

Robust Random Number Generation for Peer-to-Peer Systems

Baruch Awerbuch*
Dept. of Computer Science
Johns Hopkins University
Baltimore, MD 21218, USA

Christian Scheideler†
Institut für Informatik
Technische Universität München
85748 Garching, Germany

May 1, 2007

Abstract

We consider the problem of designing an efficient and robust distributed random number generator for peer-to-peer systems that is easy to implement and works even if all communication channels are public. A robust random number generator is crucial for avoiding adversarial join-leave attacks on peer-to-peer overlay networks. We show that our new generator together with a light-weight rule recently proposed in [4] for keeping peers well-distributed can keep various structured overlay networks in a robust state even under a constant fraction of adversarial peers.

1 Introduction

Due to their many applications, peer-to-peer systems have recently received a lot of attention both inside and outside of the research community. Most of the structured peer-to-peer systems are based on two influential papers: a paper by Plaxton et al. on locality-preserving data management in distributed environments [20] and a paper by Karger et al. on consistent hashing and web caching [13]. The consistent hashing approach is a very simple and elegant approach that assigns to each peer a (pseudo-)random point in the $[0, 1)$ -interval. Based on this approach, various local-control rules have been proposed to decide how to interconnect the peers so that they form a well-connected network with good routing properties that is easy to maintain (see, e.g., [18] for a general framework).

In open peer-to-peer systems, the presence of adversarial peers cannot be avoided. Hence, not only scalability but also robustness against adversarial behavior is an important issue. The key to scalability and robustness for peer-to-peer networks based on the consistent hashing approach is to keep the honest and adversarial peers well-distributed in the $[0, 1)$ -interval. However, just assigning a random or pseudo-random point to each new peer (by using some random number generator or cryptographic hash function) does not suffice to keep the honest and adversarial peers well-spread [2]. People in the peer-to-peer community are well aware of this problem [8, 9] and various solutions have been proposed that may help alleviating it in practice [6, 7, 19, 27, 28, 30] but until recently no mechanism was known that can *provably* keep the peers in a well-distributed state without sacrificing the openness of the system.

Various light-weight perturbation rules that can keep the honest and adversarial peers well-distributed have recently been proposed in [4, 10, 26]. These rules do not need to be able to distinguish between the honest and adversarial peers, but a crucial prerequisite for them to work is a robust distributed random number

*Email: baruch@cs.jhu.edu. Partially supported by NSF grants CCF 0515080, ANIR-0240551, CCR-0311795, and CNS-0617883.

†Contact author. Email: scheideler@in.tum.de

generator. This random number generator has to work correctly in a system without mutual trust relationships and must be robust against arbitrary adversarial behavior to be applicable to peer-to-peer systems. Certainly, designing such a random number generator is not an easy task.

1.1 Robust distributed random number generation

How can we generate random numbers in a peer-to-peer system with adversarial presence? The most naive approach is to let every peer generate its own random numbers. This approach is problematic since in a dynamic peer-to-peer system it is impossible to collect sufficient statistical evidence to accuse a particular peer of generating non-random numbers. Yet, somewhat surprisingly, it is still possible to use this approach to maintain a robust peer-to-peer network, but at the cost of losing scalability [3]. So a different approach is needed.

A more reasonable approach is the following. Suppose that we need a random number generator that generates a number by selecting a binary string uniformly at random out of $\{0, 1\}^s$ for some s . Consider the situation that a group P of the peers wants to generate a random number. Each (honest) peer p in P may then select a random number $x_p \in \{0, 1\}^s$ and commit to it to all other peers in P using a bit commitment scheme (a particularly secure one-way hash function h for which $h(x)$ does not reveal anything about x) [12, 17]. Once all commitments have been made, the peers will reveal their random numbers, and if they all do, every peer computes $x = \bigoplus_{p \in P} x_p$, where \oplus is the bit-wise XOR operation. The XOR operation has the nice property that as long as at least one x_p is chosen uniformly at random and the other numbers are independent of it, x is distributed uniformly at random in $\{0, 1\}^s$. Hence, *if* the scheme succeeds and at least one honest peer participates in it, a random number x will be generated. But the adversarial peers can easily let the scheme fail, and this not only in an oblivious manner but also in an adaptive manner (by just waiting for enough numbers x_p to be revealed before revealing their own numbers). Thus, in order to avoid a significant bias on the successfully generated random numbers, the fraction of adversarial peers in the system would have to be so small that no adversarial peer will be present in most of the groups P that are used for the random number generation. Such an approach was pursued in [2].

To avoid the problems above, we recently proposed a distributed random number generator that is based on verifiable secret sharing [4]. This random number generator can still fail if the peer initiating it does not behave correctly, but it has the advantage that if the peer initiating it is honest, then the random number generation is guaranteed to succeed, and whenever the random number generation succeeds, the number generated will be random.

Yet, using this scheme is not completely satisfying. First of all, an adversary can let it fail in an adaptive manner (i.e., it can let it fail after knowing the final key), which is sufficient to create a significant bias, even though the adversary cannot undermine the randomness of the generated key. It just has to run sufficiently many attempts until a key is generated that falls into a desired range. Furthermore, the scheme is not easy to implement and private channels are needed between the peers. So the question that led to this paper was:

Is it possible to design an elementary and sufficiently unbiased distributed random number generator that even works for public channels and a constant fraction of adversarial peers?

Remarkably, this paper shows that this is possible.

1.2 Related work on random number generation

Surprisingly little has been published about robust random number generators for distributed systems. Random number generators have mostly been studied in the context of pseudo-random number generators (PRNGs) with small seed or cryptographically secure random number generators (CSRNG). The main difference between a PRNG and a CSRNG is that a CSRNG should be indistinguishable from random on any

examination, whereas a PRNG is normally only required to look random to standard statistical tests. For foundations and surveys on random number generators see, e.g., [11, 16, 23, 32].

There are many protocols for distributed systems with adversarial presence that need random numbers for atomic broadcasting, leader election and almost-everywhere agreement (e.g., [14, 21] for recent results), but in these it is sufficient that every peer chooses its own random numbers.

Unbiased random numbers can be computed via verifiable secret sharing or secure multiparty computation schemes (e.g., [5, 29]), but these are not easy to implement (since they need error correction techniques), and they require private channels.

