



HAL
open science

Brief Announcement: Breaking the $f + 1$ Barrier: Executing Payment Transactions in Parallel with Less than $f + 1$ Validations

Rida Bazzi, Sara Tucci

► **To cite this version:**

Rida Bazzi, Sara Tucci. Brief Announcement: Breaking the $f + 1$ Barrier: Executing Payment Transactions in Parallel with Less than $f + 1$ Validations. PODC '23: Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing, Jun 2023, Orlando, United States. pp.274-277, 10.1145/3583668.3594580 . cea-04488894

HAL Id: cea-04488894

<https://cea.hal.science/cea-04488894>

Submitted on 4 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Breaking the $f + 1$ Barrier: Executing Payment Transactions in Parallel with Less than $f + 1$ Validations

RIDA BAZZI, Arizona State University, USA

SARA TUCCI-PIERGIOVANNI, Université Paris-Saclay, CEA, List, USA

We consider the problem of supporting payment transactions in an asynchronous system in which up to f validators are subject to Byzantine failures under the control of an adaptive adversary. It was shown that this problem can be solved without consensus by using byzantine quorum systems (requiring at least $2f + 1$ validations per transaction in asynchronous systems). We show that it is possible to validate transactions in parallel with less than f validations per transaction if each transaction spends no more than a small fraction of a balance. Our solution relies on a novel quorum system that we introduce in this paper and that we call (k_1, k_2) -quorum systems. In the presence of a non-adaptive adversary, these systems can be used to allow up to k_1 transactions to be validated concurrently and asynchronously but prevent more than k_2 transactions from being validated. If the adversary is adaptive, these systems can be used to allow k_1 transaction to be validated and prevent more than $k'_2 > k_2$ transactions from being validated, the difference $k'_2 - k_2$ being dependent on the quorum system's *validation slack*, which we define in this paper. Using (k_1, k_2) -quorum systems, a payer can execute multiple partial spending transactions to spend a portion of its initial balance with less than full quorum validation (less than f validations per transaction) then reclaim any remaining funds using one fully validated transaction, which we call a *settlement* transaction.

CCS Concepts: • **Theory of computation** → Distributed algorithms; • **Computer systems organization** → Reliability; • **Security and privacy** → Distributed systems security.

Additional Key Words and Phrases: Distributed Systems, Blockchain, Quorums, Fault Tolerance

1 INTRODUCTION

Existing cryptocurrencies solve the consensus problem to maintain a shared ledger of all transactions, but it was shown that maintaining a shared ledger is not always strictly needed to support exchanging funds [11]. In an asynchronous permissioned system in which at most one third of the servers are subject to Byzantine failures, Byzantine quorums can be used to allow different parties to exchange funds through the system – the asset transfer task. The fact that solving consensus is not needed for asset transfer, when the asset has a single owner, was rediscovered by others who formalized the problem, gave it the name *asset transfer task* and generalized it to multi-owner objects [10] as well as generalized the approach to work in a permissionless proof-of-stake system [18]. The key insight of [11] is that (1) Byzantine quorums can prevent double spending because any two quorums must have at least one correct server in their intersection and (2) ordering unrelated transactions is not needed to exchange funds. Consensus-less solutions are important theoretically, but also in practice for their ability to achieve higher throughput and to reduced transaction latency [9].

Despite the increased efficiency of the consensus-less approach, transaction processing is still fundamentally sequential; If one has a balance of ten coins and wants to pay two coins of the ten coins separately to two different recipients, both transactions would be required to be processed by a common correct server to avoid double spending. This, in turn, necessitates every request to be processed by a full quorum so that any two quorums have at least $f + 1$ servers in common, where f is an upper-bound on the number of faulty servers. The $f + 1$ -intersection requirement ensures *safety* so that no two conflicting payments can be simultaneously validated.

