

HASHGRAPH CONSENSUS: FAIR, FAST, BYZANTINE FAULT TOLERANCE

LEEMON BAIRD
BAIRD@SWIRLDS.COM
MAY 31, 2016
SWIRLDS TECH REPORT TR-2016-01

ABSTRACT. A new system, the *Swirls hashgraph consensus algorithm*, is proposed for replicated state machines with guaranteed Byzantine fault tolerance. It achieves *fairness*, in the sense that it is difficult for an attacker to manipulate which of two transactions will be chosen to be first in the consensus order. It has complete asynchrony, no leaders, no round robin, no proof-of-work, eventual consensus with probability one, and high speed in the absence of faults. It is based on a gossip protocol, in which the participants don't just gossip about transactions. They *gossip about gossip*. They jointly build a *hashgraph* reflecting all of the gossip events. This allows Byzantine agreement to be achieved through *virtual voting*. Alice does not send Bob a vote over the Internet. Instead, Bob calculates what vote Alice would have sent, based on his knowledge of what Alice knows. This yields fair Byzantine agreement on a total order for all transactions, with very little communication overhead beyond the transactions themselves.

Keywords: Byzantine, Byzantine agreement, Byzantine fault tolerance, replicated state machine, fair, fairness, hashgraph, gossip about gossip, virtual voting, Swirls

CONTENTS

List of Figures	2
1. Introduction	2
2. Core concepts	4
3. Gossip about gossip: the hashgraph	5
4. Consensus algorithm	6
5. Proof of Byzantine fault tolerance	11
6. Fairness	16
7. Generalizations and enhancements	17
8. Conclusions	21
9. References	22
References	22
10. Appendix A: Consensus algorithm in functional form	23

LIST OF FIGURES

1	Gossip history as a directed graph	5
2	The hashgraph data structure	7
3	Illustration of strongly seeing.	8
4	Pseudocode: the Swirls hashgraph consensus algorithm	10
5	Pseudocode: the divideRounds procedure	11
6	Pseudocode: the decideFame procedure	12
7	Pseudocode: the finalOrder procedure	13

1. INTRODUCTION

Distributed databases are often required to be replicated state machines with Byzantine fault tolerance. Some authors have used the term “Byzantine” in a weak sense, such as assuming that attackers will not collude, or that communication is weakly asynchronous [1]. In this paper, “Byzantine” will be used in the strong sense of its original definition [2]: up to just under 1/3 of the members can be attackers, they can collude, and they can delete or delay messages between honest members with no bounds on the message delays. The attackers can control the network to delay and delete any messages, though at any time, if an honest member repeatedly sends messages to another member, the attackers must eventually allow one through. It is assumed that secure digital signatures exist, so attackers cannot undetectably modify messages. It is assumed that secure hash functions exist, for which collisions will never be found. This paper proposes and describes the Swirls hashgraph consensus algorithm, and proves Byzantine fault tolerance, under the strong definition.

No deterministic Byzantine system can be completely asynchronous, with unbounded message delays, and still guarantee consensus, by the FLP theorem [3]. But it is possible for a nondeterministic system to achieve consensus with probability one. The hashgraph consensus algorithm is completely asynchronous, is nondeterministic, and achieves Byzantine agreement with probability one.

Some systems, such as Paxos [4] or Raft [5] use a leader, which can make them vulnerable to large delays if an attacker launches a denial of service attack on the current leader [6]. Many systems can even be delayed by just a single bad client [7]. In fact, the latter paper suggests that systems with such vulnerabilities might better be described as “Byzantine fault survivable” rather than “Byzantine fault tolerant”. Hashgraph consensus does not use a leader, and is resilient to denial of service attacks on small subsets of the members.

Other systems, such as Bitcoin, are based on proof-of-work blockchains [8]. This avoids all the above problems. However, such systems cannot be Byzantine, because a member never knows for sure when consensus has been achieved; they only have a probability of confidence that continues to rise over time. If two blocks are mined simultaneously, then the chain will fork until the community can agree on which branch to extend. If the blocks are added slowly, then the community can always add to the longer branch, and eventually the other branch will stop growing, and can be pruned and discarded because it is “stale”. This leads to inefficiency, in the sense

that some blocks are mined properly, but discarded anyway. It also means that it is necessary to slow down how fast blocks are mined, so that the community can jointly prune branches faster than new branches sprout. That is the purpose of the proof-of-work. By requiring that the miners solve difficult computation problems to mine a block, it can ensure that the entire network will have sufficiently long delays between mining events, on average. The hashgraph consensus algorithm is equivalent to a block chain in which the “chain” is constantly branching, without any pruning, where no blocks are ever stale, and where each miner is allowed to mine many new blocks per second, without proof-of-work, and with 100% efficiency.

Proof-of-work blockchains also require that electricity be wasted on extra computations, and perhaps that expensive mining rigs be bought. A proof-of-expired-time system [9] can avoid the wasted electricity (though perhaps not the cost of mining rigs) by using trusted hardware chips that delay for long periods, as if they were doing proof-of-work computations. However, that requires that all participants trust the company that created the chip. Such trust in chip vendors exists in some situations, but not in others, such as when FreeBSD was changed to not rely solely on the hardware RDRAND instruction for secure random numbers, because “we cannot trust them any more” [10].

Byzantine agreement systems have been developed for Byzantine agreement that avoid the above problems. These systems typically exchange many messages for the members to vote. For n members to decide a single YES/NO question, some systems can require $O(n)$ messages to be sent across the network. Other systems can require $O(n^2)$, or even $O(n^3)$ messages crossing the network per binary decision [11]. An algorithm for a single YES/NO decision can then be extended to deciding a total order on a set of transactions, which may further increase the vote traffic. Hashgraph sends no votes at all over the network, because all voting is virtual.

2. CORE CONCEPTS

The hashgraph consensus algorithm is based on the following core concepts.

- *Transactions* - any member can create a signed transaction at any time. All members get a copy of it, and the community reaches Byzantine agreement on the order of those transactions.
- *Fairness* - it should be difficult for a small group of attackers to unfairly influence the order of transactions that is chosen as the consensus.
- *Gossip* - information spreads by each member repeatedly choosing another member at random, and telling them all they know
- *Hashgraph* - a data structure that records who gossiped to whom, and in what order.
- *Gossip about gossip* - the hashgraph is spread through the gossip protocol. The information being gossiped is the history of the gossip itself, so it is “gossip about gossip”. This uses very little bandwidth overhead beyond simply gossiping the transactions alone.
- *Virtual voting* - every member has a copy of the hashgraph, so Alice can calculate what vote Bob *would have* sent her, if they had been running a traditional Byzantine agreement protocol that involved sending votes. So Bob doesn’t need to actually her the vote. Every member can reach Byzantine agreement on any number of decisions, without a single vote ever being sent. The hashgraph alone is sufficient. So zero bandwidth is used, beyond simply gossiping the hashgraph.
- *Famous witnesses* - The community could put a list of n transactions into order by running separate Byzantine agreement protocols on $O(n \log n)$ different yes/no questions of the form “did event X come before event Y?” A much faster approach is to pick just a few events (vertices in the hashgraph), called *witnesses*, and define a witness to be *famous* if the hashgraph shows that most members received it fairly soon after it was created. Then its sufficient to run the Byzantine agreement protocol only for witnesses, deciding for each witness the single question “is this witness famous?” Once Byzantine agreement is reached on the exact set of famous witnesses, it is easy to get from the hashgraph a fair total order for all events.
- *Strongly seeing* - given any two vertices X and Y in the hashgraph, it can be immediately calculated whether X can strongly see Y, which is defined to be true if they are connected by multiple directed paths passing through enough members. This concept allows the key lemma to be proved: that if Alice and Bob are both able to calculate Carol’s virtual vote on a given question, then Alice and Bob get the same answer. That lemma forms the foundation for the rest of the mathematical proof of Byzantine convergence with probability one.




