm, we manage to get a *binary Byzantine agreement* protocol with $O(n^3 \log n)$ communication complexity, resilient against an adaptive adversary, and tolerating the optimal number $f < n/3$ of failures without trusted setup or PKI. To the best of our knowledge, the best communication complexity for binary Byzantine agreement achieved so far in this setting is $O(n^4)$. We also show how the *approximate common coin*, combined with a variant of Gray code, can be used to solve an interesting problem of Intersecting Random Subsets, which we introduce in this paper.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Asynchronous, approximate agreement, weak common coin, consensus, Byzantine agreement

Digital Object Identifier 10.4230/LIPIcs.DISC.2022.24

Related Version *Full Version*: <https://arxiv.org/abs/2205.11878> [18]

Funding *Luciano Freitas*: Nomadic Labs.

Petr Kuznetsov: TrustShare Innovation Chair.

Andrei Tonkikh: TrustShare Innovation Chair.

1 Introduction

Generating randomness in distributed systems is an essential part of many protocols, such as Byzantine Agreement [4], Distributed Key Generation [21] or Leader Election [35]. Any application that needs an unpredictable or unbiased result will most likely rely on randomness. Although sometimes local sources of randomness are enough for some protocols [34], having access to a common random number can guarantee faster termination [3]. Producing a common unpredictable random number has been extensively studied in the literature on cryptography and distributed systems under the names of *random beacon*, distributed (multi-party) *random number generation* or *common coin* (even if the result is not binary). In essence, these protocols ensure:



© Luciano Freitas, Petr Kuznetsov, and Andrei Tonkikh;
licensed under Creative Commons License CC-BY 4.0
36th International Symposium on Distributed Computing (DISC 2022).
Editor: Christian Scheideler; Article No. 24; pp. 24:1–24:21



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Termination: every correct process eventually outputs some value;

Agreement: no two correct processes output different values;

Randomness: the value output by a correct process must be uniformly distributed over some domain \mathcal{D} , $|\mathcal{D}| \geq 2$.

We call a protocol that ensures the three properties (*Termination*, *Agreement*, and *Randomness*) a **perfect common coin**. There are many message-passing protocols without trusted setup that implement a perfect common coin in the presence of Byzantine adversary [40, 11, 25, 9, 7, 29, 19]. These protocols are either *synchronous*, meaning that every message sent by a correct process is delivered within a certain (known *a priori*) bound of time, or *partially-synchronous*, meaning that such a bound exists but is unknown.

In contrast, one can also consider an *asynchronous* system, where no bounds on communication delays can be assumed. In a seminal work of Fischer, Lynch, and Paterson, it has been shown that the problem of *consensus* has no asynchronous fault-tolerant solutions [17]. As we show in Appendix A, this impossibility also holds for perfect common coins: no algorithm can implement a perfect common coin in a message-passing asynchronous system where at least one process might crash. Note that this statement cannot be proven by a simple black-box reduction from consensus to a perfect common coin and a reference to FLP [17]. Indeed, if such a reduction existed, the resulting protocol would have to always terminate in a bounded number of steps, even with unfavourable outputs of the black-box common coin. Hence, if we were to replace the common coin protocol by a protocol that always returns 0, it would still provide termination as well as all other properties of consensus, violating [17].

Note that this impossibility applies even to systems with *trusted setup*, such as the one assumed in [9]. Such protocols typically do not satisfy the Randomness property of a perfect common coin. The outputs of these protocols follow deterministically from the information received by the processes during the setup.

In light of the impossibility of a perfect coin, one might look for *relaxed* versions of the common-coin problem that allow asynchronous fault-tolerant solutions. For example, one can relax the Agreement property by only requiring the output to be common with some constant probability, which results in an abstraction sometimes called *weak common coin*¹ or *δ -matching common coin*. In this paper, we call this abstraction a **probabilistic common coin**, in order to avoid confusion with other relaxations we introduce. More precisely, probabilistic common coins replace the Agreement property above with *probabilistic δ -consistency*.

Probabilistic δ -Consistency: with probability at least δ , no two correct processes output different values.

We also introduce the concept of a *Monte Carlo common coin* which is a probabilistic common coin whose success rate δ can be parameterized as follows: the more rounds of the protocol are executed, the more reliable the outcome is. In our case δ starts at $\frac{2}{3}$ in the first round of the protocol and converges to 1 at an exponential rate in the number of rounds. In most probabilistic common coins, δ could be increased by decreasing the resilience level (the allowed fraction of Byzantine processes). However, to the best of our knowledge, this paper is the first to present an implementation of a Monte Carlo common coin that can achieve arbitrarily small (but non-zero) δ by increasing the running time of the protocol (Sections 5 and 6).

¹ Many weak common coin protocols such as the one in [10] also relax Randomness.

We also propose a novel, alternative relaxation of Agreement: instead of ensuring that the *same* output is produced (with some probability), we may require that the produced outputs are *close* to each other according to some metric. For this variant, we have to also slightly relax randomness so that only one correct process is guaranteed to obtain a truly random value. More precisely, assume a discrete range of possible outputs $[0..D-1]$, and let $d_q(x, y)$ denote the distance between x and y in the algebraic ring \mathbb{Z}_q .² The **Approximate common coin** abstraction then satisfies Termination and the following two properties:

Approximate ϵ -Consistency: if one correct process outputs value x and another correct process outputs y , then $d_D(x, y) \leq \lceil \epsilon D \rceil$, for a given parameter $\epsilon \in (0, 1]$;

One Process Randomness: the value output by *at least one* correct process must be uniformly distributed over the domain $[0..D-1]$.³

Our implementations of Monte Carlo and Approximate common coins build upon the abstraction of *Approximate Agreement* [15]. It appears that the abstraction perfectly matches the requirements exposed by our relaxed common-coin definitions: it naturally grasps the notion of outputs being close where the precision can be related to the execution time. Building atop existing asynchronous Byzantine fault-tolerant implementations [15, 1], we introduce and discuss an efficient implementation of the *bundled* version of this abstraction which is, intuitively, equivalent to n parallel instances of Approximate Agreement, but is much more efficient.

We discuss two applications of our protocols. First, we observe that our Monte Carlo common coin can be plugged into many existing Byzantine agreement protocols [6, 10, 13, 33]. This helps us to obtain a *binary Byzantine agreement* protocol with $O(n^3 \lambda \log n)$ communication complexity, where λ is the security parameter. The protocol exhibits optimal resilience of $f < n/3$, tolerates adaptive adversary, and assumes no trusted setup or PKI. In this setting, the best prior protocols for binary Byzantine agreement we are aware of have communication complexity of $O(n^4 \lambda)$ [28, 2].

We also introduce *Intersecting Random Subsets*, a new problem that can be used to asynchronously choose random committees with large intersections. Using elements of coding theory, namely Gray Codes [23, 36], we show how our Approximate common coin can be used to solve this problem without additional communication overhead.

We present our model definitions in Section 2 and describe the building blocks used in our constructions in Section 3. We present our protocols in Sections 4–6, provided with implementation details and complexity analysis. In Section 7, we show how the Monte Carlo common coin can be used in solving binary Byzantine agreement. In Section 8, we overview the related work and in Section 9, we conclude the paper.

We provide some of the necessary complementary material in the appendices and we refer the reader to the full version of the paper [18] for the remaining parts. In Appendix A, we prove the impossibility of an asynchronous perfect common coin. Appendix B introduces and shows how to solve the problem of Intersecting random subsets. In addition to the contents of this paper, the full version [18] contains the proof of correctness of our common coin constructions, it contains a modular presentation of the coin proposed by Canetti and Rabin in 1993 [10], and it also presents two implementations of Random Secret Draw, one of the major building blocks of our common coins.

² I.e., $d(x, y) = \min\{|x - y|, q - |x - y|\}$.