The work on this paper started by considering the possibility of relaxing the intersection requirement so that payments from the same single-owner account could be made in parallel while avoiding double spending. In the example above, the two single-coin payments are not conflicting, so we would like to allow them to go through without interfering with each other. At the same time, if there were more than ten such single-coin payments, we would like to prevent some of them from going through. These two requirements seem to be conflicting. The question that this paper addresses is the following: *In an asynchronous system, can we validate non-conflicting asset-transfer transactions in parallel with less than a full quorum of validators while still being able to prevent double spending?*

The surprising answer to this question is in the affirmative, albeit with a some caveats. The main, and provably unavoidable, caveat is that the total balance amount cannot be fully spent with such non-interfering parallel transactions: the owner can make payments in parallel up to a given threshold. In order to spend the remainder of the balance, the owner needs to use a larger quorum for validation. In particular, the owner can issue a *settlement* transaction to pay the unspent balance to self. The other main caveat is that the solution is probabilistic. Again, this is provably unavoidable. Finally, our current solution requires $n > 8f$ and it is not yet clear how to improve this requirement. An important insight of the solution is captured in the coin example above. If each transaction only spends a fraction of the total balance, then we are able to execute multiple such *partial spendings* in parallel, but the total number of such transactions should allow for a buffer between the amount spent and the total balance so that the balance is not exceeded. Even more surprising, it turns out that amounts received through partial spending (without full quorum validation) can also be spent partially without full quorum validation if the original payer is correct. If the original payer is not correct, then payees might not be able to spend the amounts received before settlement, but at no time is double spending possible even if all clients are faulty.

Answering the question in the case of a single owner lead us also to reconsider the *proven* need for solving consensus in order to support transfers from accounts with multiple owners [10]. Intuitively, the need for consensus is to handle double spending from the same account by two different owners. It turns out that partial spending from a multi-owner account is not fundamentally different from the single owner case. Our solution allows multiple owners to spend parts of the balance independently in parallel up to a threshold. The settlement procedure for multi-owner accounts would require solving consensus, which is to be expected given the impossibility result. Our protocols only treat the single owner case.

Our solution is based on a new quorum system that we introduce in this paper and that we call (k_1, k_2) -quorum system. Unlike traditional quorums in which the main requirement is an intersection requirement, (k_1, k_2) -quorum systems have both a *non-intersection* as well as an intersection requirements. The non-intersection requirement is an upper bound on the size of the overlap between a newly selected quorum and up to $k_1 - 1$ previously selected quorums. The intersection requirement is a lower bound on the size of overlap between a newly selected quorum and a set of k_2 or more previously selected quorums (the actual definition is more subtle and should account for previously corrupted validators).

The solution for executing payment transactions would allow a client using a (k_1, k_2) -quorum system to execute with high probability k_1 partial spending transactions, each of which spends at most $1/k'_2$ of the client's balance (for a well specified $k'_2 > k_2$), so the total guaranteed spending is k_1/k'_2 of the clients balance. It is possible that the client can succeed in executing more than k_1 partial spending transactions but that is not guaranteed. Finally, with high probability the client cannot successfully execute more than k'_2 partial transactions so the total amount cannot exceed the client's balance. This is achieved even in the presence of corrupt clients that can collude with the servers that validate transactions (the validators).

The combination of intersection and non-intersection requirements of (k_1, k_2) -quorum systems is novel and can potentially be applicable to other settings in which we need to allow a limited amount of concurrent activities. In particular, we expect that these systems could be applied to some classes of smart contracts in which available funds significantly exceed actual spending which is done in small increments. To summarize the main contributions of this paper are the following:

- (1) We introduce (k_1, k_2) -quorum systems, a new quorum system that has both intersection and non-intersection requirements. The non-intersection requirements allow multiple operations in parallel and the intersection requirements limit the number of operations that can go through without conflict.
- (2) We introduce the *partial spending problem* which formalizes the requirements on partial spending transactions and their corresponding settlements.
- (3) We present the first protocol that allows multiple payment transactions from the same account to be executed in parallel.

The rest of this paper is organized as follows. Section 2 discusses related work. Section 3 presents the system model. Section 5 introduces (k_1, k_2) -Byzantine quorum systems. Section 7 presents our solutions. Correctness proofs are presented in Section ???. Section 8 concludes the paper and discusses how the solution can be simplified for synchronous systems.

2 RELATED WORK

The fundamental bottleneck of blockchains is the underlying consensus protocol used to add blocks to the replicated data structure. To improve blockchain scalability in the context of cryptocurrencies, two approaches emerged: asynchronous on-chain solutions that attempt to create consensus-less blockchain protocols [10, 11, 18] and off-chain solutions, as channels and channels factories (e.g. [3, 7, 15, 16]), that move the transaction load offline while resorting to a consensus-based blockchain only for trust establishment and dispute resolution.