FIGURE 1. Gossip history as a directed graph. The history of any gossip protocol can be represented by a graph where each member is one column of vertices. When Alice receives gossip from Bob, telling her everything he knows, that gossip event is represented by a vertex in the Alice column, with two edges going downward to the immediately preceding gossip events by Alice and Bob.

3. GOSSIP ABOUT GOSSIP: THE HASHGRAPH

Hashgraph consensus uses a gossip protocol. This means that a member such as Alice will choose another member at random, such as Bob, and then Alice will tell Bob all of the information she knows so far, and Bob will tell her all he knows. Alice then repeats with a different random member, and all other members do the same. In this way, if a single member becomes aware of new information, it will spread exponentially fast through the community until every member is aware of it.

The history of any gossip protocol can be illustrated by a directed graph like Figure 1. Each vertex in the Alice column represents a gossip event. For example, the top event in the Alice column represents Bob performing a gossip sync with Alice in which Bob sent her all of the information that he knew. That vertex has two downward edges, connecting to the immediately-preceding gossips for Alice and Bob. Time flows up the graph, so lower vertices represent earlier events in history. In a typical gossip protocol, a diagram such as this is merely used to discuss the protocol; there is no actual graph like that stored in memory anywhere.

In hashgraph consensus, that graph is an actual data structure. Figure 2 illustrates this data structure. Each *event* (vertex) is stored in memory as a sequence of bytes, signed by its creator. For example, one event by Alice (red) records the fact that Bob performed a gossip sync in which he told her everything he knew. This event is signed by her, and contains the hashes of two other events: her last event and Bob's last event prior to that gossip sync. The red event can also contain a *payload* of any transactions that Alice chose to create at that moment, and perhaps a timestamp which is the time and date that Alice claims to have created it. Other than those two, the ancestors of that event (gray) are not contained within it, but are determined by the set of cryptographic hashes. Data structures with graphs of hashes have been used for other purposes, such as in Git where the vertices are versions of a file tree, and the edges represent changes, and there is no record of how

members communicated. But the hashgraph is for a different purpose. It records the history of how the members communicated.

Gossip protocols are widely used to transfer a variety of types of information. They can involve gossiping about user identities, or gossiping about transactions, or gossiping about blockchain blocks, or gossiping about any other information that needs to be distributed. But what if the protocol were to gossip about gossip? What if the members were gossiping to transfer the hashgraph itself? When Bob gossiped to Alice, he would give her all of the events which he knew and she did not.

Gossiping a hashgraph gives the participants a great deal of information. If a new transaction is placed in the payload of an event, it will quickly spread to all members, until every member knows it. In addition, Alice will not only know of the transaction, she will know exactly when Bob learned of the transaction. And she will know exactly when Carol learned of the fact that Bob had learned of that transaction. Deep chains of such reasoning become possible when all members have a copy of the hashgraph. As the hashgraph grows upward, the different members may have slightly different subsets of the new events near the top, but they will quickly converge to having exactly the same events lower down in the hashgraph. Furthermore, if Alice and Bob happen to both have a given event, then they are guaranteed to also both have all its ancestors. And they will agree on all the edges in the subgraph of those ancestors. All of this allows powerful algorithms to run locally, including for Byzantine fault tolerance.

This power comes with very little communication overhead. If a community is simply gossiping signed transactions that they create, there is a certain amount of bandwidth required. If they instead gossip a hashgraph, and if there are enough transactions that a typical event contains at least one transaction, then the overhead is minimal. Instead of Alice signing a transaction she creates, she will sign the event she creates to contain that transaction. Either way, she is only sending one signature. And either way, she must send the transaction itself. The only required overhead is that she must send the two hashes. But even that can be greatly compressed. In figure 2, Alice will not send Carol the red event until Carol already has all its ancestors. So Alice does not need to send the two hashes of the two blue parent events. It is sufficient to tell Carol that this event is the next one by Alice, and that its other-parent is the third one by Bob. With appropriate compression, this can be sent in very few bytes, adding only a few percent to the size of the message being sent.

4. CONSENSUS ALGORITHM

It is not enough to ensure that every member knows every event. It is also necessary to agree on a linear ordering of the events, and thus of the transactions recorded inside the events. Most Byzantine fault tolerance protocols without a leader depend on members sending each other votes. So for n members to agree on a single YES/NO question might require $O(n^2)$ voting messages to be sent over the network, as every member tells every other member their vote. Some of these protocols require receipts on votes, making them $O(n^3)$. And they may require multiple rounds of voting, which further increases the number of voting messages sent.




FIGURE 2. The hashgraph data structure. Alice creates an *event* (red) recording the occurrence of Bob doing a gossip sync with her and telling her everything he knows. The event contains a hash of two parent events (blue): the self-parent (dark blue) by the same creator Alice, and the other-parent (light blue) by Bob. It also contains a payload of any new transactions that Alice chooses to create at that moment, and a digital signature by Alice. Other than those two, the ancestor events (gray) are not stored in the red event, but they are determined by all the hashes. The self-ancestors (dark gray) are those reachable by sequences of self-parent links, and the others (light gray) are not.

Hashgraph consensus does not require any votes to be sent. Every member has a copy of the hashgraph. If Alice and Bob both have the same hashgraph, then they can calculate a total order on the events according to any deterministic function of that hashgraph, and they will both get the same answer. Therefore, consensus is achieved, even without sending vote messages.

Of course, Alice and Bob may not have exactly the same hashgraph at any given moment. They will typically match in the older events. But for the very recent events, each may have events that the other has not yet seen. Furthermore, there may occasionally be a new event released to the community that should be placed in a lower (earlier) location in the hashgraph. The hashgraph consensus algorithm deals with these issues using a system that is best thought of as *virtual voting*.

Suppose Alice has hashgraph A and Bob has hashgraph B . These hashgraphs may be slightly different at any given moment, but they will always be *consistent*. Consistent means that if A and B both contain event X , then they will both contain exactly the same set of ancestors for X , and will both contain exactly the same set of edges between those ancestors. If Alice knows of X and Bob does not, and both of them are honest and actively participating, then we would expect Bob to learn of X fairly quickly, through the gossip protocol. But the consensus algorithm does not make any assumptions about how fast that will happen. The protocol is completely asynchronous, and does not make assumptions about timeout periods, or the speed of gossip, or the rate at which progress is made.