1.3 Details of our random number generator

The basic idea behind our random number generator is the insight that generating a *single* random number is difficult with public channels but generating a *batch* of random numbers is doable. An *m-random number generator* (or *m-RNG*) is a random number generator that generates a batch of up to m random numbers. We assume that every random number is represented as a binary string in $\{0, 1\}^s$ for some fixed s . Given an *m-RNG* G and any subset $S \subseteq \{0, 1\}^s$, let $E_G(S)$ be the expected number of keys y generated by G with $y \in S$. Ideally, G should satisfy $E_G(S) = m \cdot |S|/2^s$ for all $S \subseteq \{0, 1\}^s$. Let $E(S) = m \cdot |S|/2^s$. Then we define the *bias* $\beta(G)$ of G as

$$\beta(G) = \max_{S \subseteq \{0,1\}^s} \max \left\{ \frac{E_G(S)}{E(S)}, \frac{E(S)}{E_G(S)} \right\}$$

The *m-RNG* that we present in this paper is called *round-robin random number generator* (or short round-robin RNG). Let P be the group of m peers this protocol is applied to. The basic ideas of the protocol can be summarized as follows:

- When correctly initiated, every peer in P will supervise the generation of one random number in $\{0, 1\}^s$. A peer whose random number generation fails can send an accusation to the peers in P in which it can accuse exactly one other peer. Honest peers will run the random number generation one after the other (using a proper timing scheme) so as to maximize the effect of the accusations and thereby minimize the number of times an adversarial peer can cause the failure of a random number generation supervised by an honest peer.
- A single random number is generated by the supervising peer taking over the role of a dealer and the others being a group of players. Both the players and the dealer commit to a key. However, as we will see, the dealer key is a special master key that is committed to first and revealed last. In this way, the dealer is the only one that can *adaptively* decide whether to let the random number generation fail or not. However, this is the only way in which the dealer can bias the random number generation. It cannot make its probability distribution non-uniform if at least one honest player is participating in it.

More details are given in Section 2. For this protocol, the following theorem is shown.

Theorem 1.1 *Suppose that $|P| = m$ and there are $t < m/6$ adversarial peers in P . Then the round-robin RNG generates random keys $y_1, y_2, \dots, y_k \in \{0, 1\}^s$ with $m - 2t \leq k \leq m$ and the property that for all subsets $S \subseteq \{0, 1\}^s$ with $\sigma = |S|/2^s$,*

$$E[\{i \mid y_i \in S\}] \in [(m - 2t)\sigma, m \cdot \sigma]$$

The worst-case message complexity of the protocol is $O(m^2)$.

Hence, the bias of our *m-RNG* is just $1 + \frac{2t}{m-2t}$, which is a constant. It turns out that this bias is small enough in order to maintain a scalable and robust peer-to-peer network.

1.4 Application to robust peer-to-peer networks

In the area of peer-to-peer systems, work on robustness in the context of overlay network maintenance has mostly focused on how to handle a large fraction of faulty peers (e.g., [1, 24, 31]) or churn, that is, peers frequently enter and leave the system (e.g., [15, 22]). However, none of these approaches can protect a peer-to-peer network against adaptive join-leave attacks. In an adaptive join-leave attack, adversarial peers repeatedly join and leave a network in order to occupy certain areas of the network. To prevent them from doing this, proper join and leave protocols have to be found so that the honest and adversarial peers are kept well-spread in the $[0, 1)$ -interval. More precisely, what we would like to aim for is that at any time point with n peers in the system the following two conditions can be met for every interval $I \subseteq [0, 1)$ of size at least $(c \log n)/n$ for a constant $c > 0$:

- *Balancing condition*: I contains $\Theta(|I| \cdot n)$ peers.
- *Majority condition*: the honest peers in I are in the majority.

If this is the case, then proper region-based overlay networks and routing rules can be defined to guarantee connectivity and correct routing (e.g., [4]). However, maintaining the two conditions under adaptive adversarial join-leave attacks turns out to be quite tricky. Just assigning a random or pseudo-random point to each new peer (by using some random number generator or cryptographic hash function) does not suffice to preserve the balancing and majority conditions [2]. Fortunately, just recently we found a join operation, called cuckoo rule, that can solve this problem [4].

1.5 The cuckoo rule

In the following, a *region* is an interval of size $1/2^r$ in $[0, 1)$ for some integer r that starts at an integer multiple of $1/2^r$. Hence, there are exactly 2^r regions of size $1/2^r$. A *k-region* is a region of size (closest from above to) k/n , and for any point $x \in [0, 1)$, the *k-region* $R_k(x)$ is the unique *k-region* containing x .

Cuckoo rule: If a new node v wants to join the system, pick a random $x \in [0, 1)$. Place v into x and move all nodes in $R_k(x)$ to points in $[0, 1)$ chosen uniformly and independently at random (without replacing any further nodes).

Suppose that we have n honest peers and ϵn adversarial peers in the system for some $\epsilon < 1$. For the situation that the adversary adaptively rejoins the system with its peers in a one-by-one fashion, it was shown [4] that as long as $\epsilon < 1 - 1/k$, the *k-cuckoo rule* satisfies the balancing and majority conditions for a polynomial number of rejoin operations, with high probability. However, for the cuckoo rule to be implementable in a distributed system, a robust distributed random number generator is needed. Furthermore, the cuckoo rule may need up to $O(\log^2 n)$ random bits in the worst case (for $O(\log n)$ peers that need to be replaced).

1.6 The round-robin cuckoo rule.

The problem with $O(\log^2 n)$ bits is solved by proposing a slight adaptation of the cuckoo rule that we call the *de Bruijn cuckoo rule*. The new rule has the benefit that only $O(\log n)$ random bits are needed in the worst case (for two random points in $[0, 1)$).

In order to solve the problem with the random number generator, we combine the round-robin RNG with the de Bruijn cuckoo rule to the so-called *round-robin cuckoo rule*. It works in a way that for every successful random number generation in the round-robin RNG, the de Bruijn cuckoo rule is used. The protocol has the following performance.

Consider adversarial join-leave attacks in a system with n honest peers and ϵn adversarial peers. Let β be the bias of the round-robin RNG. Then it holds:

Theorem 1.2 For any constants ϵ , k and $\beta \geq 1$ with $\epsilon < (1/\beta)(1/\beta - 1/k)$, the round-robin cuckoo rule with bias β satisfies the balancing and majority conditions for a polynomial number of rounds, with high probability, for any adversarial strategy within our model.

Hence, Theorem 1.2 is a natural extension of the result in [4], which assumes a bias of $\beta = 1$.

1.7 Structure of the paper

In Section 2, we present the round-robin random number generator, and in Section 3 we show how to use it to counter join-leave attacks in peer-to-peer networks. The paper ends with conclusions.

2 Robust random number generation

In this section we consider the situation that we have a set P of m players denoted p_1, \dots, p_m . We distinguish between honest and adversarial players. The honest players follow the protocol in a correct and timely manner, whereas the adversarial players may behave in an arbitrary way, including arbitrary collusion among the adversarial players. Our goal is to find *elementary* protocols that construct random numbers with a uniform distribution in $\{0, 1\}^s$ for some given s , even under adversarial presence.

First, we state some assumptions, and then we present the round-robin random number generator. After its analysis, we discuss some extensions for peer-to-peer systems.

2.1 Assumptions

We assume that only point-to-point communication is available and that all information sent out by a player can be seen by the adversary. Thus, no broadcasting primitive and no private channels are given, which is often the case in other robust distributed protocols like verifiable secret sharing. We just need a mechanism that allows the players to verify the sender of a message. For this, we assume the existence of a proper signature scheme. A message m signed by player p will be denoted as $(m)_p$.