Asynchronous on-chain solutions rely on the assumption that at most f out of $n > 3f$ servers are Byzantine. In such solutions, payments from the same buyer to different sellers cannot be parallelized and they require at least $f + 1$ validations in common (so, at least $2f + 1$ validation per transaction). As for off-chain solutions, a buyer can set more than one channel to parallelize payments to different sellers, however, all the parties need to agree a priori on the set of participants and cooperate to operate the channel. Cooperation is needed to open the channel (by locking funds as initial balances), update the state of the channel (by signing transactions that attest of the new allocation of balances) and close the channel by sending the last state update to the blockchain and unlock funds. Frauds are avoided by constantly monitoring the state of the blockchain, in the case the other party tries to close the channel with a state update different from the last one. Interestingly, to limit cooperation failures and frauds, [3] uses a set of n processes, called wardens, to proactively validate state updates associated with increasing timestamps agreed by both parties. To get a validation, the buyer needs to contact a full quorum of t of wardens, with $t = 2f + 1$, under the assumption of at most f out of the $n = 3f + 1$ wardens are Byzantine and the non-Byzantine wardens are rational. In our setting, we allow a buyer to pay potential sellers in parallel, without prior agreement with the sellers and we can support partial spending with validations from less than f validators, all of which can be Byzantine in the case of a corrupt seller.

The closest relative of the (k_1, k_2) -quorum systems that we propose are k -quorum systems [1, 2]. Traditional k -quorum systems relax the intersection requirement of quorum systems. A read-quorum is not required to intersect every write-quorum, but they are required to intersect at least one quorum out of any sequence of k successively accessed write-quorums. Successive write-quorums are not required to have a non-empty intersection, but, unlike (k_1, k_2) -quorum systems,

there are no requirement that prevents the intersection of two write quorums from being large in size. In systems with Byzantine failures, quorums have more than f elements, and the intersection of a read quorum with the k previous write quorums should be large enough to ensure that some correct servers are in the intersection [2].

Other quorum systems that relate to our proposed quorum system are probabilistic quorum systems [12].

3 SYSTEM MODEL

We consider an asynchronous message passing system of n servers and an unbounded number of clients. Servers act as *validators* for client transactions. Up to f servers and any number of clients can be subject to Byzantine failures under the control of an adaptive adversary. The adversary can corrupt any server up to the limit f and any number of clients, but the adversary does not have access to the local memory of participants not under its control. If the adversary corrupts a client or server, then the adversary has full control of the corrupted party including the contents of its memory before it is corrupted. We refer to the parties under the adversary's control as corrupt or faulty and to the parties not under the adversary's control as honest or correct. Corrupt parties can deviate arbitrarily from their protocols. While communication is asynchronous, we make the assumption that messages cannot be selectively delayed to a large enough set of correct servers unknown to the adversary. In other words, if S is a large enough set of correct servers selected by the client's protocol but unknown to the adversary, then any client that sent a request to all servers and is waiting for replies from $n - f$ or $n - 2f$ servers, will receive with high probability a reply from an element of S . The justification for this is that the adversary can only guess the identities of a few elements of S and selectively delay their communication but cannot do so for all elements of S . This is consistent with the goals of our solution whose guarantee should hold with high probability and not in the worst case.

We assume that clients and servers use public-key signature and encryption schemes to sign and encrypt messages and that they are identified by their public keys [8, 17]. The public keys of servers are assumed to be known to clients. We assume, but do not explicitly show in the protocols, that messages are signed and signatures are verified. The adversary is computationally bounded and cannot break the encryption or signature schemes. In particular, the adversary cannot read encrypted messages between parties not under its control. Finally, parties have access to a Hash function in the random oracle model [5].

The adversary can monitor communication between clients and determine which clients are communicating together. We assume that a client can batch messages for multiple transactions together, so that validators for different transactions are contacted together and the adversary cannot tell, just by observing the communication, which of the contacted servers validate which transactions. This point is discussed further when we present the protocols.