Alice will calculate a total order on the events in A by calculating a series of *elections*. In each election, some of the events in A will be considered to cast a *vote*, and some of the events in A will be considered to receive that vote. Alice will




FIGURE 3. Illustration of strongly seeing. In each hashgraph, the yellow event at the top can *strongly see* one of the orange events on the bottom row. There are $n = 5$ members, so the least integer greater than $2n/3$ is 4. In (d), one event (orange) is an ancestor of each of 4 intermediate events by different creators (red), each of which is an ancestor of the yellow event. Therefore, the yellow event can strongly see the orange event. Each of the other hashgraphs is colored to show the same for a different orange event on the bottom row, which the yellow event see through at least 4 red events.

calculate multiple elections, and a given event might participate in some elections but not others, and might cast different votes in different elections. If the event was created by Bob, we will talk of Bob voting a certain way in a given election. But the actual member Bob is not involved. This is purely a calculation that Alice is performing locally, where she is calculating what vote Bob *would have* sent her, if the real Bob were actually sending votes over the internet to her.

This virtual voting has several benefits. In addition to saving bandwidth, it ensures that members always calculate their votes according to the rules. If Alice is honest, she will calculate virtual votes for the virtual Bob that are honest. Even if the real Bob is a cheater, he cannot attack Alice by making the virtual Bob vote incorrectly.

Bob can try to cheat in a different way. Suppose Bob creates an event x with a certain self-parent hash pointing to his previous event z . Then Bob creates a new event Y , but gives it a self-parent hash of z , instead of giving it a self-hash of x as he should. This means that the events by Bob in the hashgraph will no longer be a chain, as they should be. They will now be a tree, because he has created a *fork*. If Bob gossips x to Alice and y to Carol, then for a while, Alice and Carol may not be aware of the fork. And Alice may calculate a virtual vote for x that is different from Carol's virtual vote for y .

The hashgraph consensus algorithm prevents this attack by using the concept of one state *seeing* another, and the concept of one state *strongly seeing* another.

If Bob creates two events x and y , neither of which is a self-ancestor of the other, then Bob has cheated by forking. If some event w has x as an ancestor but doesn't have y as an ancestor, then the event w can *see* event x . However, if both x and y are ancestors of w , then w is defined to not see either of them. In other words, w can see x if x is known to it, and no forks with x are known to it.

If there are n members, then an event w can *strongly see* an event x , if w can see more than $2n/3$ events by different members, each of which can see x . This concept

is illustrated in Figure 3. Four copies of the same hashgraph are shown, each with a different event on the bottom row colored orange. In (d), the yellow event at the top can see 4 red events by different members, each of which can see the orange event at the bottom. This is also true in (a), (b), and (c), with (a) actually having 5 red events. But only 4 are needed for strongly seeing, because this example has $n = 5$ members, and the least integer greater than $2n/3$ is 4.

This concept allows an agreement protocol to achieve Byzantine fault tolerance without any actual voting, just through local virtual voting.

In virtual voting, a vote is calculated for an event x purely as a function of its ancestors. That vote will only be sent to its descent w if w can strongly see x . It is proved in section 5 that if x and y are on different branches of an illegal fork, then w can strongly see at most one of x and y , but not both. Furthermore, if hashgraphs A and B are consistent, then it is not possible for one event to strongly see x in A and another event strongly see y in B . That lemma is the cornerstone of the Byzantine proof. It ensures that even if an attacker tries to cheat by forking, they will still be unable to cause different members to decide on different orders. Historically, some Byzantine agreement algorithms have required members to send out “receipts” for each vote they receive, to defend against Alice sending inconsistent votes to Bob and Carol. There are some similarities between that attack and a hashgraph forking attack, and between the use of receipts and the use of strongly seeing.

Given those definitions, the complete hashgraph consensus protocol can be given by the algorithms in Figures 4, 5, 6, and 7.

The main algorithm in Figure 4 shows that the communication is very simple: members randomly pick other members, and gossip all events they know that the other doesn't, and vice versa. A new event is created to record the fact of each gossip. Any simple gossip protocol can be used. For example, Alice might first give Bob all her events, then Bob creates an event to record that. Bob then sends Alice all his events (including that new one). Alice then creates a new event to record that, and sends it back to Bob.

After each sync, the member calls the three procedures to determine the consensus order for as many events as possible. These involve no communication; purely local computations are sufficient. In these procedures, each *for* loop visits events in *topological order*, where an event is always visited after its parents. In the first *for* loop of the algorithm, if x is the first event in all of history, then it won't have parents or previous rounds, so it should be set to $x.\text{round}=1$ and $x.\text{witness}=\text{TRUE}$. The algorithm also uses a constant n , which is the number of members in the entire population, and c which is a small integer greater than 2, such as $c = 6$. In the following algorithm, Byzantine agreement is guaranteed with probability one.

In `divideRounds` (Figure 5), every known event is assigned an integer *round number* as a function of the round numbers of its ancestors. The hashgraphs in Figure 3 show how this is done. If all the events on the bottom row were round r , then all the rest of the events in those figures would also be round r , except for the yellow event, which would be round $r + 1$. The yellow event is advanced to the next round, $r + 1$, because it is able to strongly see more than $2n/3$ events from round r . The first event in history is defined to be round 1, so all future rounds are determined by this.

For any given member, their first event in each round is called a *witness*. It is only the witness events that send and receive the virtual votes. This occurs in the

```

loop
  sync with a randomly chosen member
  create a new event
  call divideRounds
  call decideFame
  call findOrder
end loop

```

FIGURE 4. The Swirls hashgraph consensus algorithm. Each member runs this in parallel. Each sync brings in new events, which are added to the hashgraph. All known events are then divided into rounds. Then the first events in each round are decided as being famous or not (though purely local Byzantine agreement with virtual voting). Then the total order is found on those events for which enough information is available. If two members independently assign a position in history to an event, they are guaranteed to assign the same position, and guaranteed to never change it, even as more information comes in. Furthermore, each event is eventually assigned such a position, with probability one.

decideFame procedure shown in Figure 6. This procedure is where the Byzantine agreement occurs. For each witness, it decides whether it is *famous*. A witness is famous if many of the witnesses in the next round can see it, and it is not famous if many can't. The Byzantine agreement protocol runs an election for each witness, to determine if it is famous. For a witness x in round r , each witness in round $r + 1$ will vote that x is famous if it can see it. If more than $2n/3$ agree on whether it is famous, then the community has decided, and the election is over. If the vote is more balanced, then it continues for as many rounds as necessary, with each witness in a normal round voting according to the majority of the witnesses that it can strongly see in the previous round. To defend against attackers who can control the internet, there are periodic *coin rounds* where witnesses can vote pseudorandomly. This means that even if an attacker can control all the messages going over the internet to keep the votes carefully split, there is still a chance that the community will randomly cross the $2n/3$ threshold. And so agreement is eventually reached, with probability one.