Honest players are supposed to act not only in a correct but also a timely manner (which is important to maintain dynamic systems such as peer-to-peer networks). We assume that any message sent from one honest player to another honest player needs at most δ time steps to be received and processed by the recipient for some fixed δ , and we assume that the clock speeds of the honest players are roughly the same. However, the clocks do not have to be synchronized (i.e., show the same time) nor do we require the protocols to run in a synchronous mode (i.e., all players must send their messages at exactly the same time). The latter assumption makes it hard to generate unbiased random keys even though there is a notion of time because the adversarial players can always choose to be the last to send out messages, thereby maximizing the control they have on the generation of the random number.

For the random number generation, we need a bit commitment scheme h , i.e., a scheme where $h(x)$ does not reveal anything about x . In practice, a cryptographic hash function might be sufficient for h so that the protocols below can be easily implemented. Furthermore, we assume that all honest players have a perfect random number generator. In practice, pseudo-random number generators that pass a certain collection of statistical tests (such as the DIEHARD tests) might be sufficient here.

2.2 Round-robin random number generator

Suppose that we have a set P of m players, p_1, \dots, p_m , that know each other and their indexing, with any t of them being adversarial for some $t < m/6$. The round-robin random number generator works as follows for some player $p^* \in P$ initiating it.

1. p^* sends a signed request to initiate the random number generation to all players in P .
2. Once player $p_i \in P$ receives p^* 's signed initiation request for the first time (from anywhere), it forwards it to all other players in P . Afterwards, it sets $P_i := P \setminus \{p_i\}$ and waits for $i \cdot 8\delta$ time steps. Each time it receives an accusation $(p_k)_{p_j}$ from a player $p_j \in P$ it has not received an accusation from yet, it sets $P_i := P_i \setminus \{p_k\}$. Once the $i \cdot 8\delta$ steps are over, p_i initiates step (3). p_i terminates after $(m + 1)8\delta$ steps.
3. If $|P_i| \geq 2m/3$, then p_i chooses a random $x_i \in \{0, 1\}^s$ and sends $(h(x_i), P_i)_{p_i}$ to all players in P_i . Otherwise, p_i aborts the protocol (which will not happen if $t < m/6$).
4. Each player $p_j \in P_i$ receiving a message $(h(x_i), P_i)_{p_i}$ for the first time from p_i with $|P_i| \geq 2m/3$ chooses a random $x_j \in \{0, 1\}^s$ and sends the message $(p_i, h(x_j), P_i)_{p_j}$ to p_i . Otherwise, it does nothing.
5. If all players in P_i reply within 2δ time steps, then p_i sends $(\{(p_i, h(x_j), P_i)_{p_j} \mid p_j \in P_i\})_{p_i}$ to all players in P_i . Otherwise, p_i sends an accusation $(p_j)_{p_i}$ for any $p_j \in P_i$ that did not reply correctly or in time to all players in P and stops its attempt of generating a random number.
6. Once $p_j \in P_i$ receives $(\{(p_i, h(x_k), P_i)_{p_k} \mid p_k \in P_i\})_{p_i}$ from p_i , p_j sends $(x_j)_{p_j}$ to p_i .
7. If p_i gets a correct reply back from all players in P_i within 2δ time steps, then it sends $(x_i, \{(x_j)_{p_j} \mid p_j \in P_i\})_{p_i}$ to all players in P_i and computes $y_i = x_i \oplus \bigoplus_{p_j \in P_i} x_j$ where \oplus is the bit-wise XOR operation. Otherwise, p_i sends an accusation $(p_j)_{p_i}$ to all players in P for any $p_j \in P_i$ that did not reply correctly or in time and stops.
8. Once $p_j \in P_i$ receives $(x_i, \{(x_k)_{p_k} \mid p_k \in P_i\})_{p_i}$, p_j verifies that all keys are correct. Then p_j computes $y_j^{(i)} = x_i \oplus \bigoplus_{p_k \in P_i} x_k$ and sends the message $(y_j^{(i)})_{p_j}$ to p_i .
9. If p_i receives y_j from at least $2m/3$ players in P within 2δ time steps, it accepts the computation and otherwise sends an accusation $(p_j)_{p_i}$ to all players in P for any $p_j \in P_i$ that did not reply correctly or in time.

We define the random number generation of p_i to be *successful* if p_i receives the same key from at least $2m/3$ many players in step (9). This is important for p_i since it will need the support of at least $2m/3$ other players for further operations that we will discuss in the next section.

2.3 Analysis of the round-robin RNG

The round-robin RNG has the following performance.

Theorem 2.1 *Suppose that $|P| = m$ and there are $t < m/6$ adversarial players in P . Then the round-robin RNG generates random keys $y_1, y_2, \dots, y_k \in \{0, 1\}^s$ with $m - 2t \leq k \leq m$ and the property that for all subsets $S \subseteq \{0, 1\}^s$ with $\sigma = |S|/2^s$,*

$$\mathbb{E}[|\{i \mid y_i \in S\}|] \in [(m - 2t)\sigma, m \cdot \sigma].$$

The worst-case message complexity of the protocol is $O(m^2)$.

In order to prove the theorem, we start with some simple claims.

Some basic facts.

Because of the flooding strategy in step (2) and the definition of δ it holds:

Claim 2.2 *No matter whether p^* is adversarial or not, all honest players start the protocol within δ steps.*

Since each honest player p_i needs at most 7δ time steps to complete the protocol from step (3) to (9) and starts after waiting for $i \cdot 8\delta$ steps, the claim above implies the following claim.

Claim 2.3 *No two honest players execute their random number generation scheme (steps (3) to (9)) at the same time.*

Hence, honest player p_i can make use of the accusations of all honest players p_j with $j < i$ in order to keep its own problems with the random number generation as small as possible.

Next, we bound the size of any P_i for an honest player p_i . Recall that honest players are supposed to work in a correct and timely manner. Hence, honest players will never accuse other honest players of any wrongdoing but only adversarial players. Since every adversarial player can issue at most one accusation to any honest player, there will be at least $m - 2t$ honest players left in every set P_i of an honest player p_i throughout the protocol. Hence, we get:

Lemma 2.4 *If $t < m/6$ then $|P_i| \geq 2m/3$ throughout the protocol for every honest player p_i .*

Moreover, every player p_i can only be successful for one key. This is because all players in P_i have to see commitments to the same P_i for all players in P_i and $|P_i| \geq 2m/3$ before revealing their random keys in step (6). Since $t < m/6$, this means that there must be more than $m/2$ honest players in P_i , which can only be possible for at most one P_i . Hence, we get.

Lemma 2.5 *If $t < m/6$ then every player can be successful for at most one key.*

Analysis of steps (3) to (9).

Next, we focus on the execution of steps (3) to (9) by some fixed peer p_i . First, we consider the case that p_i is honest, and then we consider the case that p_i is adversarial.