4 THE PARTIAL SPENDING PROBLEM

4.1 Payments and Settlements

Transactions have two clients, the payer and the payee. We refer to them as the *buyer* and the *seller*, respectively. We use pkb to denote a buyer b and pks to denote a seller s . Spending is done from funds. A fund F is identified by a unique fund identifier $F.id$. It has a balance $F.bl$, where bl is a non-negative amount of money, and one or more owners $F.owners$ associated with it (in the protocols we only consider single owner funds). Funds are *certified* by validators. Each fund has a certificate $F.certificate$ that consists of a set of validations where each validation has the form $(\langle F \rangle, \sigma, pkv)$ such that $\sigma = \text{Sign}_{skv}(\langle F \rangle)$, where $\langle F \rangle$ is an encoding of the fund information (id,

balance and owners) and pkv and skv are the public and private keys of the validator v that signed the validation.

Depending on the size of $F.certificate$, we distinguish between fully certified and partially certified funds. A fund is fully certified if the certificate has validations from at least $f + 1$ validators. A fund is partially certified otherwise. Fully certified funds can be initially fully certified, through external means, or are the result of settlement transactions.

To keep the presentation simple, we only consider partial spending transactions from fully certified funds in the formal problem definition and solution and we don't consider partial spending from partially certified funds. In the appendix, we explain how spending from partially certified funds can be supported.

For fully certified funds, our protocol allows an honest client (using (k_1, k_2) -quorums), to execute at least k_1 partial spending transactions and at most k_2 such transactions, for some global constants k_1 and k_2 that apply to all partial spending transactions from fully certified funds. In order to avoid double spending by corrupt clients under the control of an adaptive adversary, the amount spent per transaction is $F.bl/k'_2$, for some $k'_2 \geq k_2$. The value of k'_2 is fully defined as a function of k_2 , n and f .

A *partial spending transaction* $tx : (F, pkb, pks, PAY)$ is a transfer of money from fund F with $pkb \in F.owners$ to a recipient pks . The fund resulting from a transaction (F, pkb, pks, PAY) from fully certified fund F is a partially certified fund F' such that $F'.id$ is uniquely determined by $F.id$, pks , pkb and a payment number N_s which is uniquely generated by the seller for each payment transaction between pks and pkb to allow for multiple payments from a buyer to the same seller. The balance of F' is a fraction of the balance of F : $F'.bl = F.bl/k'_2$, for a global constant k'_2 , and $F'.owners = pks$. The value $1/k'_2$ is the *partial spending fraction*. To keep the presentation simple, we don't support the aggregation of partial spending transactions from separate funds. We note that aggregation can be supported by separate partial spending transactions from separate funds (which can be optimized if the separate funds have the same owners). Also, we note here that the model is different from the UTXO model [14] in which no balance remains in the inputs used for payment. In our model, partial spending transactions leave a balance in the fund from which the payments are made, but the balance is not explicitly maintained. In addition to payment transactions, there are *settlement transactions*. A settlement transaction $(F, SETTLE)$ specifies a fund F to be settled. The fund being settled can be a fully validated or a partially validated fund, but the fund resulting from the execution of a settlement transaction is fully validated. The identifier of the fund F' resulting from executing $(F, SETTLE)$ is uniquely determined by the identifier of F . A validated settlement transaction from fund F results in a fund whose owners are the same as those of F and whose balance is equal to the unspent amount in F and is specified in the safety requirements below.

4.2 Problem Definition

The formal problem definition below only considers requirements for partial spending from initially certified funds (level 1 payments) and settlements associated with those payments. The problem definition does not consider spending from funds resulting from level 1 payments (level 2 payments) or higher level payments that we referred to in the introduction. The problem requirements ensure that no double spending is possible. In particular, funds resulting from partial spending transactions to honest sellers can always be settled successfully. The same is not true for funds resulting from partial spending transactions to corrupt sellers. In all cases, the total balances of all the funds resulting from settling funds resulting from partial spending transactions from F does not exceed the balance of F . Finally, we need for honest owners the settlement amount is at least the initial balance minus all payments from the fund (requirement 5); it is not guaranteed to be equal because payment to corrupt sellers could be erased by the adversary. These requirements

are straightforward to capture. Before presenting the requirements, we introduce some notation first. For a given fund F , we denote by $funds_F$ the set of funds resulting from payment transactions of the form (F, pkb, pks, PAY) and we denote by $funds_F^H$ the set of such funds for which pks is honest. We denote by $settled_F$ the set of fully certified funds resulting from transactions of the form $(F_{pay}, SETTLE)$, where $F_{pay} \in funds_F$; i.e. the set of settlement transactions of funds resulting from spending money from F .