Once consensus has been reached on whether each witness in a given round is famous, it is then easy to use that to determine a total order on older events. This is done by procedure findOrder, found in Figure 7. All of the events are sorted by their *received round*. Event x has a received round of r if that is the first round in which at least half the famous witnesses could see it. If two events have the same received round, they are sorted by their *received time*. The received time for event x is the median of the timestamps of the events where each famous witness's creator first had an event that saw x . If there are still ties there, they are broken by simply sorting by signature, after the signature is whitened by XORing with the XOR of the signatures of all the famous witnesses in the received round.

```

procedure divideRounds

for each event  $x$ 
  if  $x$  has no parents ( $x$  is first event in history)
     $x$ .round  $\leftarrow$  1
     $x$ .witness  $\leftarrow$  TRUE
  else
     $x$ .round  $\leftarrow$  the maximum round number  $r$  such that
       $x$  can strongly see a set of events
      whose round is  $r-1$ , and where events
      in the set were created by more than
       $2n/3$  members
     $x$ .witness  $\leftarrow$  is  $x$ .round strictly greater than
      the round of each of its parents?

```

FIGURE 5. The divideRounds procedure. As soon as an event x is known, it is assigned a round number x .round, and the boolean value x .witness is calculated, indicating whether it is the first event that a member created in that round.

5. PROOF OF BYZANTINE FAULT TOLERANCE

In all of the following proofs it is assumed that there are n members, more than $2n/3$ are honest, and less than $n/3$ are not honest. It is also assumed that the digital signatures and cryptographic hashes are secure, so signatures cannot be forged, signed messages cannot be changed without detection, and hash collisions can never be found.

Definition 5.1. Event x is a *fork with* event y if they have the same creator, but neither is a self-ancestor of the other.

Definition 5.2. An *honest* member syncs infinitely often with every other member, creates a valid event after each sync with hashes of the latest self-parent and other-parent, and never creates two events that are forks with each other.

Definition 5.3. Hashgraphs A and B are *consistent* if for any event x , if both hashgraphs contain x , then they both contain the same set of ancestors for x , with the same subgraph for those ancestors.

Lemma 5.4. *All members have consistent hashgraphs.*

Proof: If two members have hashgraphs containing x , then they have the same two hashes contained within x . A member will not accept an event during a sync unless that member already has both parents for that event, so both hashgraphs must contain both parents for x . The cryptographic hashes are assumed to be secure, therefore the parents must be the same. By induction, all ancestors must be the same. Therefore the two hashgraphs are consistent. \square

The purpose of the concept of *strongly seeing* is to make the following lemma true. This lemma is the foundation of the entire proof, because it allows for consistent

```

procedure decideFame

for each event x
  for each event y
    if x.witness and y.witness and y.round>x.round
      d ← y.round - x.round
      s ← the set of witness events in round
          y.round-1 that y can strongly see
      v ← majority vote in s (is TRUE for a tie)
      t ← number of events in s with a vote of v
      if d = 1 // first round of the election
        y.vote ← can y see x?
      else if d mod c > 0 // this is a normal round
        if t > 2*n/3 // if supermajority, then decide
          x.famous ← v
          break out of the y loop
        else // else, just vote
          y.vote ← v
      else // this is a coin round
        if t > 2*n/3 // if supermajority, then vote
          y.vote ← v
        else // else flip a coin
          y.vote ← middle bit of y.signature

```

FIGURE 6. The decideFame procedure. For each witness event (i.e., an event x where $x.witness$ is true), decide whether it is famous (i.e., assign a boolean to $x.famous$). This decision is done by a Byzantine agreement protocol based on virtual voting. Each member runs it locally, on their own copy of the hashgraph, with no additional communication. It treats the events in the hashgraph as if they were sending votes to each other, though the calculation is purely local to a member's computer. The member assigns votes to the witnesses of each round, for several rounds, until more than $2/3$ of the population agrees.

voting, and for guarantees that different members will never calculate inconsistent results, even with purely virtual voting.

Lemma 5.5. *If event x is a fork with event y , and x is strongly seen by event z in hashgraph A , then y will not be strongly seen by any event in any hashgraph B that is consistent with A .*

Proof: The proof is by contradiction. Suppose event w in B can strongly see y . By the definition of strongly seeing, there must exist a set S_A of events in A that z can see, and that all can see x . There must be a set S_B of events in B that w can see, and which all see y . Then S_A must contain events created by more than

```

procedure findOrder

for each event x
  if there is a round r where there is no event x
    such that x.witness and x.famous=UNDECIDED
  and at least 1/2 of round r famous witnesses see x
  and this is not true of any round earlier than r
    x.roundReceived ← r
    s ← set of round r famous witnesses that see x
    t ← set of the earliest self-ancestor that
      sees x, for each event in s that sees x
    x.timestamp ← median of the timestamps of
      all the events in t

return all events that have roundReceived not UNDECIDED,
sorted by roundReceived, then ties sorted by timestamp,
then those ties sorted by whitened signature

```

FIGURE 7. The findOrder procedure. Now that it is known exactly which witnesses are famous in each round, they act as the judges to assign every other event a round received and consensus timestamp. An event is said to be “received” in the first round where at least half the famous witnesses have received it. And its timestamp is the median of the timestamps of those events where each of those members first received it. Once these have been calculated, the events are sorted by round received. Any ties are subsorted by consensus timestamp. Any remaining ties are subsorted by signature.

$2n/3$ members, and so must S_B , therefore there must be an overlap of at least $n/3$ members who created events in both sets. It is assumed that less than $n/3$ members are not honest, so there must be at least one honest member who created events in both S_A and S_B . Let M be such a member, and their events $q_A \in S_A$ and $q_B \in S_B$. Because M is honest, q_A and q_B cannot be forks with each other, so one must be the self-ancestor of the other. Without loss of generality, let q_A be the self-ancestor of q_B . Then x is an ancestor of q_A , which is an ancestor of q_B , therefore x is an ancestor of q_B . But y is also an ancestor of q_B , so both x and y are ancestors of q_B and are forks of each other, so q_B cannot see either of them. But that contradicts the assumption that q_B can see y . That is a contradiction, so the lemma is proved. \square

Lemma 5.6. *If hashgraphs A and B are consistent and both contain event x , then both will assign the same round number to x .*

Proof: If the consistent hashgraphs both contain x , then they both contain all its ancestors, including the first event in history. Then the proof is by induction: they agree on the round number of that first event, which is 1 by definition. And if

they both contain an arbitrary state y , and agree on the round numbers of all its ancestors, then they will agree on the maximum round number r for which y can strongly see more than $2n/3$ ancestors with round r by different members, and therefore will agree on the round number of y . Therefore they will agree on the round number of all events they share, including x . \square

Different members may have slightly different hashgraphs, and so may have slightly different elections. However, all the votes will be consistent. If one hashgraph shows Alice sending Bob a given vote in a given round, then any consistent hashgraph must show either the same vote, or no vote at all from Alice to Bob in that round. It is impossible for two consistent hashgraphs to show two different votes for Alice in that round. This is shown in the following lemma.