Lemma 2.6 *If p_i is honest and $|P_i| \geq 2m/3$, then no matter how many adversarial players there are in P_i , if the protocol terminates successfully, then the key y_i generated by p_i is distributed uniformly at random in $\{0, 1\}^s$ and all honest players in P_i compute the same key as p_i .*

Proof. p_i will not reveal x_i before the keys in P_i have all been revealed. Hence, the probability distribution on $z = \bigoplus_{p_j \in P_i} x_j$ must be independent of x_i . But for any probability distribution on $z = \bigoplus_{p_j \in P_i} x_j$ that is independent of x_i it holds that if x_i is chosen uniformly at random in $\{0, 1\}^s$, then also $y_i = x_i \oplus z$ is distributed uniformly at random in $\{0, 1\}^s$. Moreover, also the decision of the adversarial players to let the random number generation fail must be independent of x_i and can only be a function of z because x_i will not be revealed before. Hence, it holds for any adversarial strategy and any $y^* \in \{0, 1\}^s$ that

$$\Pr[y_i = y^* \mid \text{generation of } y_i \text{ successful}] = \Pr[y_i = y^*] = \frac{1}{2^s}$$

If p_i succeeds with computing y_i , then it informed all players in P_i about the revealed keys, and all honest among them will accept these keys since they match the message sent out by p_i in step (5). Hence, all honest players in P_i compute the same key as p_i . \square

Notice that if the adversarial players knew about x_i *before* deciding to let the random number generation fail, they can create a significant bias, even if the other keys were chosen independent of x_i . A simple example for this would be:

Focus on any fixed $y^* \in \{0, 1\}^s$. If $y_i = y^*$, then let the attempt fail, and otherwise let it be successful.

It is easy to see that this would make it very unlikely for the round-robin RNG to generate y^* (since it would have to be generated more than t times to be successful at least one). Hence, it is crucial that x_i is only revealed *after* all the other keys have been revealed. Next, we consider the case that p_i is adversarial.

Lemma 2.7 *If p_i is adversarial, then no matter what p_i and the other adversarial players in P_i do, whenever an honest player p_j reveals its key x_j , $y_i^{(j)}$ has a uniform distribution on $\{0, 1\}^s$.*

Proof. An honest player p_j will only reveal x_j once it receives $(\{(p_i, h(x_k), P_i)_{p_k} \mid p_k \in P_i\})_{p_i}$ from p_i and $p_j \in P_i$ (so that $y_j^{(i)}$ is well-defined). In this case, x_j is a random number that is independent of $z = x_i \oplus \bigoplus_{p_k \in P_i \setminus \{p_j\}} x_k$, and since x_j is independent of z and chosen uniformly at random, $y_j^{(i)} = x_j \oplus z$ has a uniform distribution. \square

Notice, however, that p_i can commit to different sets P_i to different honest players without being detected, so the keys $y_j^{(i)}$ can differ among the honest players. Nevertheless, if p_i wants to be successful (i.e., collect commitments to the same key from at least $2m/3$ many players), it must let more than $m/2$ honest players p_j succeed with computing the same $y_j^{(i)}$, which has a uniform distribution.

Still, the adversarial players can create a bias on the *successfully* computed keys since after knowing y_i , an adversarial player p_i still has the option to let the key generation be successful or not. Fortunately, this bias cannot be too large, as shown in the following lemma.

Analysis of the entire protocol.

Lemma 2.8 *If $t < m/6$ then at least $m - 2t$ of the $m - t$ random number generations initiated by the honest players are successful, irrespective of whether p^* is adversarial or not. Furthermore, it holds for all subsets $S \subseteq \{0, 1\}^s$ with $\sigma = |S|/2^s$ that $\mathbb{E}[|\{i \mid y_i \in S \text{ for a successful } y_i\}|] \in [(m - 2t)\sigma, m \cdot \sigma]$*

Proof. According to Lemma 2.7, every key y that an honest player p commits to must be distributed uniformly at random in $\{0, 1\}^s$. However, whereas the adversarial players can adaptively abort the random number generation initiated by adversarial players, it follows from the protocol that they can only do this in an oblivious way for the honest players. We know from Claim 2.3 that the adversarial players can only sabotage the random number generation of at most t honest players. This insight together with Lemma 2.4 implies that at least $m - 2t$ random number generations of honest players p_i will be successful, and their success does not depend on their values. Thus, the probability for any of these players p_i that $y_i \in S$ is equal to σ and, therefore, the expected number of successful p_i 's with $y_i \in S$ is at least $(m - 2t)\sigma$.

On the other hand, we know from Lemma 2.5 that at most m key generations can be successful, and since every successfully generated key y_i is distributed uniformly at random in $\{0, 1\}^s$, the probability for any y_i to be in S is equal to σ . Hence, the expected number of successful p_i 's with $y_i \in S$ is at most $m \cdot \sigma$. \square

The next lemma follows immediately from the protocol.

Lemma 2.9 *The message complexity of the round robin-random RNG is $O(m^2)$.*

2.4 Extensions

In our random number generator we assumed that the players in P know each other and the indexing. This assumption can be problematic in peer-to-peer systems since there might be a disagreement among the honest players about the set of adversarial players in P . Fortunately, it is not too difficult to address this issue, as we will show in the following.

First, suppose that for every player p_i there is an agreement among the honest players about its index $i \in \{1, \dots, m\}$. Furthermore, every honest player knows every other honest player in P but not necessarily all the adversarial players in P , and therefore the honest players may have different estimates of $|P|$. Then the round robin RNG can easily be modified so that Theorem 2.1 still holds:

- In lines 3 and 9, replace “ $|P_i| \geq 2m/3$ ” by “ $|P_i| \geq 2m_i/3$ ” where m_i is the number of players that p_i knows in P .
- In line 4, replace “ $|P_i| \geq 2m/3$ ” by “ $|P_i \cap P_j| \geq 2m_j/3$ ”.
- An honest player only accepts random number initiations and accusations from players that it knows about.

Given these rules, the following lemma can be used to replace Lemma 2.4:

Lemma 2.10 *If $t < m/6$ then $|P_i \cap P_j| \geq 2m_j/3$ throughout the protocol for every pair of honest players p_i and p_j .*

Proof. Consider some fixed honest player p_i . Let $m_h \geq 5m/6$ be the number of honest players and $m_a < m/6$ be the number of adversarial players that p_i knows about. Then, initially, $|P_i| = m_h + m_a$, and P_i can be reduced by at most $2m_a$ accusations. Hence, throughout the protocol, $|P_i| \geq m_h - m_a \geq 2m/3 \geq 2m_i/3$. Since the number of honest players in P_i will be at least $m_h - m_a$ as well and honest players know each other, it follows that also $|P_i \cap P_j| \geq 2m_j/3$ for any pair of honest players p_i and p_j . \square

Furthermore, Lemma 2.5 can be replaced by the following lemma.

Lemma 2.11 *If $t < m/6$ then every player can be successful for at most one key.*

Proof. A prerequisite for an adversarial player p_i to be successful is that for a sufficient number of honest players p_j , $|P_i \cap P_j| \geq 2m_j/3$. Let $P_j = m_h + m_a$, where m_h is the number of honest players and m_a is the number of adversarial players p_j is aware of. In the worst case, the adversarial player may select $P_i \subseteq P_j$ so that $|P_i \cap P_j| = 2m_j/3$ and P_i contains all m_a adversarial players that p_j knows about. Even then, there are still $m_h - m_a \geq 2m/3 > m/2$ honest players left in $P_i \cap P_j$, which means that more than half of the honest players are needed to successfully generate a key, which implies the lemma. \square

With these lemmas, the rest of the results in Section 2.3 still holds.