³ With a small modification to our protocol, we can easily achieve $(f+1)$ -Process Randomness instead of One Process Randomness. However, we do not know if it is possible to guarantee that all outputs of correct processes are random without relaxing other properties.

2 System Model

We consider a system of n processes able to communicate using reliable communication channels. Among the participants, at most $f < \frac{n}{3}$ are Byzantine and might display arbitrary behaviour.

We assume the *adaptive* adversarial model: up to f Byzantine processes are chosen by the adversary depending on the execution. A non-Byzantine process is called *correct*. The communication complexity of our baseline protocols can be improved by a factor of n using Aggregatable Publicly Verifiable Secret Sharing (APVSS) [24]. However, as we are not aware of APVSS implementations that are secure against the adaptive adversary, the improved protocols can only be proved correct in the presence of the static one.

The adversary can control the time the messages sent by correct processes take to arrive, as well as reorder them. However, it cannot drop a message sent by a correct process unless it corrupts this process before the message has arrived.

We assume that each process has access to a local random number generator that can be accessed as follows:

RandomInt(D): produces a uniformly distributed random integer number in the range $[0..D-1]$.

The proposed protocols as well as some of the building blocks rely on the use of cryptographic hash functions. The hash of an arbitrary string s is denoted $H(s)$ and has length λ that we call the *security parameter*. It is computationally infeasible to find two strings $s \neq s'$ such that $H(s) = H(s')$, as well as inverting a hash without knowing which input was used a priori.

We assume a computationally bounded adversary, so that it is incapable of breaking cryptographic primitives with all but negligible probability. However, since such a probability exists, we allow the properties of all our protocols as well as all building blocks to be violated with a negligible in λ probability.

3 Building Blocks

Our protocols make use of a wide range of building blocks. None of them are completely new, but some of them are modified according to our needs. In particular, we introduce the *Random Secret Draw* abstraction inspired by the ideas from [10] and [2] (Section 3.3). We also provide a *bundled* version of the *Approximate Agreement* [15, 1] abstraction (Section 3.5). In addition, we use *Byzantine Reliable Broadcast* [6, 14] (Section 3.1), *Asynchronous Verifiable Secret Sharing* [8, 14] (Section 3.2), and *Gather* [10, 39] (Section 3.4).

3.1 Byzantine Reliable Broadcast

Byzantine Reliable Broadcast (BRB) [6] allows a designated leader to communicate a single message to all processes in such a way that, if any correct process delivers a message, then every other correct process eventually delivers exactly the same message (even if the leader is Byzantine). More precisely, a BRB protocol must satisfy the following properties:

Validity: if the leader is correct and it broadcasts message m , then every correct process will eventually deliver m ;

Consistency: if two correct processes j and k deliver messages m_j and m_k , then $m_j = m_k$.

Totality: if a correct process j delivers some message m , then eventually all correct processes will deliver m .

The performance of reliable broadcast is of crucial importance to our protocols. We believe that the BRB implementation recently proposed by Das, Xiang, and Ren [14] will be the most suitable option. It has total communication complexity of just $O(n|M| + n^2\lambda)$, where $|M|$ is the size of the message and λ is the security parameter, total message complexity of $O(n^2)$, and the latency of 3 message delays in case of a correct leader and 4 message delays in case of a Byzantine leader.

In this paper, we always use BRB in groups of n instances, with each process being the leader of one. We use the following notation:

BRB_{*i*}.Broadcast(*m*): allows process i to broadcast a message in an instance of BRB where i is the leader;

BRB_{*i*}.Deliver(*m*): an event indicating that message m from process i has been delivered.

3.2 Asynchronous Verifiable Secret Sharing

Asynchronous Verifiable Secret Sharing (AVSS) [8] allows a process to securely share information with other participants and to keep its contents secret until the moment a threshold of participants agree to open it.

In our protocols, AVSS is used with the following interface:

AVSS_{*i*}.ShareSecret(*x*): allows process i to share a secret x among the participants;

AVSS_{*i*}.SharingComplete(): an event issued when a secret is correctly shared by process i ;

AVSS_{*i*}.EnableRetrieve(): enables responses to retrieval requests;

AVSS_{*i*}.Retrieve(): returns x if it was previously shared and all correct processes invoked **AVSS_{*i*}.EnableRetrieve()**.

An AVSS implementation must satisfy the following properties:

Validity: if a correct process i invokes **AVSS_{*i*}.ShareSecret(*x*)**, then every correct process eventually receives the **AVSS_{*i*}.SharingComplete()** event and no value other than x can be returned from the **AVSS_{*i*}.Retrieve()** operation invoked by a correct process;

Notification Totality: if one correct process receives the **AVSS_{*i*}.SharingComplete()** event, then every correct process eventually receives it;

Retrieve Termination: if all correct processes invoke **AVSS_{*i*}.EnableRetrieve()** and any correct process invokes **AVSS_{*i*}.Retrieve()**, then this operation will eventually terminate and the process will obtain the shared secret;

Binding: if some correct process receives the **AVSS_{*i*}.SharingComplete()** notification, then there exists a fixed secret x such that no value other than x can be returned from the **AVSS_{*i*}.Retrieve()** operation invoked by a correct process;

Secrecy: if process i is correct and no correct process invoked **AVSS_{*i*}.EnableRetrieve()**, then the adversary has no information about the secret shared by i .

Das, Xiang, and Ren [14] proposed an AVSS protocol with quadratic communication complexity, constant latency, and without assuming trusted setup. Notice that in order to secretly share a long string s , it is better to follow the method proposed in [30]: encrypt s using a much shorter secret key sym , reliably broadcast the encrypted value $\{s\}_{sym}$ and then perform secret sharing of the key sym . Thus, we shall assume that the total communication complexity of secret sharing of string s is $O(n|s| + n^2\lambda)$.

3.3 Random Secret Draw

One of the key ideas of the weak common coin protocol of Canetti and Rabin [10] is to *assign* each process a random number in a given domain $[0..D-1]$ in such a way that:

Assignment Termination: if a correct process i participates, then it is eventually assigned a value. Moreover, everyone will eventually receive a notification that i has been assigned a random value;

Notification Totality: if process i receives a notification that some process j has been assigned a value, then every correct process will eventually receive such a notification;

Randomness: the assigned numbers are independent and distributed uniformly over the domain $[0..D-1]$. The distribution of the value assigned to process j cannot be affected by the adversary even if j itself is Byzantine;

Unpredictability: until at least one correct process agrees to reveal the assigned values, the value assigned to each process j remains secret, even to process j itself;

Retrieve Termination: if all correct processes invoke `EnableRetrieve()` and any correct process invokes `RetrieveValue(j)`, having received `ValueAssigned(j)`, then this operation eventually returns a value.

Although this idea has been widely used as part of the implementation of asynchronous consensus protocols, to the best of our knowledge, it was never considered a separate primitive and assigned a name. Hence, we shall call it *Random Secret Draw (RSD)*.

This abstraction resembles a well known concept of a Verifiable Random Function (VRF) [32]. However, the important difference is that process j itself cannot know the value it is assigned until the reveal phase. Hence, a Byzantine process cannot choose whether it wants to participate or not based on the random value it is assigned. Moreover, unlike Random Secret Draw, VRF schemes typically require a seed chosen at random *after* the process chose the public key for its pseudo-random function. In fact, a variant of RSD has been recently used to generate such seeds [20].

We use the following interface for the RSD abstraction:

RSD.Start(): allows a process to start participating in RSD and, eventually, to be assigned a random number. We assume that this function is non-blocking, i.e., that an invocation of this function terminates after 0 message delays;

RSD.EnableRetrieve(): used by the processes to start participating in the process of reconstructing the assigned values;

RSD.RetrieveValue(j): returns the value assigned to process j if all correct processes invoked `RSD.EnableRetrieve()` and process j has been assigned some value.