Lemma 5.7. *If hashgraphs A and B are consistent, and the algorithm running on A shows that a round r event by member m_0 sends a vote v_A to member m_1 in round $r+1$, and the algorithm running on B shows that a round r event by member m_0 sends a vote v_A to an event by member m_1 in round $r+1$, then $v_A = v_B$.*

Proof: The algorithm only sends a vote from event x to event y if y can strongly see x . It is not possible for consistent hashgraphs to have two events that are forks of each other and that are both strongly seen, by lemma 5.5. Therefore, the two votes must be coming from the same event x in both hashgraphs. An event's vote is calculated purely as a function of its ancestors, so the two hashgraphs must agree on the vote, and $v_A = v_B$. \square

Byzantine agreement on a particular YES/NO question is achieved by multiple rounds of virtual voting. A given member will end their election calculations in round r if it is a normal round (not a coin round) and some round $r+1$ event strongly sees at least $2n/3$ of the members voting the same way in round r . If that happens, then every other member will end their election in round r or $r+1$ (or $r+2$ if $r+1$ is a coin round), and will decide the same way. In other words, the following lemma proves that if anyone decides on a YES/NO question, then everyone achieves Byzantine consensus almost immediately thereafter.

Lemma 5.8. *If hashgraphs A and B are consistent, and A decides a Byzantine agreement election with result v in round r and B has not decided prior to r , then B will decide v in round $r+2$ or before.*

Proof: Decisions can't happen in coin rounds, so r must be a regular round. If A decides a vote v , that means some witness in round r received votes of v from a set of members S that contains than $2n/3$ members. Because voting is consistent (by the previous lemma), all other round r events in A and B will receive votes from at least $2n/3$ members, a majority of whom will also be in S , because two subsets of size greater than $2n/3$ drawn from a set of size n must each have a majority of their elements in common with the other. Therefore, every round r witness in both A and B will either vote for v or decide for v . If round $r+1$ is a regular round, then every event in A and B in that round will receive unanimous votes of v and will decide v . If round $r+1$ is a coin round, then all will receive unanimous votes of v , so none will flip coins, and all will vote v , and then all will decide v in round $r+2$. \square

The following theorem shows that Byzantine fault tolerance is achieved for any single YES/NO question.

Theorem 5.9. *For any single YES/NO question, consensus is achieved eventually with probability 1.*

Proof: If any member decides the question, then all members will decide the same way within 2 rounds, by the last lemma. So the only way consensus could fail is if no member ever decides, because no witness ever receives more than $2n/3$ matching votes. However, in a coin round, if such a supermajority has not yet been achieved, then all the honest members randomly choose their vote, and will unanimously choose the vote with nonzero probability. Coin rounds occur periodically forever, so eventually the honest members will become unanimous, with probability one, and then consensus will be reached within 2 rounds. \square

In the hashgraph consensus algorithm, Byzantine agreement is used to decide whether each witness in a given round is famous or not. Every round is guaranteed to have at least one witness that is famous, by the following lemma.

Lemma 5.10. *In any round r , at least one witness will be decided to be famous by the consensus algorithm, and this decision will be made within 2 rounds.*

Proof: Each round $r + 1$ event can strongly see more than $2n/3$ round r events, by the definition of round number. Therefore, each round r event is strongly seen by more than $2/3$ of the round $r + 1$ events, on average. Not all of the round r events can be below average, so there must be at least one round r event x that is strongly seen by more than $2/3$ of the round $r + 1$ events. For the vote on whether x is famous, those round $r + 1$ events will vote YES because they can see x . Therefore, every round $r + 2$ event will receive a majority of its votes saying x is famous, because it receives votes from at least $2n/3$ round $r + 1$ events, and more than $2/3$ of the round $r + 1$ events are voting YES. Therefore, every member will decide YES in round $r + 2$ or $r + 3$. \square

Given the last lemma and the theorem before it, we know that every round will eventually have all its witnesses classified as famous or not famous by universal consensus, with at least one of the witnesses being famous. This set of famous witnesses can then act as a judge to find a total order on all the events that have reached most of them.

Theorem 5.11. *Each event x created by an honest member will eventually be assigned a consensus position in the total order of events, with probability 1.*

Proof: All honest members will eventually learn of x , by the definition of honest and the assumptions that the attackers who control the internet must eventually allow any two honest members to communicate. Therefore, there will eventually be a round where all the famous witnesses have received x . Therefore in that round, or possibly earlier, there will be a round r where at least half of the famous witnesses have received x , and x is assigned a received round of r , and a consensus timestamp of the median of when those members first received it, and its place in history will be fixed. Furthermore, it is not possible to later discover a new event y that will come before x in the consensus order. Because y would have to have a received round less than or equal to r , which means that at least half the famous witnesses in round r must have received y . But once the set of famous witnesses is known for a round, all of their ancestors are known, so there is no way to discover new ancestors for them in the future. Furthermore, it isn't possible for a round to gain

new famous witnesses in the future, once the famousness of all the known witnesses in that round are known. Because any new round r witness that is discovered in the distant future will *not* be an ancestor of the known round $r + 1$ witnesses (of which there are more than $2n/3$, and so the consensus will immediately be reached that it is not famous. Therefore, once an event is assigned a place in the total order, it will never change its position, either by swapping with another known event, nor by new events being discovered later and being inserted before it. \square

6. FAIRNESS

Most existing systems for distributed consensus can fail to be “fair” in their consensus ordering of transactions. To see this, first consider a stock market that is run by a single server. Alice and Bob each submit a bid to that server, with Alice submitting it just before Bob. If the server is *fair*, then it will count Alice’s transaction as occurring before Bob’s. For some applications, the exact order does not matter, but for a stock market it can be critically important that this decision be made fairly.

Now consider a distributed peer-to-peer system, where there is no single server, but there is a community that will reach consensus on whose transaction was first. It may still be critically important that the consensus decision is fair. But what should be the definition of “fair”?

The “fair” decision on transaction order could be defined as favoring whichever transaction was created first. But that would be bad. Alice might have created her transaction one second before Bob, while she was in a cabin in the woods, disconnected from the internet. Then the community would only hear of Bob’s transaction, and would assume that Bob was first. A year later, when Alice finally emerges from the woods and rejoins the internet, the community would have to revise history in order to be “fair”. That would cause a host of problems. So that wouldn’t be an ideal definition of fairness. There needs to be a requirement that the transaction actually be sent to the community, in order to count as being first.

The “fair” decision could be defined as reflecting the order in which the transactions reached the current *leader*. But that would also be bad. The leader might be a member chosen by the Paxos algorithm. Or it might be whichever member currently has a turn in a round-robin system. In a proof-of-work system, it would be whichever miner manages to solve a puzzle first. In any case, the leader could arbitrarily decide to ignore either Alice’s or Bob’s transaction for a period of time, delaying one of them, to force their transaction to come after the other. If the goal is distributed trust, then no single individual can be trusted.