Another problem is how to fix the indexing issue. When there is disagreement about P , it may not be possible for the honest players to agree on a common indexing scheme. Instead, they can use the following simple trick. Each player p_i picks a random slot out of $c \cdot m_i$ many slots for generating a random number, where c is a sufficiently large constant. Then it holds for every honest player p_i that $\Pr[p_i \text{ avoids the slots taken by other honest players}] \geq (cm_i - m)/(cm_i) \geq 1 - 8/(7c)$. Hence, if $m = \Theta(\log n)$ is sufficiently large (for some parameter n), then the Chernoff bounds can be used to prove that the number of slots occupied by the honest players is at least $(1 - 1/(2c))m_h$, w.h.p. (with respect to n), where m_h is the number of honest players. Thus, the adversarial players would only manage to let up to $t + m_h/(2c)$ random number generations of honest players fail, w.h.p., instead of just t , which is still acceptable if c is sufficiently large.

3 Application to robust peer-to-peer networks

In this section we show how to use the round-robin random number generator above to satisfy the balancing and majority conditions for any adversarial join-leave strategy for a polynomial number of rejoin operations, with high probability. We start with a formal model. Then we present the de Bruijn cuckoo rule, and afterwards we combine it with the round-robin RNG to obtain the round-robin cuckoo rule.

3.1 Model

Recall that we want to associate all peers with points in $[0, 1)$. These points can be encoded as binary strings from $\{0, 1\}^s$ (in a sense that $b = (b_1, \dots, b_s)$ represents $x_b = \sum_{i \geq 1} b_i/2^i$) for a sufficiently large s (in SHA-1, which is used by the Chord system, for example, $s = 160$).

There are n blue (or honest) nodes and ϵn red (or adversarial) nodes for some fixed constant $\epsilon < 1$. There is a rejoin operation that, when applied to node v , lets v first leave the system and then join it again from scratch. The leaving is done by simply removing v from the system and the joining is done with the help of a join operation to be specified by the system. We assume that the sequence of rejoin requests is controlled by an adversary. The adversary can only issue rejoin requests for the red nodes, but it can do this in an arbitrary adaptive manner. That is, at any time it can inspect the entire system and select whatever red node it likes to rejoin the system. The goal is to find an *oblivious* join operation, i.e., an operation that does not distinguish between the blue and red nodes, so that for *any* adversarial strategy above the balancing and majority conditions can be kept for any polynomial number of rejoin requests.

3.2 The de Bruijn cuckoo rule

Recall the original cuckoo rule in Section 1.5. We present a slight but crucial modification to this rule, called the *de Bruijn cuckoo rule*, which only needs two random numbers in $\{0, 1\}^s$, irrespective of k . The prefix *de Bruijn* was chosen because the rule can be easily implemented in dynamic de Bruijn graphs, as will be demonstrated in Section 3.3.

de Bruijn cuckoo rule: If a new peer v wants to join the system, pick random $x, y \in [0, 1)$. Place v into x and replace all peers in $R_k(x)$ in the following way. If $|R_k(x)| = 0$, we are done, and if $|R_k(x)| = 1$, then the peer in $R_k(x)$ is moved to position y . Otherwise, let $b = \lceil \log |R_k(x)| \rceil$. Given that y is represented by a binary string $(y_1, \dots, y_s) \in \{0, 1\}^s$, peer $i \geq 0$ in $R_k(x)$ is moved to position $((y_{s-b+1}, \dots, y_s) \oplus (i)_2) \circ (y_1, \dots, y_{s-b})$ where $(i)_2$ represents the binary representation of i and \circ the concatenation.

For example, suppose that $y = 0100110$ and $|R_k(x)| = 3$. Then the new positions of the three peers are $(10 \oplus 00) \circ 01001 = 1001001$ for peer 0, $(10 \oplus 01) \circ 01001 = 1101001$ for peer 1, and $(10 \oplus 10) \circ 01001 = 0001001$ for peer 2. This rule of mapping peers to new points has the following property:

Lemma 3.1 *Every replaced peer is moved to a position that is distributed uniformly at random in $\{0, 1\}^s$.*

Proof. Consider peer i in $R_k(x)$ for any fixed i and suppose that y is distributed uniformly at random in $\{0, 1\}^s$. Then $(y_{s-b+1}, \dots, y_s) \oplus (i)_2$ is distributed uniformly at random in $\{0, 1\}^b$ and (y_1, \dots, y_{s-b}) is distributed uniformly at random in $\{0, 1\}^{s-b}$, resulting in the lemma. \square

Moreover, any two peers in a region $R_k(x)$ with p peers have a distance of at least $(1/2)^{\log p - 1} \geq 1/(2p)$ of each other. Hence, when looking at the analysis in [4], it turns out that all results still hold when using a perfect random number generator (though in Lemma 2.6 and Lemma 2.10 the independence property of the new node positions has to be replaced by negative correlation, but the negative correlation is so small that it is negligible – more precise arguments on this will be given in Section 3.4).

Theorem 3.2 For any constants ϵ and k with $\epsilon < 1 - 1/k$, the de Bruijn cuckoo rule with parameter k satisfies the balancing and majority conditions for a polynomial number of rounds, with high probability, for any adversarial strategy within our model. The inequality $\epsilon < 1 - 1/k$ is sharp as counterexamples can be constructed otherwise.

3.3 Applying the de Bruijn cuckoo rule in a dynamic de Bruijn graph

To illustrate the application of the de Bruijn cuckoo rule in a dynamic overlay network, we will introduce the continuous-discrete variant of the de Bruijn graph proposed by Naor and Wieder [18]. We start with the definition of the classical de Bruijn graph.

Definition 3.3 The d -dimensional de Bruijn graph $DB(d)$ is an undirected graph $G = (V, E)$ with node set $V = [2]^d$ and edge set E that contains all edges $\{v, w\}$ with the property that $v = (v_1, \dots, v_d)$ and $w = (x, v_1, \dots, v_{d-1}) : x \in \{0, 1\}$.

Let us view each binary label (v_1, \dots, v_d) as a point $x \in [0, 1)$ with $x = \sum_{i=0}^{d-1} v_i/2^i$ and let d go to infinity. Then we obtain a continuous form of the de Bruijn graph with node set $U = [0, 1)$ and edge set $E = \{\{x, y\} \in U \mid y = x/2 \text{ or } y = (1+x)/2\}$. In order to convert this back into a discrete graph, one can use the continuous-discrete approach of Naor and Wieder [18].

For $i \in \{0, 1\}$ let $f_i(x) = (i+x)/2$. Given any finite set of points $V \subset [0, 1)$, we define the *home region* U_v of point v as the interval $[v, \text{succ}(v))$ where $\text{succ}(v)$ is the closest successor of v in V on the $[0, 1)$ -ring, i.e. the minimum point $w \in V$ so that $w > v$. We say that a point $v \in V$ *owns* a point $x \in [0, 1)$ if $x \in U_v$.