The original RSD implementation by Canetti and Rabin [10] used n^2 instances of AVSS. To the best of our knowledge, to this day, there is no known AVSS protocol that would allow to do it with less than $\Omega(n^4)$ bits of communication in total. Hence, in the full version of the paper [18], we give two possible implementations of RSD. The first one is secure against an adaptive adversary and does not rely on PKI, while the second one uses the implementation from [2] that relies on *Aggregatable Publicly Verifiable Secret Sharing* instead of AVSS. While saving a linear factor in communication complexity, this solution lacks security against adaptive adversary and requires PKI. Since both [10] and [2] did not consider RSD as a separate abstraction and did not provide separate pseudocode for it, in [18], we present both RSD implementations.

3.4 Gather

Yet another important contribution made by Canetti and Rabin in their weak common coin construction [10] is a multi-broadcast protocol that has been recently given the name *Gather* [39, 2]. In this protocol, every process starts by broadcasting a single message through Byzantine Reliable Broadcast. The processes then do a few more rounds of message exchanges

and, in the end, each participant i outputs a set of process ids \mathcal{S}_i such that for all $j \in \mathcal{S}_i$: i has received the message of j through reliable broadcast.⁴ Moreover, the sets output by correct processes satisfy a strong intersection property:

Binding Common Core: There exists a set \mathcal{S}^* of process ids of size at least $n - f$, called the *common core*, such that for every correct process i : $\mathcal{S}^* \subseteq \mathcal{S}_i$. Moreover, once the first correct process outputs, \mathcal{S}^* is fixed and the adversary cannot manipulate it anymore.

The fact that the adversary cannot affect the common core once a single correct process outputs will be important in our protocols. The adversary should not be able to choose the common core based on the generated random numbers after some of the correct processes invoked `EnableRetrieve`.

We slightly generalize the interface of `Gather` by using it in conjunction with BRB, but also with other similar primitives (in particular, AVSS and RSD) and their combinations. When a process invokes `GATHER`, it passes to it an arbitrary callable function `GatherAccept` that takes a process id j and returns *true* if the message from this process is considered to be delivered (not necessarily through BRB). We assume that `Gather` exports the following interface:

Gather.Start(GatherAccept): allows a process to start participating in the `Gather` protocol;
Gather.DeliverSet(S): provides the output of the `Gather` protocol.

In order for the protocol to terminate, the `GatherAccept` function has to satisfy properties similar to those of reliable broadcast.

Accept Validity: if a correct process i invoked `GATHER.Start`, then for every correct process j , `GatherAccept(i)` invoked by process j must eventually return *true*. Moreover, for all i , once `GatherAccept(i)` returned to *true* to some correct process, it must never switch back to *false*;

Accept Totality: if `GatherAccept(i)` invoked by one correct process returned *true*, then eventually it must return *true* to all correct processes.

Thanks to the properties of AVSS and Random Secret Draw (in particular, to the Notification Totality property), in our protocols, this assumption is trivially satisfied. For `Gather`, we use the original protocol of [10] (to the best of our knowledge, it was first described as a separate primitive in [39]).

3.5 Bundled Approximate Agreement

The last building block that we shall need is *Approximate Agreement (AA)* [16]. In a (one-dimensional) AA instance, the processes propose inputs and produce outputs (real values) so that the following properties are satisfied:

Validity: the outputs of correct processes must be in the range of inputs of correct processes.

Approximate $\tilde{\epsilon}$ -consistency: the values decided by non-faulty processes must be at most a distance $\tilde{\epsilon}$ apart from each other.⁵

Termination: every non-faulty process eventually decides.

In our algorithms, Approximate Agreement is always executed in *bundles* of n parallel instances. For the sake of efficiency, one can treat it as a bundled problem with an input vector of size n , corresponding to the different instances and then, for every message, send

⁴ In [39] and [2], `Gather` returns a set of pairs $(id, value)$. However, for our purposes, working with sets of ids is more convenient. The values will be delivered through normal `BRB.Deliver` event.

⁵ We use $\tilde{\epsilon}$ to distinguish it from ϵ used in the definition of Approximate Common Coin.

24:8 Distributed Randomness from Approximate Agreement

information about all instances at the same time, but treat them separately as before. We call this abstraction *Bundled Approximate Agreement* (BAA). BAA **should not be confused** with *Multidimensional Approximate Agreement* [31], which is a stronger abstraction than the one we rely upon.

Assuming binary inputs, the processes access BAA via the following interface:

BAA.Run($[x_1, x_2, \dots, x_n]$): Launches n instances of Approximate Agreement protocol, where the input for the i -th instance is x_i . For a given parameter $\tilde{\epsilon}$, the protocol is executed until $\tilde{\epsilon}$ -approximation is satisfied in every instance and then returns a vector of outputs $[y_1, y_2, \dots, y_n]$.

For implementing BAA, we suggest using the Approximate Agreement protocol proposed in [1] with resilience $f < \frac{n}{3}$. Since in our protocols, the inputs are either 0 or 1, we do not need the termination detection techniques described in [1] neither do we need the “init” phase of the protocol. With the aforementioned *BRB* and some trivial changes⁶, this implementation will give us the communication complexity of $O(n^3\lambda)$ and latency $4 \cdot \log_2(1/\tilde{\epsilon})$.

4 Approximate Common Coin

■ **Algorithm 1** Approximate common coin.

```

1: Instance parameters: domain  $D$ , precision  $\epsilon$ 

2: Distributed objects:
3:    $\forall j \in [n] : \text{AVSS}_j$  – instance of Asynchronous Verifiable Secret Sharing with leader  $j$ 
4:   GATHER – instance of Gather
5:   BAA – instance of Bundled Approximate Agreement with precision  $\tilde{\epsilon} = \epsilon/f$ 

6: function GatherAccept( $j$ )
7:   return true if received  $\text{AVSS}_j.\text{SharingComplete}()$ 

8: operation Toss() returns integer
9:    $x = \text{RandomInt}(D)$ 
10:   $\text{AVSS}_i.\text{ShareSecret}(x)$ 
11:  GATHER.Start(GatherAccept)

12: wait for GATHER.DeliverSet( $S$ )
13:   $\forall j \in [n] : \text{let } w_j = \begin{cases} 1, & j \in S, \\ 0, & \text{otherwise} \end{cases}$ 
14:   $[w'_1, \dots, w'_n] = \text{BAA.Run}([w_1, \dots, w_n])$ 
15:   $\forall j \in [n] : \text{AVSS}_j.\text{EnableRetrieve}()$  // only after BAA completes
16:   $\forall j \in [n] : \text{let } x_j = \begin{cases} \text{AVSS}_j.\text{Retrieve}(), & \text{if } w'_j \neq 0 \\ 0, & \text{otherwise} \end{cases}$ 

17:  // Compute and return the final random number
18:  return  $\left( \left[ \sum_{j \in [n]} x_j \cdot w'_j \right] \bmod D \right)$ 

```

The main idea of this protocol is to aggregate numbers locally generated by enough different processes so that at least one of them is correct and the number it generated is truly random and uniformly distributed. With a good aggregation function, the resulting value will also be uniformly distributed. An example of such an aggregation function is addition modulo the size of the domain. Indeed, it is easy to see that, if x is uniformly distributed

⁶ In the “report” messages, hashes of the values should be sent instead of the values themselves.

over $[0..q-1]$ and y is any number chosen independently of x , then $(x + y) \bmod q$ will also be uniformly distributed over $[0..q-1]$. Another example of an aggregation operation that satisfies a similar property is bit-wise xor (as long as the domain size is a power of 2).

However, without being able to solve consensus, we cannot just elect $f + 1$ or $n - f$ processes from whom we shall take these values. Thankfully, unlike xor, addition has one more useful property: it is continuous. If we take two numbers, x and y , such that $d_q(x, y) \leq \alpha$, then for any number z , $d_q(z + x, z + y) \leq \alpha$.⁷ Hence, a natural idea is to *approximately* elect the set of processes to provide the random inputs.