The “fair” decision could be defined as reflecting when each transaction first reached a certain fraction of the entire community. This is a little better. The community is then ordering transactions by when the transactions first reached a virtual server, where “reaching the server” means reaching some fraction of the community as a whole. However, there are still issues. If the fair choice is defined as whichever transaction reached more than half the community first, then there will be problems if Carol saw Alice first, Dave saw Bob first, and everyone else is evenly split on the question. This fails if Carol and Dave are both attackers who turn off their computers permanently before telling the community what they saw. In that case, the community could never reach a fair consensus, because they would be waiting forever on Carol and Dave.

A better definition might be to say it is “fair” to consider Alice as being first if a significant fraction of the community received Alice’s transaction before Bob’s, and that fraction of the community then went on to communicate with most of the others quickly. Under this definition, if Alice and Bob are releasing their transactions to the gossip network at almost the same time, and both spread at about the same rate, then the consensus could go either way, and still be considered fair. However, if Alice gossips her transaction for Bob, beating him by just over the duration of a single gossip sync, then it might be expected that as both transactions spread exponentially, doubling the number of members reached on each sync, in the end over $2/3$ of the population will hear of Alice before Bob, and less than $1/3$ will favor Bob. So in that case, it might be fair to favor Alice but not Bob in the consensus. The hashgraph consensus algorithm is fair in this sense. The members who are online and regularly participating will generate a set of events called *famous witnesses*, and the consensus decision will be that the “first” transaction is whichever transaction reached the majority of that set first. If a small set of members are offline, or are partitioned so that they cannot communicate with the rest, then they will not have famous witnesses, and so having a transaction reach them will not count as having reached the community as a whole. But if the members in that set are communicating with the rest, then they will count as famous witnesses, and they will help decide who reach “the community” first.

There are attacks against this system that would not be considered to be a failure of the consensus system, because they would be equally effective against a single-server solution. For example, the Byzantine proofs assume the attackers control the internet, and can delay arbitrary messages. If attackers actually had that power, they could simply disconnect Alice from the internet for as long as it takes for Bob to send a transaction and have it recorded. This could be done on the real internet by launching a denial of server attack, flooding every computer with packets from Bob in an attempt to prevent Alice from communicating. Of course, this would also be effective if Alice were communicating with a central server, so it could be considered more a failure of the internet than a failure of the consensus system.

Similarly, Bob could gain an advantage over Alice by buying more bandwidth, so that his gossips reach more people. If he has 8 times the bandwidth of Alice, so that he can send his transaction initially to 8 members in the time Alice sends to 1, then he can gain an advantage of the time of about 3 gossip syncs. This is not considered a failure. If his message actually reaches the world before hers, then he should have the credit for it. This is similar to the current stock markets, where companies spend large sums of money for slightly faster connections. So the consensus algorithm would not be considered “unfair” in this case, because it is no worse than a central server.

7. GENERALIZATIONS AND ENHANCEMENTS

7.1. proof-of-stake. So far, it has been assumed that every member is equal. The above algorithms refer to things depending on “more than $2N/3$ of the members” and “more than half of the famous events”. They also use the idea of a “median” of set of numbers. The proof shows Byzantine convergence when more than $2N/3$ of the members are honest.

It is easy to modify the algorithm to allow members to be unequal. Each member can be assumed to have some positive integer associated with them, known as their “stake”. Then, the votes would be replaced with weighted voting, and the medians with weighted medians, where votes are weighted proportional to the voter’s stake. In all of the above definitions, algorithms, and proofs, define “more than $2N/3$ members” to mean “a set of members whose total stake is more than $2N/3$, where N is the total stake of all members”. Similarly “more than half of the famous events” would mean “a set of famous events created by members who total stake is more than half of the total stake of the creators of all the famous events”. The “median of the timestamps of events in S ” would become “the weighted median of the timestamps in S , weighted by the stake of the creator of each event in S ”. The weighted median can be thought of taking each event Y in S , and putting multiple copies of the timestamp of Y into a bag, where the number of copies equals the stake of the member who created Y . Then take the median of the timestamps in the bag.

The Byzantine proof applied as long as the attackers constituted less than $1/3$ of the population. With these new definitions, it will now apply when the attackers together have a stake that is less than $1/3$ of the total stake of all members.

This new proof-of-stake system is more general than the unweighted system. It can still be used to emulate the unweighted system, by simply giving every member a stake of 1. But it can also be used to provide better behavior. For example, the stake might be proportional to the degree to which a member is trusted. Perhaps members who have been investigated in some way should be trusted more than others. Or it could be used to give greater weight to members who have a greater interest in the system as a whole working properly. A cryptocurrency might use each member’s number of coins as their stake, on the grounds that those with more coins have a greater interest in ensuring the system runs smoothly. Or a community could be started by a group of members with mutual trust, each of which is given an equal stake. Then, each existing member could be allowed to invite arbitrarily many new members to join, subject to the constraint that the inviter must split their stake with the invitee. This would discourage a Sybil attack, where one member invites a huge number of sock puppet accounts, in order to control the voting.

the “stake record” is the list of members and the amount of owned by each member. So far, it has been assumed that the stake record is universally known, and is unchanging. It is easy to relax that assumption.

Assume that there is a particular form of transaction that changes the stake record. The community might set up rules at the beginning, governing which such transactions are valid. For example, each member could be allowed to invite other members, up to a total of at most 10 new members. Or perhaps anyone inviting a new member must simultaneously give the new member a portion of their own stake. The validity of such a transaction might depend on the exact order of the transactions in the consensus order. For example, if the rule is that only one new member can be invited, and Alice and Bob both try to invite a new member simultaneously, then whichever invitation comes first in the consensus order will succeed, and the other will fail.

All of this can be accommodated. When the consensus algorithm finishes deciding the question of which round R firsts are famous, at that moment it becomes possible to find exactly which events will have a received round of R , and to calculate

their exact position in the consensus order. At that time, each of the transactions in those events can be processed, and the rules can be consulted to see which are valid, and the valid transactions can be applied. This may change the stake record.

If the stake record does change, then the algorithm should be re-run for all events in round R and later. This may change the calculations of which events are strongly seen, of event round numbers, of which events are firsts, and of which are famous firsts.

Note that when deciding which round R firsts are famous, the calculations are done using the “old” stake record. The voting for round R may continue several rounds into the future, all using the old stake record. Once round R is settled, the future rounds will reshuffle, and the calculations for round $R+1$ firsts will be done using the “new” stake record.

This approach allows all members to be in agreement on exactly what stake record is being used for any given calculation. That ensures that they will always agree on the results of those calculations. And Byzantine agreement will still be guaranteed with probability one.