The *de Bruijn graph* $DB(V)$ of a point set V is an undirected graph with node set V that contains an edge $\{v, w\}$ for every two points $v, w \in V$ with $f_0(U_v) \cap U_w \neq \emptyset$ or $f_1(U_v) \cap U_w \neq \emptyset$. This definition immediately implies the following fact:

Fact 3.4 For any set of points V and any two points $x, y \in [0, 1)$ with $y = f_i(x)$ for some i it holds for the owner v of x and the owner w of y that $\{v, w\}$ is an edge in $DB(V)$.

If V is well-spread over U (i.e., the balancing condition holds), then one can also show that every peer in V has at most a logarithmic degree in $DB(V)$ and the diameter of $DB(V)$ is logarithmic, that is, $DB(V)$ is a scalable network and therefore useful for dynamic peer-to-peer systems.

Now, recall the de Bruijn cuckoo rule. Given that all peers are reliable, then in order to forward a peer i in $R_k(x)$ to position $((y_{s-b+1}, \dots, y_s) \oplus (i)_2) \circ (y_1, \dots, y_{s-b})$, we first move it along the (owners of the) points (x_1, \dots, x_s) , (y_s, x_1, \dots, x_s) , \dots , $(y_1, y_2, \dots, y_s, x_1, \dots, x_s)$ and from there along b further moves to $((y_{s-b+1}, \dots, y_s) \oplus (i)_2) \circ (y_1, \dots, y_s) \circ (x_1, \dots, x_s)$ whose owner is in direct proximity (if not the same) of $((y_{s-b+1}, \dots, y_s) \oplus (i)_2) \circ (y_1, \dots, y_{s-b})$. If not all peers are reliable, then a region-based routing as sketched in [4] has to be used to ensure the correct movement of the peers to their new positions.

3.4 The round-robin cuckoo rule

Finally, we show how to combine the de Bruijn cuckoo rule and the round-robin random number generator into a simple and efficient join protocol called *round-robin cuckoo rule* that achieves a result similar to Theorem 3.2.

Recall the definition of a region in Section 1.5. Given a node $v \in [0, 1)$, we define its *quorum region* R_v as the unique region of size closest from above to $(\gamma \log n)/n$, for a fixed constant $\gamma > 1$, that contains v .

We demand that whenever a new node u wants to join the system, it has to do so via a node v already in the system. v then initiates the following protocol:

1. v initiates the round-robin RNG in R_v (i.e., v acts as p^*).

2. For each successful node $v_i \in R_v$, v_i initiates the de Bruijn cuckoo rule by sending a message $(y_i, \{(y_j^{(i)})_{p_j} \mid v_j \in P_i\})_{p_i}$ with $1 + 2m/3$ signed keys to all nodes in R_v .
3. Once node $w \in R_v$ receives a correctly signed $(y_i, \{(y_j^{(i)})_{p_j} \mid v_j \in P_i\})_{p_i}$ containing more than $2m/3$ keys, it forwards it to all other nodes in R_v and initiates the de Bruijn cuckoo rule.

In the de Bruijn cuckoo rule, majority decision is used to execute the proper actions (see [4] for more details). Since step (3) ensures the “all or nothing” principle concerning the honest peers, the de Bruijn cuckoo rule is guaranteed to be executed in a correct and timely manner once a single honest peer has received a correct $(y_i, \{(y_j^{(i)})_{p_j} \mid v_j \in P_i\})_{p_i}$ message. We assume that the new node u can choose to assume any one of the new positions of a successfully executed de Bruijn cuckoo rule. It just needs to commit to one to R_v . If the node v just wants to rejoin the system (like in the adversarial strategies considered here), then we identify v with u .

3.5 Perturbation with biased randomness

Next we analyze the round-robin cuckoo rule. Recall that we consider adversarial join-leave attacks in a system with n honest nodes and ϵn adversarial nodes. Let β be the bias of the round-robin RNG. Then it holds:

Theorem 3.5 *For any constants ϵ , k and β with $\epsilon < (1/\beta)(1/\beta - 1/k)$, the round-robin cuckoo rule with the round-robin RNG with bias β satisfies the balancing and majority conditions for a polynomial number of rounds, with high probability, for any adversarial strategy within our model.*

Proof. We recall the analysis of the cuckoo rule in [4]. In the following, the time is counted in *rounds*. Each successful application of the de Bruijn cuckoo rule defines a round. Thus, an application of the round-robin cuckoo rule can have between $m - 2t$ and m rounds, where m is the number of nodes in the region in which the join operation is initiated and t is the number of adversarial nodes in that region.

Let \hat{R} be any fixed region of size $(c \log n) \cdot k/n$, for some constant c , for which we want to check the balancing and majority conditions over polynomial in n many rejoin operations. Thus, \hat{R} contains exactly $c \log n$ many k -regions. The *age* of a k -region is the difference between the current round and the last round when a new node was placed into it (and all old nodes got evicted), and the age of \hat{R} is defined as the sum of the ages of its k -regions. A node in \hat{R} is called *new* if it was placed in \hat{R} when it joined the system, and otherwise it is called *old*.

We assume that before the adversary starts with its rejoin operations, only the n blue nodes were in the system, and sufficiently many rejoin operations have been executed on the blue nodes so that every k -region has been entered by a new node at least once. Afterwards, the adversary enters with its ϵn red nodes one by one, using the round-robin cuckoo rule, and then it starts executing rejoin operations on the red nodes as it likes. The assumption of acting on a sufficiently old system significantly simplifies the proofs.

The next lemma follows directly from the round-robin cuckoo rule because every k -region can have at most one new node at any time.

Lemma 3.6 *At any time, \hat{R} contains at most $c \log n$ new nodes.*

In order to bound the number of old nodes in \hat{R} , we first have to bound the age of \hat{R} (Lemma 3.7). Then we bound the maximum number of nodes in a k -region (Lemma 3.8) and use this to bound the number of evicted blue and red nodes in a certain time interval (Lemma 3.9). After that, we can combine all lemmas to bound the number of old blue and red nodes in \hat{R} (Lemma 3.10).

Lemma 3.7 *At any time, \hat{R} has an age between $(1 - \delta)(1/\beta)(c \log n)(n/k)$ and $(1 + \delta)\beta(c \log n)(n/k)$, with high probability, where $\delta > 0$ is a constant that can be made arbitrarily small depending on the constant c .*

Proof. Let R_1, \dots, R_C be the k -regions of \hat{R} , where $C = c \log n$. For every k -region R_i , let the random variable X_i denote the age of R_i at the beginning of the given round, and let $X = \sum_{i=1}^C X_i$. We focus on a particular k -region R_i and consider the two extreme cases for it. In case 1, the adversary aims at minimizing the lifetime of R_i , which it achieves by adaptively letting all of its key generations fail that do not cause R_i to be replaced, and in case 2, the adversary aims at maximizing the lifetime of R_i , which it achieves by adaptively letting all of its key generations fail that do cause R_i to be replaced. In addition to these adaptive failures, the adversary also aims at causing failures of as many key generations by honest nodes as possible in order to maximize the impact of its adaptive strategy. When considering these two cases in the following, we will assume for simplicity that m (the number of nodes) and t (the number of adversarial nodes) are the same for all regions in which the adversary initiates a join operation. This will be fine for our calculations as long as we assume the worst case ratio between m and t .