The following holds with high probability (we use *with high probability*, abbreviated *w.h.p.*, to denote a probability of the form $1 - \text{negl}(n)$, where $\text{negl}(n)$ is a negligible function in the relevant parameters):

- (1) (**progress**) All partially certified funds with honest owners can be settled successfully: Let F' be a partially validated fund resulting from transaction (F, pkb, pks, PAY) where F is a fully certified fund and pks is honest. If pks executes transaction $(F', SETTLE)$, the transaction will terminate resulting a fully certified fund.
- (2) (**safety**) If F'' is a fully certified fund resulting from executing $(F', SETTLE)$ for partially validated fund F' , then $F''.bl = F'.bl$.
- (3) (**safety**) There can be at most k'_2 partial spending transactions from a given fund F for a total spending not exceeding $F.bl$.
- (4) (**safety**) Settlement amounts for payments from F are subtracted from settlement for F : If executing transaction $(F, SETTLE)$ results in a fully certified fund F_R : $F_R.bl \leq F.bl - \sum_{F' \in settled_F} F'.bl$
- (5) (**safety**) Payments to honest sellers are subtracted from settlement amount: If executing transaction $(F, SETTLE)$ results in a fully certified fund F_R : $F_R.bl \leq F.bl - \sum_{F' \in funds_F^H} F'.bl$
- (6) (**safety**) If the owner of a fund F is honest, settlement amount is no less than the the initial balance of F minus payments made from F : If executing transaction $(F, SETTLE)$ results in a fully certified fund F_R and $F.owners$ is honest: $F_R.bl \geq F.bl - \sum_{F' \in funds_F} F'.bl$
- (7) (**progress**) Non-interference: If a total of $k \leq k_1$ payment transactions are initiated by $pkb \in F.owners$ from fully certified fund F and no additional payment or settlement transactions are initiated by $F.owners$ and pkb is honest, then, every one of the k transactions whose seller (payee) is honest will be validated.
- (8) (**progress**) Successful settlement for fully certified funds: If the owner of fully certified fund F is honest and executes an $(F, SETTLE)$ transaction, the settlement transaction for F will terminate.

Note that for the non-interference requirement, the guarantee of termination for each transaction holds even if the messages for the remaining transactions are arbitrarily delayed by slow sellers for example. In other words, executing one transaction does not interfere with the completion of other transactions if the total number of transactions does not exceed the threshold k_1 .

5 (k_1, k_2) -BYZANTINE QUORUMS

As we explained in the introduction, (k_1, k_2) -Byzantine quorum systems are what enables our solution. We first introduce the properties of our new quorum system, then we propose a construction of a particular quorum system that satisfies these properties. In this section, lemmas and theorems are given without proof. Proofs are given in the Appendix.

5.1 (k_1, k_2) -quorums: intersection and non-intersection properties

Our proposed (k_1, k_2) -quorum systems have both lower bounds and upper bounds on the sizes and composition of quorum intersections. We discuss their requirements informally before giving the formal definitions. In our definitions, we assume, and our protocols ensure, that quorums are




Fig. 1. Illustration of (k_1, k_2) -quorums intersection requirements: (a) the intersection of up to k_1 previously selected quorums with Q together with previously corrupted servers in Q does not exceed $\alpha|Q|$; (b) the intersection of k_2 (or more) previously selected quorums with Q minus any previously corrupted quorums in the intersection is at least $\beta|Q|$;

selected according to the uniform access strategy in which all quorum sets are equally likely to be selected. We do not distinguish between read-quorums and write-quorums. For the upper and lower bounds on the sizes of intersections, we only consider previously corrupted servers, as opposed to servers that are corrupted in response to quorum selection. Previously corrupted servers are chosen randomly according to a distribution that is induced by the probabilistic access strategy. Even though we only consider previously corrupted servers in our definition of (k_1, k_2) -quorum systems, in our protocols that use these systems, we also consider corruption by an adaptive adversary that might decide on what servers to corrupt based on quorum selection and other information it might have. The k_1, k_2 parameters are such that the following requirements are satisfied:

- (*Non-intersection requirement*): For any newly selected quorum Q and any $\ell \leq k_1$ previously accessed quorums, we require that, with high probability, the size of the union of the previously corrupted servers in Q together with the intersection of Q with the ℓ previously selected quorums does not exceed a fraction α of the size of Q .
- (*Intersection requirement*) For any newly selected quorum Q and any $\ell \geq k_2$ previously accessed quorums, we require that, with high probability, the intersection of Q with the union of the ℓ previously selected quorums minus the previously corrupted servers in Q has a size that exceeds a fraction β of the size of Q . In other words, with high probability, the number of correct servers in Q that have been contacted in the previous ℓ accesses is more than a fraction β of size of Q .

The non-intersection and intersection requirements are used to ensure that up to k_1 accesses create no conflicts and k_2 or more accesses create a conflict. The idea is to allow a transaction to go through if less than α fraction of the quorum have handled potentially conflicting transactions and to deny it if more than β fraction of the quorum has handled conflicting transactions. The definition takes into consideration previously corrupted servers who can attempt to coordinate their replies to cause the most damage either to prevent a valid transaction from being validated by claiming to have handled conflicting transactions or to allow an invalid transaction to be validated by validating conflicting transactions. The requirements do not say anything about what holds if $k_1 < \ell < k_2$, which is not an issue for the protocols using these systems. The reason for using two parameters α and β instead of only α has to do with the adaptive adversary who can corrupt a randomly chosen quorum after it is chosen. By corrupting selected servers, the adversary can flip a decision from denying a request to accepting a request. As we will see below, the difference

between β and α determines how costly it is for the adversary to flip a decision and places an upper bound on the number of such *flipped* decisions that the adaptive adversary can achieve. In what follows we give the formal definition of (k_1, k_2) -quorum systems.

Definition 5.1. A $(k_1, k_2, \alpha, \beta, \epsilon, \delta)$ -quorum system is a collection Q of sets such that:

- (1) Lower bound: $\Pr_{Q, Q_j \leftarrow Q; j \in J_1; |J_1| \leq k_1} [|Q \cap (\mathcal{F}_{pr} \cup \bigcup_{j \in J_1} Q_j) | > \alpha |Q|] \geq 1 - \epsilon$
- (2) Upper bound: $\Pr_{Q, Q_j \leftarrow Q; j \in J_2; |J_2| \geq k_2} [|(Q \cap (\bigcup_{j \in J_2} Q_j)) - \mathcal{F}_{pr}| \leq \beta |Q|] \geq 1 - \delta$

where \mathcal{F}_{pr} is the set of faulty servers that existed prior to the random choice of Q .

The lower and upper bound requirements in the definition correspond to the non-intersection and intersection requirements stated above and can be best understood by referring to Figure 1.

The definition contains many parameters. The ϵ and δ parameters are essentially security parameters that need to be small enough (negligible, in the formal security sense) for the given application. The α and β parameters are used for separating the lower and upper bounds. The main parameters from a client point of view are the k_1 and k_2 parameters, assuming ϵ and δ are negligible, hence we refer to these systems as (k_1, k_2) -quorum systems.

In (k_1, k_2) -quorum systems, individual quorums can be of size less than f , which means that even if the randomly chosen quorum does not have any faulty servers, it can be completely corrupted by the adaptive adversary if the adversary knows which servers are in a quorum. In our protocols, this can happen if the seller is corrupt. The definition of the system guarantees that if no more than k_1 accesses are made, less than α fraction of servers in a quorum will deny a request. If more than k_2 access are made, then more than β fraction of the servers (none of them previously corrupted) will deny the request. In order for the adaptive adversary to flip the decision, it will need to corrupt $(\beta - \alpha)|Q|$ servers. In fact, by definition, the $\beta|Q|$ servers consisting of previously non-corrupt servers will deny the request with high probability. This number will need to be reduced to no more than $\alpha|Q|$ servers and that would require corrupting $(\beta - \alpha)|Q|$ servers, so that only $\alpha|Q|$ remains. This is the motivation for defining the *validation slack*.