7.2. signed state. Another enhancement to the system is to have signed states. Once consensus has been reached on whether each witness in round r is famous or not, a total order can be calculated for every event in history with a received round of r or less. It is also guaranteed that all other states (include any that are still unknown) will have a received round greater than r . In other words, at this point history is frozen and immutable for all events up to the start of round r . A member can therefore take all the transactions from those events, and feed them into a database in the consensus order, and calculate the state that is reached after processing those transactions. Every member will calculate the same consensus order, so every member will calculate the same state. This is a *consensus state*. Each member can take the hash of this state and digitally sign it, and put the signature into a new transaction. Soon after, every member will have received by gossip many signatures for the consensus state. Once signatures are collected from more than $2/3$ of the population, that consensus state along with the set of signatures constitutes a *signed state* that is an official consensus state for the system at the start of round r . It can be given to people outside the community, and they can check the signatures, and therefore trust the state. At this point, a member can feel free to delete all the transactions that were used to create the state, and delete all the events that contained those transactions. Only the state itself needs to be kept. It might be possible to do this every few minutes, so there will never be a huge number of transactions and events stored. Only the consensus state itself.

7.3. Efficient gossip. The gossip protocol makes very efficient use of bandwidth. Suppose there are enough transactions being created that every event contains at least one transaction. In any replicated state machine, using a point-to-point network such as the internet, it will be necessary for each member to receive each signed transaction, and to also send each signed transaction on average once. For the hashgraph gossip, the same is true, except that the signature is for the event containing the transaction, rather than on the transaction itself. The only additional overhead is the two hashes and the timestamp, plus the array of counts at the start of the sync. However, the hashes themselves don't have to be sent over

the internet. It is sufficient to merely send the identity of the creator of the event, and its sequence number (for other-hashes).

For example, in Figure <hashgraph>, the upper-right corner is the event of Carol receiving gossip from Bob. During that sync, Bob will send her three events: the central one, the bottom-middle, and the lower-left. After exchanging count arrays, both of them will know that Bob is sending 1 event by Alice and 2 by Bob and 0 by Carol. So, during that sync, they can assign the 3 transferred events sequence numbers, sorting them by creator, and then by order from earliest to latest. If the sync will transfer a total of 256 events or fewer, then the sequence numbers will fit in a single byte. Or 2 bytes for up to 65536 events. If syncing is frequent, it is likely that the members will reach bandwidth limitations before they start sending more than that many events. Therefore, each 256-bit hash can be replaced with a sequence number of just 1 or 2 bytes when Bob sends it, and Carol can expand that back to a full hash in memory after receiving it.

If the average signed transaction is 100 bytes, of which 64 bytes are the 512-bit signature, then the hashgraph protocol adds only 2 to 4 percent to the size of the message being sent. And because voting is virtual, there is no other bandwidth cost at all in order to achieve consensus. In this sense, the bandwidth required for hashgraph consensus is very close to the theoretical limit, which would be the bandwidth needed to simply send the signed and dated transactions themselves.

A system that merely sent the transactions could save bandwidth by not attaching timestamps to the transactions, if the application didn't need timestamps. Hashgraph consensus can do the same. In that case, the "timestamp" within an event would simply be an integer that is its self-parent's "timestamp" plus one.

A system that only sent transactions could also save bandwidth by grouping together several transactions by the same creator, and attaching on a single signature to the list, rather than one per transaction. Hashgraph can do the same, by putting several transactions into a single event, and so having only a single signature for the list.

So the bandwidth requirements of hashgraph consensus are very close to the theoretical minimum in all cases.

7.4. Fast elections. That second part of the algorithm is a Byzantine agreement step for deciding fame. It has an interesting property. When a group of members are all online and all participating regularly, the Byzantine agreement will be applied to a set of elections where almost all the voters start with identical votes. That is because a round R+1 first will strongly see many of the round R firsts, so a round might be expected to last about two "gossip periods", where a gossip period is the time it takes for a message to propagate through the entire community. This should be the time to do $\log_2(N)$ syncs, when there are N members online. For a round R+1 first X to vote YES on the fame of a round R first Y, it isn't necessary for X to strongly see Y. It can merely see Y. It would be expected that Y would propagate to all the online members in a single gossip period. So there is an overwhelmingly high probability it will propagate to them in two gossip periods. In practice, this is unlikely to happen if the members who created X and Y are both online and both participating throughout the round.

Similarly, if Y is a round R first, but was created by a member who was asleep and then awoke just before the end of round R, then it is likely that almost all round R+1 firsts will vote NO, and the election will again end immediately. There

is a small window of time, on the order of the duration of a single sync, in which a member awakening and creating Y can cause the round $R+1$ firsts to start with a 50/50 vote split. If the online members are all choosing each other randomly and syncing frequently, then such a result will converge to a decision in about 3 rounds, with a probability of only a few percent for more than 3 rounds, and of less than a tenth of a percent for more than 6 rounds. If an attacker completely controls the internet, they can cause this to drag on for exponentially many rounds. This can be reduced to a constant expected number of rounds by using a cryptographic “shared coin” protocol, rather than the “middle bit of the signature” described in the above algorithm. The middle bit is intended to be like each member having a random coin flip that the attacker couldn’t predict ahead of time. A shared coin protocol is same, but ensures all members end up with the same “random” result. This addition would reduce the theoretical worst-case expected time. But such an addition seems unlikely to be worth the effort in practice. If an attacker can truly control the internet enough to keep the honest members from syncing randomly with each other for a long period, then during that period, the attacker likely has the power to simply block the honest users from accessing the internet at all. So a shared coin seems to be of only theoretical interest here. But using a shared coin is always an option.

7.5. Efficient calculations. The first part step of the algorithm is to assign a round of either R or $R+1$ to an event, based on whether it can strongly see enough round R events. So it is necessary to calculate whether a round R first event X can be strongly seen by an arbitrary event Y . The following is one way to calculate that answer.

Give each event a sequence number that is one greater than the sequence number of its self-parent. Store an array for Y and an array for X . The Y array remembers the sequence number of the last event by each member that is an ancestor of Y . The array for X remembers the sequence number of the earliest event by each member that is a descendent of X . Compare the two arrays, and find how many elements in the Y array are greater than or equal to the corresponding element of the X array. If there are more than $2N/3$ such matches, then Y strongly sees X .

The array for X is only calculated for a small fraction of all events: those that are round firsts. The array for Y is quickly calculated from the corresponding arrays of its two parents: take the max of corresponding elements of the parents’ arrays, and then set one element of the result to Y ’s sequence number. The comparison of the X and Y arrays can be sped up by multithreading (to use more cores), packing multiple elements into one integer (to use the ALU more efficiently), using assembly language (to access the CPU vector instructions) or using the GPU (for more vector parallelism).

8. CONCLUSIONS

A new system has been presented, based on the Swirls hashgraph data structure, and the Swirls hashgraph consensus algorithm. It is fair, fast, Byzantine fault tolerant, and extremely bandwidth efficient due to virtual voting. The algorithm is given in pseudocode in the figures, using an imperative language, but it is also very natural to describe it in a functional language. The appendix gives the algorithm in functional language, which is concise, and may be of interest.