Case 1: In order to prove a lower bound for the lifetime of R_i , let us assume for the moment that there are $m - 2t$ successful key generations by the blue nodes and t successful key generations by the red nodes, independent on whether the generated keys affect R_i or not. Then it holds that $\Pr[X_i = t] = (k/n)(1 - (k/n))^{t-1}$ and, therefore, $E[X_i] = n/k$. Furthermore, using the proof of Lemma 2.5 in [4], it follows that the age of \hat{R} would be at least $(1 - \delta)(c \log n)(n/k)$, w.h.p. However, of the t adversarial key generations, only those will succeed that generate a key in R_i , which implies that of the $m - t$ key generations considered above, up to t may get canceled. Thus, the age of \hat{R} can reduce up to a $(m - 2t)/(m - t) = 1/\beta$ -factor. Combining this with the idealistic lower bound on \hat{R} gives a lower bound of $(1 - \delta)(1/\beta)(c \log n)(n/k)$ that holds w.h.p.

Case 2: In order to prove an upper bound for the lifetime of R_i , let us assume as a worst case that none of the keys generated by adversarial nodes will hit R_i . If we only focus on the $m - 2t$ successful key generations by blue nodes, then it holds that $\Pr[X_i = t] = (k/n)(1 - (k/n))^{t-1}$ and, therefore, $E[X_i] = n/k$. Furthermore, using the proof of Lemma 2.5 in [4], it follows that the age of \hat{R} would be at most $(1 + \delta)(c \log n)(n/k)$, w.h.p. However, when also considering the t key generations by red nodes, the age of \hat{R} can increase up to a $(m - t)/(m - 2t) = \beta$ -factor. Combining this with the idealistic upper bound on \hat{R} gives an upper bound of $(1 + \delta)\beta(c \log n)(n/k)$ that holds w.h.p. \square

Lemma 3.8 *For any k -region R in \hat{R} it holds at any time that R has at most $O(k \log n)$ nodes, with high probability.*

Proof. Following the arguments of Lemma 2.6 in [4], there are at most $(1 + \epsilon)n \cdot (1 + \delta)\gamma \ln n$ node replacements during the lifetime of a k -region, w.h.p., for some constants $\epsilon, \delta > 0$. Hence, the expected number of nodes in a k -region can be at most

$$\frac{k}{n} \cdot ((1 + \epsilon)n \cdot (1 + \delta)\gamma \ln n + 1) = (1 + \epsilon)(1 + \delta)k \cdot \gamma \ln n + 1.$$

Since the locations of the node replacements are independent (if they are due to different rounds) or negatively correlated (if they happen at the same round), it follows from the Chernoff bounds for negatively correlated variables [25] that the number of nodes in a k -region is at most $O(k \log n)$ at any time, w.h.p. \square

Next we bound the number of blue and red nodes that are evicted in a certain time interval.

Lemma 3.9 For any time interval I of size $T = (\gamma/\epsilon) \log^3 n$, the number of blue nodes that are evicted in I is within $(1 \pm \delta)T \cdot k$, with high probability, and the number of red nodes that are evicted in I is within $(1 \pm \delta)T \cdot \epsilon k$, with high probability, where $\delta > 0$ can be made arbitrarily small depending on γ .

Proof. The proof is exactly the same as the proof of Lemma 2.9 in [4]. \square

Combining Lemmas 3.7 to 3.9, we obtain the following lemma.

Lemma 3.10 At any time, \hat{R} has between $(1 - \delta)(1/\beta)(c \log n) \cdot k$ and $(1 + \delta)\beta(c \log n) \cdot k$ old blue nodes and between $(1 - \delta)(1/\beta)(c \log n) \cdot \epsilon k$ and $(1 + \delta)\beta(c \log n) \cdot \epsilon k$ old red nodes, with high probability, where the lower bound on the red nodes holds if none of the red nodes has rejoined.

Proof. Consider any age distribution t_1, \dots, t_C for the k -regions R_1, \dots, R_C of \hat{R} , where $C = c \log n$. Let $T = (\gamma/\epsilon) \log^3 n$ be selected as in Lemma 3.9. Under the assumption that the locations of the replaced nodes are chosen independently at random, it follows from Lemma 3.9 and the Chernoff bounds that \hat{R} has at least

$$\frac{(1 + \delta)T \cdot k}{n/k} \sum_{i=1}^C \lfloor t_i/T \rfloor \geq \frac{(1 + \delta)k^2}{n} \left(\sum_{i=1}^C t_i - C \cdot T \right)$$

blue nodes and at most

$$\frac{(1 - \delta)T \cdot k}{n/k} \sum_{i=1}^C \lceil t_i/T \rceil \leq \frac{(1 - \delta)k^2}{n} \left(\sum_{i=1}^C t_i + C \cdot T \right)$$

blue nodes, w.h.p. Unfortunately, due to the use of the de Bruijn cuckoo rule, the locations of the nodes are not completely independent. However, from Lemma 3.8 we know that a k -region contains at most $bk \log n$ nodes for some constant b , w.h.p., and under the assumption of independent replacements, the probability that any two of these nodes are placed into \hat{R} when that k -region gets evicted is at most $\binom{bk \log n}{2} ((c \log n) \cdot \frac{k}{n})^2$. Using this bound, it is easy to show that over $O((n/k) \log n)$ many rounds (the maximum lifetime of a k -region, w.h.p.) there will only be $o(\log n)$ many pairs of nodes from the same k -region that end up in \hat{R} , w.h.p. This number affects the upper and lower bounds above only in a negligible way when choosing the bounds for $\sum_{i=1}^C t_i$ as given in Lemma 3.7. Hence, up to a $(1 + o(1))$ factor the bounds above also apply to the de Bruijn cuckoo rule. Since $\sum_{i=1}^C t_i$ is between $(1 - \delta')(1/\beta)Cn/k$ and $(1 + \delta')\beta Cn/k$ according to Lemma 3.7, Lemma 3.10 follows for the blue nodes.

The same calculations (with an additional ϵ factor) apply to the red nodes. \square

Combining Lemmas 3.6 and 3.10, we can now prove when the balancing and majority conditions are satisfied.