More precisely, in order to produce an *approximate common coin* in the range $[0..D-1]$, each process locally generates and secret-shares a random number in this range. Then each process gathers a set of ids of processes that have completed the sharing (line 12).

The next step is to create a binary vector with n positions, where each position j is set to 1 if and only if j is present in the gathered set (line 13). This vector is then used as an input for the BAA protocol (line 14), which outputs a vector W of *weights* such that for each position j , the weights received by different processes are at most $\tilde{\epsilon}$ apart.

The value of $\tilde{\epsilon} = \epsilon/f$ is chosen such that the final outputs of the coin are at most ϵ apart from each other. For the details on how this particular value was computed, see [18].

Recall that BAA ensures that the output values lie within the range of inputs of correct processes. Moreover, by the properties of Gather and AVSS, if at least one correct process has j in its gathered set, then j has correctly shared its value and it can be later retrieved by the correct processes. Therefore if the j -th value is irretrievable, the j -th component will always be assigned weight 0. On the other hand, due to the common core property of Gather, at least $n - f$ values will have weight 1, which guarantees that the result is uniformly distributed in the desired range.

Finally, the processes reveal the secrets (lines 15 and 16) and compute the resulting random number.

► **Theorem 1.** *Algorithm 1 implements an approximate common coin.*

The proof of Theorem 1 is presented in the full version of this paper [18].

Complexity analysis. The communication complexity of our *approximate common coin* can be broken down into:

1. n instances of AVSS.ShareSecret in parallel $\Rightarrow O(n^3\lambda)$;
2. One instance of GATHER $\Rightarrow O(n^3\lambda)$;
3. One instance of BAA with $\tilde{\epsilon} = \epsilon/f \Rightarrow O(n^3\lambda(\log f + \log \frac{1}{\epsilon}))$;
4. n instances of AVSS.Retrieve in parallel $\Rightarrow O(n^3\lambda)$;

Hence, the total communication complexity is $O(n^3\lambda(\log f + \log \frac{1}{\epsilon}))$ with Bundled Approximate Agreement being the bottleneck.

The time complexity of the protocol is $O(\log f + \log \frac{1}{\epsilon})$.

5 Monte Carlo Common Coin from Approximate Common Coin

In this section, we present a simple reduction from an approximate common coin to a Monte Carlo common coin. The very short pseudocode is in Algorithm 2.

⁷ Recall that $d_q(x, y)$ is the distance between x and y in the ring \mathbb{Z}_q , i.e., $d_q(x, y) = \min\{|x - y|, q - |x - y|\}$.

24:10 Distributed Randomness from Approximate Agreement

■ **Algorithm 2** Monte Carlo Common Coin from Approximate Common Coin, code for process i .

19: **Instance parameters:** domain size D , success probability δ

20: let $k = \lceil \frac{2}{1-\delta} \rceil$

21: **Distributed objects:**

22: AC – instance of approximate common coin with domain size kD and precision $\epsilon = \frac{1}{kD}$

23: **operation** Toss() **returns** integer

24: **return** $\left\lfloor \frac{\text{AC.Toss()}}{k} \right\rfloor$

The transformation requires first to generate an approximate common coin of domain kD where k is an integer number $\lceil \frac{2}{1-\delta} \rceil$ and $\epsilon = \frac{1}{kD}$. This implies that different processes shall get values at most $\lceil \frac{1}{kD} \cdot kD \rceil = 1$ apart.

The domain of the approximate coin is k times larger than the domain of the targeted Monte Carlo common coin. By dividing the result by k , we get the desired range of values and success probability δ , where k values of the *approximate common coin* are mapped to one value of *Monte Carlo common coin*.

► **Theorem 2.** *Algorithm 2 implements a Monte Carlo common coin with domain D and success probability δ .*

Proof. Termination and Unpredictability follow from the properties of *approximate common coin*, while Randomness follows from One Process Randomness since exactly k values from the larger domain (kD) are mapped to each value in the smaller domain (D). Let x' be the resulting toss of the first correct process that completes BAA in its approximate common coin toss. Then, as established, every other process will be at a distance at most 1 from it. Hence, if the remainder of the division of x' by k is neither 0 nor $k - 1$, every correct process decides the same value:

$$\Pr[\text{failure}] \leq \frac{2}{k} = \frac{2}{\lceil \frac{2}{1-\delta} \rceil} \leq 1 - \delta. \quad \blacktriangleleft$$

Complexity analysis. This protocol runs a single instance of an Approximate Common Coin, with precision $\epsilon = \frac{1}{D \lceil \frac{2}{1-\delta} \rceil}$. If used with our algorithm from Section 4, it will take $O(\log f + \log \frac{1}{\epsilon}) = O(\log f + \log D + \log \frac{1}{1-\delta})$ rounds of approximate agreement.

Hence, the overall time complexity of the protocol is $O(\log f + \log D + \log \frac{1}{1-\delta})$ and the communication complexity is $O(n^3 \lambda(\log f + \log D + \log \frac{1}{1-\delta}))$.

6 Direct Implementation of Monte Carlo Common Coin

Overview. The main idea of this protocol is to assign to each process a random value and a random ticket. Then, using approximate agreement, the protocol is able to select the process with maximum ticket with adjustable probability of success and adopt the value corresponding to this ticket as the coin value.

Tickets and values. The protocol first assigns random values and random tickets to each participant, maintaining both secret until later (lines 37 and 38). Processes then gather a list of participants who have both drawn a ticket and a value, guaranteeing that all processes will hold sets that all intersect in at least $n - f$ participants (line 39).

■ **Algorithm 3** Monte Carlo Common Coin, code for process i .

```

25: Instance parameters: domain size  $D$ , success probability  $\delta$ , security parameter  $\lambda$ 

26: Functions:
27:   Calibrate( $w$ ) – returns the weight to apply to a ticket given that BAA returned  $w$ 

28: Distributed objects:
29:   TICKETDRAW – instance of Random Secret Draw with domain size  $2^\lambda$ 
30:   VALUEDRAW – instance of Random Secret Draw with domain size  $D$ 
31:   GATHER – instance of Gather
32:   BAA – instance of Bundled Approximate Agreement with precision  $\tilde{\epsilon}$ ,
33:     where  $\tilde{\epsilon}$  depends on the calibration function (see “Weight calibration” below)

34: function GatherAccept( $j$ ) returns boolean
35:   return true iff received both TICKETDRAW.ValueAssigned( $j$ ) and VALUEDRAW.ValueAssigned( $j$ )

36: operation Toss() returns integer
37:   TICKETDRAW.Start()
38:   VALUEDRAW.Start()
39:   GATHER.Start(GatherAccept)

40:   wait for event GATHER.DeliverSet( $S$ )
41:    $\forall j \in [n] : \text{let } w_j = \begin{cases} 1, & j \in S, \\ 0, & \text{otherwise} \end{cases}$ 
42:    $[w'_1, \dots, w'_n] = \text{BAA.Run}([w_1, \dots, w_n])$ 
43:    $\text{candidates} = \{j \in [n] \mid w'_j > 0\}$ 

44:   TICKETDRAW.EnableRetrieve()
45:   VALUEDRAW.EnableRetrieve()
46:    $\text{tickets} = \{\text{TICKETDRAW.RetrieveValue}(j) \mid j \in \text{candidates}\}$ 
47:    $\text{values} = \{\text{VALUEDRAW.RetrieveValue}(j) \mid j \in \text{candidates}\}$ 


48:    $\text{winner} = \arg \max_{j \in \text{candidates}} \text{Calibrate}(w'_j) \cdot \text{tickets}[j]$ 
49:   return  $\text{values}[\text{winner}]$ 

```

Approximate Agreement. In a similar manner as in the previous protocol, each process runs Bundled Approximate Agreement inputting 1 in the dimensions corresponding to the processes it has received from Gather, and 0 in the other dimensions (line 42). Similar to the previous protocol, if a process has not made a valid draw it will always be assigned weight zero, whereas if the weight is different than zero then it is possible to recover the secretly drawn number.