Definition 5.2. The *validation slack* of a $(k_1, k_2, \alpha, \beta, \epsilon, \delta)$ -quorum system is equal to $(\beta - \alpha)|Q_{min}|$, where Q_{min} is the minimum quorum size

In total, in a payment protocol using a $(k_1, k_2, \alpha, \beta, \epsilon, \delta)$ quorum system, it is possible for at most $k_2 + f/((\beta - \alpha)|Q_{min}|)$ accesses to be made: k_2 accesses could go through without the intervention of the adversary but the additional $f/((\beta - \alpha)|Q_{min}|)$ accesses can only be achieved by adaptively corrupting validators when clients are corrupt.

5.2 (k_1, k_2) -quorums Construction

In this section, we propose a simple construction of a (k_1, k_2) -quorum system to show that they can be constructed. Our construction is not optimal and a more careful choice of the system parameters could yield better performance. Exploring the design space for (k_1, k_2) -quorum systems is beyond the scope of this work.

Definition 5.3. An (m, n) uniform balanced (k_1, k_2) -quorum system is a quorum systems in which each quorum has size m , $|U| = n = (k_1 + k_2)m$ and quorums are selected uniformly at random.

LEMMA 5.4. An (m, n) uniform balanced quorum system is a (k_1, k_2) -quorum system with $\alpha = 1/3$ and $\beta = 2/3$ has $\epsilon = \delta = \text{negl}(m)$ and validation slack = $m/3$, if $p_f + \alpha_1 < 1/3$, where $\alpha_1 = k_1 m/n$ and $p_f = f/n$.

5.3 Asynchronous (k_1, k_2) -quorum systems

The calculations and definition of (k_1, k_2) -quorum systems implicitly assume that all servers in a quorum reply to a request from a client. In order to access a quorum in an asynchronous system, the access should allow for non-responding corrupt servers that cannot be timed out. Assuming that the corrupt servers in a given quorum are randomly chosen, the expected number of corrupt servers in a quorum of size m is at most mp_f , where $p_f = f/n$ is an upper bound on the probability that a randomly selected server has been previously corrupted. With high probability, the number of corrupt servers in a quorum is less than $(1 + \mu)mp_f$ for any constant μ , $0 < \mu < 1$, and large enough m . In other words, the probability that the number of corrupt servers in a quorum exceeds $(1 + \mu)mp_f$ is a negligible function of m . It follows that, when contacting a random quorum unknown to the adversary, an honest client can wait for replies from $m(1 - (1 + \mu)mp_f)$ servers and be guaranteed *w.h.p.* to receive that many replies. Of those replies, $m(1 - 2(1 + \mu)mp_f)$ are guaranteed *w.h.p.* to be from non-corrupt servers.

These considerations affect our definition and calculations for (k_1, k_2) -quorums in asynchronous systems. The intersection and non-intersection properties should allow for the exclusion of $(1 + \mu)mp_f$ servers, which could be correct. We revise the definition as follows.

Definition 5.5. A $(k_1, k_2, \alpha, \beta, \epsilon, \delta, \mu)$ -asynchronous quorum system is a collection of Q of sets such that $\forall Q_s : |Q_s| \leq (1 + \mu)p_fm$, the following holds

- (1) Lower bound: $\Pr_{Q, Q_j \leftarrow Q; j \in J_1; |J_1| \leq k_1} [|(Q - Q_s) \cap (\mathcal{F}_{pr} \cup \bigcup_{j \in J_1} Q_j)| > \alpha|Q|] \geq 1 - \epsilon$
- (2) Upper bound: $\Pr_{Q, Q_j \leftarrow Q; j \in J_2; |J_2| \geq k_2} [|(Q - Q_s) \cap (\bigcup_{j \in J_2} Q_j) - \mathcal{F}_{pr}| \leq \beta|Q|] \geq 1 - \delta$

where \mathcal{F}_{pr} is the set of faulty servers that existed prior to the random choice of Q .

This definition requires the intersection and non-intersection properties to hold even if we exclude any set Q_s of size up to $(1 + \mu)mp_f$ servers from the intersection. Again, the definition has too many parameters. We want μ and α and β to be specific constants for which ϵ and δ are small enough. We prove in the appendix the following lemma.