9. REFERENCES

REFERENCES

- [1] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, OSDI '99, pages 173–186, Berkeley, CA, USA, 1999. USENIX Association.
- [2] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, July 1982.
- [3] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985.
- [4] Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, May 1998.
- [5] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 305–319, Philadelphia, PA, June 2014. USENIX Association.
- [6] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of bft protocols. Cryptology ePrint Archive, Report 2016/199, 2016. <http://eprint.iacr.org/>.
- [7] Allen Clement, Edmund Wong, Lorenzo Alvisi, Mike Dahlin, and Mirco Marchetti. Making byzantine fault tolerant systems tolerate byzantine faults. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'09, pages 153–168, Berkeley, CA, USA, 2009. USENIX Association.
- [8] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. posted to the internet November, 2008, 2008. <http://bitcoin.org/bitcoin.pdf>.
- [9] Giulio Prisco. Intel develops ‘Sawtooth Lake’ distributed ledger technology for the Hyperledger project. *Bitcoin Magazine*, April 2016.
- [10] Dag-Erling Smorgrav. FreeBSD quarterly status report. Posted on FreeBSD.org, 2013. <http://www.freebsd.org/news/status/report-2013-09-devsummit.html#Security>.
- [11] Miguel Miguel Correia, Giuliana Santos Veronese, Nuno Ferreira Neves, and Paulo Verissimo. Byzantine consensus in asynchronous message-passing systems: a survey. *International Journal of Critical Computer-Based Systems*, 2(2):141–161, 2011.

10. APPENDIX A: CONSENSUS ALGORITHM IN FUNCTIONAL FORM

An event is a tuple $e = \{d, h, t, c, s\}$ where:

d	$= data(e)$	$=$ the “payload” data, which may include transactions.
h	$= hashes(e)$	$=$ a list of hashes of the event’s parents, self-parent first.
t	$= time(e)$	$=$ creator’s claimed date and time of the event’s creation.
c	$= creator(e)$	$=$ creator’s ID number.
s	$= sig(e)$	$=$ creator’s digital signature of $\{d,h,t,c\}$.
n	$=$	the number of members in the population
$first$	$=$	the unique event that has no parents
E	$=$	the set of all events
\mathbb{T}	$=$	set of all possible $(time, date)$ pairs
\mathbb{B}	$=$	$\{true, false\}$
\mathbb{N}	$=$	$\{0, 1, 2, \dots\}$
ancestor	$: E \times E \rightarrow B$	
selfAncestor	$: E \times E \rightarrow B$	
see	$: E \times E \rightarrow B$	
stronglySee	$: E \times E \rightarrow B$	
parentRound	$: E \rightarrow \mathbb{N}$	
witness	$: E \rightarrow \mathbb{B}$	
round	$: E \rightarrow \mathbb{N}$	
roundDiff	$: E \times E \rightarrow \mathbb{I}$	
votes	$: E \times E \times \mathbb{B} \rightarrow \mathbb{N}$	
voteFraction	$: E \times E \rightarrow \mathbb{R}$	
vote	$: E \times E \rightarrow \mathbb{B}$	
decide	$: E \times E \rightarrow \mathbb{B}$	
allFamous	$: \mathbb{I} \rightarrow 2^E$	
famous	$: E \rightarrow \mathbb{B}$	
roundReceived	$: E \rightarrow \mathbb{N}$	
timeReceived	$: E \rightarrow \mathbb{T}$	

$$\begin{aligned}
\text{ancestor}(x, y) &= (x = y) \vee (\exists z \in \text{parents}(x) : \text{ancestor}(z, y)) \\
\text{selfAncestor}(x, y) &= \text{ancestor}(x, y) \wedge ((\text{selfParent}(x) = y) \vee \text{selfAncestor}(\text{selfParent}(x), y)) \\
\text{see}(x, y) &= \text{ancestor}(x, y) \wedge \neg(\exists a, b, c \in E : \\
&\quad (\text{ancestor}(y, a) \wedge \text{ancestor}(y, b) \wedge c \in \text{parents}(x) \wedge c \in \text{parents}(b))) \wedge \\
&\quad \text{creator}(a) = \text{creator}(b) = \text{creator}(c)) \\
\text{stronglySee}(x, y) &= \text{see}(x, y) \wedge (\exists S \in 2^E : (|S| > 2n/3) \wedge (z \in S \iff (\text{see}(x, z) \wedge \text{see}(z, y)))) \\
\text{parentRound}(x) &= \begin{cases} 0 & \text{if } x = \textit{first} \\ \max_{y \in \text{parents}(x)} \text{round}(y) & \text{otherwise} \end{cases} \\
\text{witness}(x) &= \exists S \in 2^E : (|S| > 2n/3 \wedge \\
&\quad (\forall y \in S : (\text{round}(y) = \text{parentRound}(x) \wedge \text{stronglySee}(x, y)))) \\
\text{round}(x) &= \begin{cases} 1 + \text{parentRound}(x) & \text{if witness}(x) \\ \text{parentRound}(x) & \text{otherwise} \end{cases} \\
\text{roundDiff}(x, y) &= \text{round}(x) - \text{round}(y) \\
\text{votes}(x, y, v) &= |\{z \in E \mid \text{see}(x, y) \wedge \text{roundDiff}(x, z) = 1 \wedge \\
&\quad \text{stronglySee}(x, z) \wedge \text{vote}(z, y) = v\}| \\
\text{voteFraction}(x, y) &= \text{votes}(x, \textit{true}) / (\text{votes}(x, \textit{true}) + \text{votes}(x, \textit{false})) \\
\text{vote}(x, y) &= \begin{cases} \text{see}(x, y) & \text{if } \text{roundDiff}(x, y) = 1 \\ (\text{voteFraction}(x, y) \geq 1/2) & \text{if } (\text{roundDiff}(x, y) \bmod 5 \neq 1) \vee \\ & |\text{voteFraction}(x, y) - 1/2| > 1/6 \\ (1 = \text{middleBit}(\text{signature}(x))) & \text{otherwise} \end{cases} \\
\text{decide}(x, y) &= \text{vote}(x, y) \wedge (\text{roundDiff}(x, y) \bmod 5 \neq 1) \wedge (\text{voteFraction}(x, y) > 2/3) \\
\text{allFamous}(r) &= \{x \in E \mid \text{famous}(x) \wedge \text{round}(x) = r\} \\
\text{famous}(x) &= \text{witness}(x) \wedge \exists y \in E : \text{decide}(y, x) \\
\text{roundReceived}(x) &= \min_{r \in \mathbb{N}} (|\{y \in E \mid \text{round}(y) = r \wedge \text{famous}(y) \wedge \text{see}(y, x)\}| / \\
&\quad |\{y \in E \mid \text{round}(y) = r \wedge \text{famous}(y)\}| \geq 1/2) \\
\text{timeReceived}(x) &= \text{median}(\{\text{time}(y) \mid y \in E \wedge \text{see}(y, x) \wedge \\
&\quad (\exists z \in E : \text{round}(z) = \text{roundReceived}(x) \wedge \text{selfAncestor}(z, y)) \wedge \\
&\quad \neg(\exists w \in E : \text{selfAncestor}(y, w) \wedge \text{see}(w, x))\})
\end{aligned}$$

E-mail address: baired@swirllds.com