Opening the secrets. Prior to the first decision of a correct process in BAA, no secrets are leaked, as the underlying Random Secret Draw abstraction requires at least one correct process to invoke the EnableRetrieve operation before any information about the generated numbers is revealed. After this first decision of a correct process, the secrets can be opened (line 45), but at this point the adversary can only induce other processes deciding values which are at most $\tilde{\epsilon}$ apart from the first decision, which does not undermine the safety of the protocol.

Decision. With the tickets and values now openly available, the processes calibrate the tickets by multiplying the plain ticket by a *calibration function* applied to the weights. The simplest calibration function is an identify function, the calibrated ticket of a process i is simply the product of the output i -th output of BAA and the original ticket t_i . In their final steps, processes decide the value corresponding to the highest calibrated ticket (lines 48 and 49).



■ **Figure 1** The weight calibration function.

Weight calibration. The protocol without weight calibration (i.e., with $\text{Calibrate}(w) = w$) requires $3 + \lceil \log_2(n) + \log_2\left(\frac{1}{1-\delta}\right) \rceil$ rounds of approximate agreement in order to achieve the success probability δ . A similar performance is achieved by the protocol in Section 5.

In order to get rid of the $\log_2(n)$ part in the time complexity, we use a calibration function that is linear on $(0, 1]$ with a discontinuity point at 0, as illustrated in Figure 1. If $w_1 > 0$ and $w_2 > 0$ and $|w_1 - w_2| = \tilde{\epsilon}$, then, after calibration, $|\text{Calibrate}(w_1) - \text{Calibrate}(w_2)| \approx \tilde{\epsilon} \cdot (1 - v)$. This is similar in effect to running extra $\log_2 \frac{1}{1-v}$ rounds of approximate agreement, but at no extra latency cost. We then balance the value of the parameter v in such a way that, intuitively, the discontinuity at 0 is very unlikely to cause disagreement. An example of a good value for v that achieves this goal is $1 - \frac{\ln(2/(1-\delta))}{2n/3}$. In order to achieve success probability δ , $5 + \lceil \log_2\left(\frac{1}{1-\delta}\right) + \log_2\left(\log_2\left(\frac{1}{1-\delta}\right)\right) \rceil$ rounds of approximate agreement are required.

► **Theorem 3.** *Algorithm 3 implements a Monte Carlo common coin.*

The proof of Theorem 3 is presented in the full version of this paper [18].

Complexity analysis. The communication complexity of our *Monte Carlo common coin* can be broken down into:

1. 2 instances of Random Secret Draw $\Rightarrow O(n^3\lambda)$;
2. 1 instance of GATHER $\Rightarrow O(n^3\lambda)$
3. One instance of BAA with $\tilde{\epsilon} = O(1/(1-\delta)) \Rightarrow O(n^3\lambda \log \frac{1}{1-\delta})$;
4. $2n$ secret retrievals in parallel $\Rightarrow O(n^3\lambda)$;

Hence, the total communication complexity is $O(n^3\lambda \log \frac{1}{1-\delta})$ with Bundled Approximate Agreement being the bottleneck. The time complexity of the protocol is $O(\log \frac{1}{1-\delta})$.

7 Applications

Asynchronous Binary Byzantine Agreement. Monte Carlo common coins can be plugged into any Byzantine Agreement protocol that makes a call to a probabilistic common coin, such as [6, 10, 13, 33]. Given any probabilistic common coin with constant success probability, these algorithms terminate in a constant number of rounds, exchanging at most $O(n^3\lambda)$ bits.

Using our Monte Carlo common coin obtained via the transformation from approximate common coin (Section 5), we get a protocol that is secure against an adaptive adversary, assumes no trusted setup or PKI, and exhibits communication complexity $O(n^3\lambda \log n)$ at the expense of extra $O(\log n)$ factor in time complexity (the complexity of coin flips

dominate the rest of the protocol). As far as we know, the best existing setup-free solutions that are resilient against adaptive adversary and tolerate up to $f < n/3$ failures exhibit communication complexity of $O(n^4\lambda)$ [28, 2].

Intersecting Random Subsets. An example of an interesting new application that can be solved with Approximate Common Coins but not probabilistic common coins, with direct consequences into the choice of committees among processes, is given in Appendix B.

8 Related Work

Ben-Or [4] proposed the first randomized consensus algorithm based on the “independent choice” common coin. In this algorithm, every participant tosses a local random coin and with probability 2^{-n} , the values picked by n participants are identical. Bracha [6] extended this algorithm to the Byzantine fault model with $f < n/3$ faulty participants.

Rabin [37] proposed an implementation of a perfect common coin based on Shamir’s secret sharing [38], assuming trusted setup (a trusted dealer distributes *a priori* a large number of secrets). Today, most protocols that allow trusted setup use the solution proposed by Cachin et al. [9] who have described a practical perfect common coin based on threshold pseudorandom functions (tPRF), assuming that a trusted dealer distributes a short tPRF key. These protocols use the pre-distributed randomness in a clever way to obtain multiple numbers that are computationally indistinguishable from random (much like a PRNG [5]). In contrast, our algorithms do not assume trusted setup.

In the setup-free context, Canetti and Rabin [10] proposed a weak common coin algorithm based on asynchronous verifiable secret sharing (AVSS), which resulted in an efficient randomized *binary* consensus. In designing our protocols, we make use of multiple ideas from this work, taking into account recent improvements, such as the use of Aggregatable Publicly Verifiable Secret Sharing [24], suggested in a similar context by Abraham et al. [2].

Using standard PKI cryptography, Cohen et al. [12] built two common coins relying on VRFs. The first one with resilience $f < (1/3 - \epsilon)$ achieves success rate $\delta = \frac{18\epsilon^2 + 24\epsilon - 1}{6(1+6\epsilon)}$. The second coin involves sampling committees among the processes and also guarantees a constant success rate that depends on the system’s resilience.

Kogias et al. [28] proposed a relaxed abstraction called *eventually perfect common coin*. They first build a weak distributed key generator (wDKG) which is a protocol that never terminates: each party outputs a sequence of candidate keys to be used for encryption and decryption with the property that they will eventually agree on a set of keys. This mechanism can replace the trusted setup in [9]. Moreover, it can be shown that the participants may disagree on the set of keys at most $f + 1$ times. The resulting coin eventually terminates with a perfect result, hence its name. In contrast, the challenge of our work was to devise (one-shot) unbiased common coins with provable termination.

Gao et al. [20] combined VRFs and AVSS to produce the first random coin which has $O(n^3\lambda)$ communication complexity. With the advent of new broadcast and APVSS implementations, the classic protocol of [10] gains the same complexity. They created a weak form of *Gather* called *core-set selection* in which $f + 1$ correct participants share at least $n - f$ VRFs coming from different processes. They then use AVSS to build the seeds for VRFs and whenever the highest VRF is in the common core, the nodes successfully agree in the coin outcome. Their protocol assumes the static adversary, but it can be made adaptive with a *relatively weak* form of trusted setup: a single common random number must be published after the public keys of the participants are fixed. In contrast, our protocols do not assume any form of trusted setup. Assuming static adversary, our protocols can achieve the same communication cost while additionally enabling parameterized success rate.

Our Monte Carlo coin from Section 6 was inspired by the Proposal Election protocol recently introduced by Abraham et al. [2]. Technically, it is not a common coin *per se* but it uses elements of it. In this protocol, every party inputs some externally valid value, and with a constant probability, all parties output the same value that was proposed by a non-faulty party chosen at random. Intuitively, the main contribution of the protocol in Section 6 is the use of approximate agreement to *amplify* the success probability.