- **Balancing condition:** From Lemmas 3.6 and 3.10 it follows that every region R of size $(c \log n)k/n$ has at least $(1 - \delta)(1/\beta)(c \log n) \cdot k$ and at most $(1 + \delta)\beta(c \log n) + (c \log n)k + (c \log n)\epsilon k = (1 + \delta)\beta(c \log n)(1 + (1 + \epsilon)k)$ nodes, where the constant $\delta > 0$ can be made arbitrarily small. Hence, the regions are balanced within a factor of close to $\beta^2(1 + \epsilon + 1/k)$.
- **Majority condition:** From Lemmas 3.6 and 3.10 it also follows that every region of size $(c \log n)k/n$ has at least $(1 - \delta)(1/\beta)(c \log n) \cdot k$ blue nodes and at most $(1 + \delta)(c \log n) + \beta(c \log n) \cdot \epsilon k$ red nodes, w.h.p., where the constant $\delta > 0$ can be made arbitrarily small. Hence, the adversary is not able to obtain the majority in any region of size $(c \log n)k/n$ as long as $(c \log n)(\beta \cdot \epsilon k + 1) < (1/\beta)(c \log n) \cdot k$ which is true if and only if $\epsilon < (1/\beta)(1/\beta - 1/k)$.

Hence, for $\epsilon < (1/\beta)(1/\beta - 1/k)$ the balancing and majority conditions are satisfied, w.h.p., which proves Theorem 3.5. \square

4 Conclusions

In this paper, we presented a simple and robust random number generator sufficient for keeping honest and adversarial peers well-distributed in $[0, 1)$. We only proved our results assuming a sequential execution of rejoin operations (see our model) though we expect that as long as not too many rejoin operations are executed concurrently, there should be only insignificant side effects (see also the comments in [26]).

Interesting problems for future work are how to extend our results to general β -biased m -RNGs and how to extend our RNGs to the situation in which the system can be under DoS attacks.

References

- [1] J. Aspnes and G. Shah. Skip graphs. In *Proc. of the 14th ACM Symp. on Discrete Algorithms (SODA)*, pages 384–393, 2003.
- [2] B. Awerbuch and C. Scheideler. Group Spreading: A protocol for provably secure distributed name service. In *Proc. of the 31st International Colloquium on Automata, Languages and Programming (ICALP)*, 2004.
- [3] B. Awerbuch and C. Scheideler. Robust distributed name service. In *Proc. of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS)*, 2004.
- [4] B. Awerbuch and C. Scheideler. Towards a scalable and robust DHT. In *Proc. of the 18th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, 2006. See also <http://www14.in.tum.de/personen/scheideler>.
- [5] M. Ben-Or, B. Kelmer, and T. Rabin. Asynchronous secure computations with optimal resilience. In *Proc. of the 13th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 183–192, 1994.
- [6] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. Wallach. Security for structured peer-to-peer overlay networks. In *Proc. of the 5th Usenix Symp. on Operating Systems Design and Implementation (OSDI)*, 2002.
- [7] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proc. of the 2nd Usenix Symp. on Operating Systems Design and Implementation (OSDI)*, 1999.
- [8] S. Crosby and D. Wallach. Denial of service via algorithmic complexity attacks. In *Usenix Security*, 2003.
- [9] J. R. Douceur. The sybil attack. In *Proc. of the 1st International Workshop on Peer-to-Peer Systems (IPTPS)*, 2002.
- [10] A. Fiat, J. Saia, and M. Young. Making Chord robust to Byzantine attacks. In *Proc. of the European Symposium on Algorithms (ESA)*, 2005.
- [11] O. Goldreich. *Modern Cryptography, Probabilistic Proofs and Pseudorandomness*, volume 17 of *Algorithms and Combinatorics*. Springer-Verlag, 1998.
- [12] S. Halevi and S. Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *CRYPTO 96*, pages 201–215, 1996.
- [13] D. Karger, E. Lehman, T. Leighton, M. Levine, D. Lewin, and R. Panigrahi. Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the World Wide Web. In *29th ACM Symp. on Theory of Computing (STOC)*, pages 654–663, 1997.
- [14] V. King, J. Saia, V. Sanwalani, and E. Vee. Towards a secure and scalable computation in peer-to-peer networks. In *Proc. of the 47th IEEE Symp. on Foundations of Computer Science (FOCS)*, 2006.
- [15] F. Kuhn, S. Schmid, and R. Wattenhofer. A self-repairing peer-to-peer system resilient to dynamic adversarial churn. In *Proc. of the 4th International Workshop on Peer-to-Peer Systems (IPTPS)*, 2005.
- [16] M. Luby. *Pseudorandomness and Cryptographic Applications*. Princeton University Press, 1996.
- [17] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991.
- [18] M. Naor and U. Wieder. Novel architectures for P2P applications: the continuous-discrete approach. In *Proc. of the 15th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, 2003.

- [19] S. Nielson, S. Crosby, and D. Wallach. Kill the messenger: A taxonomy of rational attacks. In *Proc. of the 4th International Workshop on Peer-to-Peer Systems (IPTPS)*, 2005.
- [20] G. Plaxton, R. Rajaraman, and A.W. Richa. Accessing nearby copies of replicated objects in a distributed environment. In *Proc. of the 9th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 311–320, 1997.
- [21] H.V. Ramasamy and C. Cachin. Parsimonious asynchronous Byzantine-fault-tolerant atomic broadcast. In *9th Conference on Principles of Distributed Systems (OPODIS)*, pages 88–102, 2005.
- [22] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz. Handling churn in a DHT. In *USENIX Annual Technical Conference*, 2004.
- [23] T. Ritter. RNG implementations: A literature survey. <http://www.ciphersbyritter.com/RES/RNGENS.HTM>.
- [24] J. Saia, A. Fiat, S. Gribble, A. Karlin, and S. Saroiu. Dynamically fault-tolerant content addressable networks. In *Proc. of the 1st International Workshop on Peer-to-Peer Systems (IPTPS)*, 2002.
- [25] C. Scheideler. *Probabilistic Methods for Coordination Problems*. HNI-Verlagsschriftenreihe 78, University of Paderborn, 2000. See also <http://www14.in.tum.de/personen/scheideler>.
- [26] C. Scheideler. How to spread adversarial nodes? Rotate! In *Proc. of the 37th ACM Symp. on Theory of Computing (STOC)*, pages 704–713, 2005.
- [27] A. Singh, M. Castro, A. Rowstron, and P. Druschel. Defending against Eclipse attacks on overlay networks. In *Proc. of the 11th ACM SIGOPS European Workshop*, 2004.
- [28] E. Sit and R. Morris. Security considerations for peer-to-peer distributed hash tables. In *Proc. of 1st International Workshop on Peer-to-Peer Systems (IPTPS)*, 2002.
- [29] K. Srinathan and C.P. Rangan. Efficient asynchronous secure multiparty distributed computation. In *Proc. of the 1st Int. Conference on Progress in Cryptology*, pages 117–129, 2000.
- [30] M. Srivatsa and L. Liu. Vulnerabilities and security threats in structured overlay networks: A quantitative analysis. In *Proc. of the 20th IEEE Computer Security Applications Conference (ACSAC)*, 2004.
- [31] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for Internet applications. In *Proc. of the ACM SIGCOMM '01*, 2001. See also <http://www.pdos.lcs.mit.edu/chord/>.
- [32] J. Viega. Practical random number generation in software. In *Proc. of the 19th Annual Computer Security Applications Conference*, 2003.