In the *full information* model, without using cryptography, King and Saia [27] observed that the strength of the Byzantine adversary is in its anonymity, but it cannot bias the coin indefinitely without being detected. Even though their Byzantine agreement algorithm with polynomial expected time does allow the adversary to bias the coin, but amended this with statistical tests aiming at detecting this kind of deviation and evicting misbehaving participants. The resilience level of this algorithm is, however, orders of magnitude lower than n ($f < n/400$ in the best case). Huang et al. [26] recently extended this work to achieve the resilience of $f < n/4$. In contrast, our algorithms use cryptographic tools to produce unbiased (approximate) outputs and maintain optimal resilience $f < n/3$.

Monte Carlo protocols are randomized algorithms that have a fixed number of rounds and yield results that are correct with a given probability, while Las Vegas protocols always give the correct results but do not have a fixed number of rounds. Notice that Las Vegas algorithms must have a fixed probability of terminating every round and thus can be converted into Monte Carlo by stopping after a fixed number of rounds and deciding a random value if termination is not attained.

Such a transformation could be applied to [28], but since their latency is a function of $O(f)$, the resulting Monte Carlo common coin would also have a latency which is a function of f , while our solution is independent of this parameter. Another option would be to create a set of keys using [2] which could be run a fixed number of rounds and then to use [9]. Since their expected number of rounds is not a function of f , this transformation would have the same asymptotic complexity as ours, but it would include many unnecessary message delays from the consensus protocol, from the verifiable gather and other parts of their ADKG that are not present in a direct implementation such as ours.

9 Conclusion

In this paper, we suggest 2 new relations of the common coin primitive implementable in a fully asynchronous environment. We provide efficient implementations based on a range of novel techniques. Our protocols are the first use of approximate agreement to generate random numbers, which is used to keep decided values close in the approximate common coin protocol and to increase the probability of agreement in the direct implementation of the Monte Carlo common coin. Moreover, we also introduced elements of coding theory that were not previously applied to the distributed computing realm in the solution of what we called intersecting random subsets.

Further studies are necessary to explore the full potential of using these new abstractions in the design of distributed protocols and to understand the theoretical limits of their performance.

Tight performance analysis for the Monte Carlo common coin. In the full version of the paper, we proved that in order to achieve success probability δ , our probably common coin protocol requires $3 + \lceil \log_2(n) + \log_2\left(\frac{1}{1-\delta}\right) \rceil$ or $5 + \left\lceil \log_2\left(\frac{1}{1-\delta}\right) + \log_2\left(\log_2\left(\frac{1}{1-\delta}\right)\right) \right\rceil$ rounds of approximate agreement, depending on whether weight calibration is used. While it seems to correctly reflect the asymptotic behaviour of the actual distribution, these bounds seem to be rather pessimistic when only a few rounds of approximate agreement are run.

For example, according to these bounds, with weight calibration, we will need 8 rounds in order to achieve $\delta = \frac{2}{3}$. However, in practice, $\delta = \frac{2}{3}$ is achieved with 0 rounds of approximate agreement as with probability at least $2/3$ a process in the common core will have the largest ticket. As we estimated empirically through simulations, the actual δ that we obtain after running 8 rounds of approximate agreement is around 0.993 instead of $2/3$.

Non-linear calibration function for the Monte Carlo common coin. Weight calibration is necessary to achieve the latency of $O(\log \frac{1}{1-\delta})$ rounds in our Monte Carlo common coin protocol. We chose a concrete linear function because it was relatively simple to analyze (as we could do a reduction to the case without the calibration). However, this function is unlikely to be optimal. The extra $\log_2(\log_2 \frac{1}{1-\delta})$ part in the resulting estimate on the number of rounds of approximate agreement is likely to be due to sub-optimal choice of the weight calibration function.

Approximate common coin without extra $\log_2(f)$ rounds. Using the magic of weight calibration, for Monte Carlo common coin, we managed to achieve $O(\log(1/\epsilon))$ time complexity, which is likely to be optimal. However, our approximate common coin protocol requires $\log_2 f + \log_2(1/\epsilon)$ rounds of approximate agreement and, hence, its time complexity depends on two variables: f and ϵ . In some applications, ϵ may be constant and $\log_2(f)$ can become the bottleneck.

Creating a protocol without this extra delay or proving a $\Omega(\log(f))$ lower bound would be an interesting contribution to the understanding of the approximate common coin abstraction. It would also mean that the transformation from approximate common to Monte Carlo common would result in more efficient coins of the latter type.

References

- 1 Ittai Abraham, Yonatan Amit, and Danny Dolev. Optimal resilience asynchronous approximate agreement. In *8th International Conference on Principles of Distributed Systems (OPODIS 2004)*, OPODIS'04, pages 229–239, Berlin, Heidelberg, 2004. Springer-Verlag. doi:10.1007/11516798_17.
- 2 Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Reaching consensus for asynchronous distributed key generation. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, pages 363–373, 2021.
- 3 Marcos K Aguilera and Sam Toueg. The correctness proof of ben-or's randomized consensus algorithm. *Distributed Computing*, 25(5):371–381, 2012.
- 4 Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract). In *PODC '83: Proceedings of the annual ACM symposium on Principles of distributed computing*, pages 27–30, 1983.
- 5 Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudorandom bits. *SIAM journal on Computing*, 13(4):850–864, 1984.
- 6 Gabriel Bracha. Asynchronous byzantine agreement protocols. *Information and Computation*, 75(2):130–143, 1987.
- 7 Benedikt Bünz, Steven Goldfeder, and Joseph Bonneau. Proofs-of-delay and randomness beacons in ethereum. *IEEE Security and Privacy on the blockchain (IEEE S&B)*, 2017.
- 8 Christian Cachin, Klaus Kursawe, Anna Lysyanskaya, and Reto Strohli. Asynchronous verifiable secret sharing and proactive cryptosystems. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 88–97, 2002.
- 9 Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, 2005.

- 10 Ran Canetti and Tal Rabin. Fast asynchronous byzantine agreement with optimal resilience. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 42–51, 1993.
- 11 Ignacio Cascudo and Bernardo David. Scrape: Scalable randomness attested by public entities. In *International Conference on Applied Cryptography and Network Security*, pages 537–556. Springer, 2017.
- 12 Shir Cohen, Idit Keidar, and Alexander Spiegelman. Not a COINcidence: Sub-Quadratic Asynchronous Byzantine Agreement WHP. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing (DISC 2020)*, volume 179 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.DISC.2020.25.
- 13 Tyler Crain. Two more algorithms for randomized signature-free asynchronous binary byzantine consensus with $t < n/3$ and $o(n^2)$ messages and $o(1)$ round expected termination, 2020. doi:10.48550/ARXIV.2002.08765.
- 14 Sourav Das, Zhuolun Xiang, and Ling Ren. Asynchronous data dissemination and its applications. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 2705–2721, 2021.
- 15 D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl. Reaching approximate agreement in the presence of faults. *Journal of the ACM*, 33(3):499–516, July 1986.
- 16 A.D. Fekete. Asynchronous approximate agreement. *Information and Computation*, 115(1):95–124, 1994. doi:10.1006/inco.1994.1094.
- 17 Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.
- 18 Luciano Freitas, Petr Kuznetsov, and Andrei Tonkikh. Distributed randomness from approximate agreement. *arXiv preprint*, 2022. arXiv:2205.11878.
- 19 Luciano Freitas de Souza, Andrei Tonkikh, Sara Tucci-Piergiovanni, Renaud Sirdey, Oana Stan, Nicolas Quero, and Petr Kuznetsov. Randsolomon: Optimally resilient random number generator with deterministic termination. In *25th International Conference on Principles of Distributed Systems (OPODIS 2021)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.
- 20 Yingzi Gao, Yuan Lu, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Efficient asynchronous byzantine agreement without private setups. *arXiv preprint*, 2021. arXiv:2106.07831.
- 21 Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 295–310. Springer, 1999.
- 22 Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles, SOSP '17*, pages 51–68, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3132747.3132757.
- 23 Frank Gray. Pulse code communication. *United States Patent Number 2632058*, 1953.
- 24 Kobi Gurkan, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Aggregatable distributed key generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 147–176. Springer, 2021.
- 25 Timo Hanke, Mahnush Movahedi, and Dominic Williams. DFINITY technology overview series, consensus system. *CoRR*, abs/1805.04548, 2018. arXiv:1805.04548.
- 26 Shang-En Huang, Seth Pettie, and Leqi Zhu. Byzantine agreement in polynomial time with near-optimal resilience. *arXiv preprint*, 2022. arXiv:2202.13452.
- 27 Valerie King and Jared Saia. Byzantine agreement in expected polynomial time. *J. ACM*, 63(2), March 2016. doi:10.1145/2837019.
- 28 Eleftherios Kokoris Kogias, Dahlia Malkhi, and Alexander Spiegelman. *Asynchronous Distributed Key Generation for Computationally-Secure Randomness, Consensus, and Threshold Signatures*, pages 1751–1767. Association for Computing Machinery, New York, NY, USA, 2020. doi:10.1145/3372297.3423364.

- 29 Mikhail Krasnoselskii, Grigorii Melnikov, and Yury Yanovich. No-dealer: Byzantine fault-tolerant random number generator. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 568–573. IEEE, 2020.
- 30 Hugo Krawczyk. Secret sharing made short. In *Annual international cryptology conference*, pages 136–146. Springer, 1993.
- 31 Hammurabi Mendes and Maurice Herlihy. Multidimensional approximate agreement in byzantine asynchronous systems. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pages 391–400, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2488608.2488657.
- 32 Silvio Micali, Michael Rabin, and Salil Vadhan. Verifiable random functions. In *40th annual symposium on foundations of computer science (cat. No. 99CB37039)*, pages 120–130. IEEE, 1999.
- 33 Achour Mostéfaoui, Hamouma Moumen, and Michel Raynal. Signature-free asynchronous binary byzantine consensus with $t < n/3$, $o(n^2)$ messages, and $o(1)$ expected time. *J. ACM*, 62(4), September 2015. doi:10.1145/2785953.
- 34 Achour Mostefaoui, Matthieu Perrin, and Michel Raynal. A new insight into local coin-based randomized consensus. In *2019 IEEE 24th Pacific Rim International Symposium on Dependable Computing (PRDC)*, pages 207–20709. IEEE, 2019.
- 35 Achour Mostefaoui and Michel Raynal. Leader-based consensus. *Parallel Processing Letters*, 11(01):95–107, 2001.
- 36 Andre Neubauer, Jurgen Freudenberger, and Volker Kuhn. *Coding theory: algorithms, architectures and applications*. John Wiley & Sons, 2007.
- 37 Michael Rabin. Randomized Byzantine generals. In *Proceedings of the 24th Symposium on Foundations of Computer Science*, pages 403–409. IEEE Computer Society Press, November 1983.
- 38 Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- 39 Gilad Stern and Ittai Abraham. Living with asynchrony: the gather protocol. URL: <https://decentralizedthoughts.github.io/2021-03-26-living-with-asynchrony-the-gather-protocol>, 2021. Accessed: 2022-02-12.
- 40 Ewa Syta, Philipp Jovanovic, Eleftherios Kokoris Kogias, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Michael J Fischer, and Bryan Ford. Scalable bias-resistant distributed randomness. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 444–460. Ieee, 2017.

A Impossibility of an Asynchronous Perfect Common Coin

A perfect common coin protocol makes sure that the correct processes agree on a random number taken uniformly over a specific domain.

By reusing the classical arguments of the impossibility of asynchronous consensus [17], we are going to show that no asynchronous perfect common coin protocol may exist, if a single process is allowed to fail by crashing. Recall that we consider protocols in which processes may non-deterministically choose its actions based on local coin tosses.

Formally, a protocol provides each process with an automaton that, given the process' state and an input (a received message and a result of a local random coin toss), produces an output (a finite set of messages to send and/or the application output). We assume that the automaton itself is deterministic, i.e., all non-determinism is delegated to the outcomes of local random coins. A step of process p is therefore a tuple (p, m, r) , where m is the message p receives (can be \perp if no message is received in this step) and r is the outcome of its local random coin.

A *configuration* of the protocol assigns a local state to each process' automaton and a set of messages in transit (we call it the *message buffer*). The *initial configuration* C_{init} assigns the initial local state to each process and assumes that there are no messages in transit. A

step $s = (p, m, r)$ is *applicable to a configuration* C if m is \perp or m is in the message buffer of C . The result of s applied to C is a new configuration $C.s$, where, based on the automaton of p , the local state of p is modified and finitely many messages are added to the message buffer.

A sequence of steps $E = s_1, s_2, \dots$ is applicable to a configuration C_0 if each $s_i, i = 1, 2, \dots$ is applicable to C_{i-1} , where, for all $i \geq 1, C_i = C_{i-1}.s_i$. We use $C_0.E$ to denote the resulting configuration; we also say that $C = C_0.E$ is an *extension of* C_0 . By convention, C is a trivial extension of C .

An *execution* of the protocol is a sequence of steps E applicable to C_{init} . A finite execution E results in a *reachable* configuration C_E . Somewhat redundantly, we call a sequence of steps applicable to a reachable configuration C_E an *extension of* E . It immediately follows that steps of disjoint sets of processes commute:

► **Lemma 4.** *Let C be a reachable configuration, and E and E' be sequences of steps of disjoint sets of processes. If both E and E' are applicable to C , then $C.E.E'$ and $C.E'.E$ are identical reachable configurations.*

In an infinite execution, a process is *correct* if it executes infinitely many steps. We assume that every message m that is sent to a correct process p is eventually received, i.e., the execution will eventually contain a step $(p, m, -)$. As we assume that at most one process is allowed to fail by crashing, we only consider infinite executions in which at least $n - 1$ out of n processes are correct.

Without loss of generality, we assume that the implemented coin is binary: the correct processes either all output 0 or all output 1.

A configuration C is called *bivalent* if it has an extension $C.E_0$ in which some process outputs 0 and an extension $C.E_1$ in which some process outputs 1. Notice that any configuration preceding a bivalent configuration must be bivalent. Also, no process can produce a random-coin output in a bivalent configuration: otherwise, we get an execution in which two processes disagree on the output.

A protocol configuration that is not bivalent is called *univalent*: 0-valent if it has no extension in which 1 is decided or 1-valent otherwise.

► **Lemma 5.** *The initial configuration C_{init} is bivalent.*

Proof. The algorithm must output each of the two values with a positive probability. Thus, for each $v \in \{0, 1\}$, there exists an assignment of local coin outcomes and a message schedule that result in an execution with outcome v . ◀

► **Lemma 6.** *Let C be a reachable configuration, and E and E' be sequences of steps of a process p applicable to C . If $C.E$ and $C.E'$ are univalent, then they have the same valencies.*

Proof. The difference between $C.E$ and $C.E'$ consists in the local state of p and the message buffer. Since the algorithm is required to tolerate a single crash fault, there must exist a sequence of steps E'' that does not include any steps of p such that some process q outputs a value $v \in \{0, 1\}$ in $C.E''$. Moreover, as E'' contains no steps of p , E'' is also applicable to both $C.E$ and $C.E'$. But as the two configurations have opposite valencies, and q decides the same value in $C.E.E''$ and $C.E'.E''$. Thus, if $C.E$ and $C.E'$ are univalent, then they must have the same valencies. ◀

Following the steps of [17], we show that the protocol must have a *critical* configuration D for which:

- there exist steps s_p and s_q of processes p and q such that both s_p and $s_q.s_p$ are applicable to D
- $D.s_p$ is 0-valent;
- $D.s_q.s_p$ is 1-valent.

► **Lemma 7.** *There must exist a critical configuration.*

Proof. Starting with the initial (bivalent by Lemma 5) configuration $C := C_{init}$, we pick process $p := p_1$ and check if there exists C' , an extension of C , and s_p , a step of p applicable to C' in which the oldest message to p in the message buffer is consumed (if any), such that $C'.s_p$ is bivalent. If this is the case, we set C to $C'.s_p$, pick up the next process $p := p_2$. Again, if there exists C' , an extension of C , and s_p , a step of p in which the oldest message to p in the message buffer is consumed, such that $C'.s_p$ is bivalent, then we set C to $C'.s_p$. We repeat the procedure, each time picking the next process (in the round-robin order, i.e., after p_n we go to p_1 , etc.), as long as it is possible.

The first observation is that the procedure cannot be repeated indefinitely. Indeed, otherwise, we obtain an infinite sequence of steps in which every process takes infinitely many steps and receives every message sent to it (and which is, thus, an execution of the protocol) that goes through bivalent configurations only. Hence, this execution cannot produce an output – a contradiction with the Termination property.

Thus, there must exist a bivalent configuration C and a process p , such that for each C' , an extension of C and each $s_p = (p, m, -)$, a step of p applicable to C' , $C'.s_p$ is univalent, where m is the oldest message addressed to p in C .

Let $s_p = (p, m, r)$ be a step of p applicable to C . Without loss of generality, let $C.s_p$ be 0-valent.

As C is bivalent, it must have an extension $E = e_1, \dots, e_k$ such that $C.E$ is 1-valent.

Let ℓ be the largest index in $\{1, \dots, k\}$ such that either s_p is not applicable to $C_\ell = C.e_1, \dots, e_\ell$ or $C_\ell.s_p$ is 1-valent. Such an index exists, as $C.s_p$ is 0-valent $C.E$ is 1-valent and for any C' , an extension of C , if s_p is applicable to C' , then $C'.s_p$ is bivalent.

Suppose first that s_p is not applicable to C_ℓ . As s_p is applicable to every configuration $C_j = C.e_1, \dots, e_j$, $j = 1, \dots, \ell - 1$, e_ℓ must be of the form (p, m, r') , i.e., e_ℓ must consume the message received in $s_p = (p, m, r)$. By our assumption $C_\ell = C_{\ell-1}.(p, m, r')$ is univalent. As C_ℓ has a 1-valent descendant C' , it must be 1-valent too.

But, by Lemma 6, $C_{\ell-1}.s_p = C_{\ell-1}.(p, m, r)$ and $C_\ell = C_{\ell-1}.(p, m, r')$, must have the same valencies – a contradiction.

Thus, $C_\ell.s_p$ is 1-valent. Hence, we have a bivalent configuration $D = C_{\ell-1}$ and steps s_p and $s_q = e_\ell$ such that both s_p and $s_q.s_p$ are applicable to D . Moreover, $D.s_p$ is 0-valent and $D.s_q.s_p$ is 1-valent. Thus, we get a critical configuration. ◀

Finally, we establish a contradiction by showing that:

► **Lemma 8.** *No critical configuration may exist.*

Proof. By contradiction, let a critical configuration D exist, and let s_p and s_q be steps of p and q applicable to D such that $D.s_p$ is 0-valent and $D.s_q.s_p$ be 1-valent.

If s_q is a step of p , then Lemma 6 establishes a contradiction.

Otherwise, consider any infinite execution going through to $D.s_p$ in which all processes but q take infinitely many steps in this execution after $D.s_p$. By the Termination property, there must exist a finite sequence of steps E such that some process outputs a value $v \in \{0, 1\}$ in $D.s_p.E$. As $D.s_p$ is 0-valent, $v = 0$. Moreover, as E contains no steps of q and p has the same state in $D.s_p$ and $D.s_q.s_p$, E is also applicable to $D.s_q.s_p$. Thus, 0 is also decided in $D.s_q.s_p.E$ – a contradiction with the assumption that $D.s_q.s_p$ is 1-valent. ◀

Lemmata 7 and 8 imply:

► **Theorem 9.** *There does not exist a 1-resilient asynchronous random coin protocol.*

B Intersecting Random Subsets

A direct application of an approximate common coin is a problem that we call *intersecting random subsets*. In this problem the following information is globally known: a set S , where $|S| = s$; and two parameters m and k , where $1 \leq k \leq m \leq s$. Each correct process i chooses S_i , a random subset of S with cardinality m , such that $|\bigcap_{j \text{ - correct}} S_j| \geq m - k$. In other words, there are at least $m - k$ elements chosen at random that appear in all subsets.

The following variation of Gray Codes [23, 36] will be instrumental:

► **Definition 10.** *Code $C_{s,m}$ is a list of $\binom{s}{m} - 1$ binary strings (called code words) satisfying the following conditions:*

- *for all i , $C_{s,m}[i]$ is a string of m ones and $s - m$ zeros;*
- *every string of m ones and $s - m$ zeros is present in $C_{s,m}$ exactly once;*
- *$\forall i \in \{1, \dots, \binom{s}{m} - 1\} : C_{s,m}[i]$ and $C_{s,m}[i-1]$ differ in two bits;*
- *$C_{s,m}[\binom{s}{m} - 1]$ and $C_{s,m}[0]$ differ in two bits.*

The following recursive equation satisfies all requirements.

$$\forall s \geq 0 \text{ and } m \in [0..s] : C_{s,m} = \begin{cases} [\text{“ ”}] & \text{if } 0 = m = s \\ [\text{“000...000”}] & \text{if } 0 = m < s \\ [\text{“111...111”}] & \text{if } 0 < m = s \\ 0 \| C_{s-1,m}, \text{reverse}(1 \| C_{s,m-1}) & \text{otherwise} \end{cases}$$

In order to avoid computing all $\binom{s}{m}$ code words, when s and m are large, we can use the following recursive formula to efficiently (with $O(\text{poly}(s))$ operations) find a code word by its index:

$$\forall s \geq 0 \text{ and } m \in [0..s], i \in [0.. \binom{s}{m} - 1] : \\ C_{s,m}[i] = \begin{cases} \text{“ ”} & \text{if } 0 = m = s \\ \text{“000...000”} & \text{if } 0 = m < s \\ \text{“111...111”} & \text{if } 0 < m = s \\ 0 \| C_{s-1,m}[i] & \text{if } 0 < m < s, i < \binom{s-1}{m} \\ 1 \| C_{s-1,m-1}[\binom{s-1}{m-1} - (i - \binom{s-1}{m}) - 1] & \text{if } 0 < m < s, i \geq \binom{s-1}{m} \end{cases}$$

For example, here is $C_{5,2}$: [00011, 00110, 00101, 01100, 01010, 01001, 11000, 10100, 10010, 10001].

Intuitively, this code is composed of binary strings of length s and it can be read in the following manner: if the i -th position of the string is 1, then i -th element is selected, otherwise it is considered to be left outside. Moreover, this code has the property that consecutive numbers differ only by swapping exactly one position set to 1 with a position marked with a 0. Therefore all consecutive subsets have the same fixed size and include $m - 1$ common elements and differ by only one. Hence, by generating an approximate common random coin over the domain $\{0.. \binom{s}{m} - 1\}$ with parameter $\epsilon \leq k \cdot \binom{s}{m}^{-1}$, processes can select subsets of size m differing by at most k elements.

This could be interesting for selecting a committee among n users in scenarios subject to a mobile Byzantine adversary, i.e. on systems where the set of processes who display malicious behaviour changes, provided the time to corrupt a majority of processes in any given committee is higher than an asynchronous round. Note that this solution provides an interesting alternative to committee elections in protocols such as Algorand [22]. It not only is completely asynchronous, but it also guarantees a fixed committee size and provides a way to control the intersection of quorums obtained by different users. Recall that in the case of Algorand, with non-zero probability, it might happen that quorums do not intersect at all.