LEMMA 5.6. *An (m, n) uniform balanced (k_1, k_2) -quorum system is a $(k_1, k_2, \epsilon, \delta, \alpha = 1/3, \beta = 2/3, \mu > 0)$ -asynchronous quorum system with $\epsilon = \delta = \text{negl}(m)$ and validation slack $= m/3$, if $n > 8f$ and $k_1 m/n < 1/24$.*

6 SOLUTION OVERVIEW

We assume that we have a buyer with a fully certified fund to be spent with partial spending transactions. To understand the difficulties in coming up with a working solution, we start with a solution that does not work. Then, we successively show how the solution should be modified until we get a working solution. As a first attempt, to make a payment, we can require the buyer to get the transaction validated by a (k_1, k_2) -quorum of validators and present the validation to the seller. The idea is that, given the properties of (k_1, k_2) -quorum systems, the buyer should not be able to validate more than k_2 partial spending transactions and thus double spending is avoided. It turns out that buyer validation is fundamentally flawed due to concerns about settlement and quorum choice. We consider that next.

Buyer shouldn't learn the identities of validators: The buyer can always contact the same corrupt quorum for validation which would allow for the validation of more than k_2 transactions. To prevent the buyer from choosing the same quorum, we can involve the seller in the choice of the quorum, a choice that should be randomized using the output of the random oracle (hash function) seeded with a combination of nonces provided by the seller and the buyer. This would ensure that the choice is random and, if the seller is honest, the buyer has little control over the identities of the validators. Choosing the quorum randomly prevents the buyer from validating more than k_2

partial spending transactions, but, if the buyer is corrupt, the adaptive adversary would know the identities of the validators and would be able to corrupt them after they are chosen. While the number of transactions that can be validated would still be bounded even if the adversary can corrupt servers in chosen quorums (discussed below), the adversary can create problems when transactions are settled. In fact, consider a fully certified fund F from which a payment is made to an honest seller resulting in partially validated fund F' . If F is settled before F' and the validators of F' are corrupted by the adaptive adversary, they can deny knowledge of F' and F' will not be counted against F' 's balance. This would either result in double spending when F' is settled or in denying the settlement of F' in violation of the problem requirements. This scenario is possible because the buyer knows the identities of the validators and the adversary is adaptive. It follows that the solution should avoid giving the buyer knowledge of the identities of the validators. In our solution, the seller, without involvement from the buyer, chooses the quorum of validators. This works because the buyer's and the seller's interests in getting proper validation are at odds. An honest seller has interest in getting as many validators as possible for a given transaction because that would ensure that there is a record of the transaction that they can claim at settlement time. A corrupt buyer has the opposite interest in that it would want to erase any record of a transaction in the hope that the transaction would not count against its balance at settlement time. Of course, the quorum should be randomly chosen and dependent on input from the seller in addition to the transaction identifier. Still this does not solve all the issues. We discuss this next.

Protecting the seller during settlement: In the previous section, we discussed how the adversary should not learn the identities of the validators of an honest seller. During settlement, the seller needs to fully validate the settlement transaction and therefore needs to prove to new validators that were not involved in validating the partially validated fund being settled that the fund is properly validated. For that the seller needs to divulge the identities of the validators that validated its partially validated fund. The seller cannot send the information to all validators because it is possible for one corrupt validator to receive the information before everyone else. At that point, the adversary learns the identities of the validators and corrupt them to erase their record of the transaction then the buyer can settle its own fund before the seller's settlement request is received by non-corrupt validators, thereby repeating the scenario above. The same scenario can occur during the settlement of a buyer's fund F when validators ncommunicate with each other to determine if there are any partial spending transactions from F that need to be subtracted from the balance owed to the buyer. If the adversary delays the validators' communication except for a pre-chosen corrupt validator, the adversary can learn the identities of all validators that validated partial spendings from F and could then corrupt them to erase the record of some of those partial spendings. To prevent this from happening, divulging the identities of validators is done using secret sharing [4]. The information is divulged in two stages. In a first stage the holder of the information (the *dealer*) shares the secret information and in the second stage, the secret information is *reconstructed*. This way, when the adversary learns the information, it is too late to erase it. In our setting, we only require that information provided by honest validators can be recovered.

Buyer should limit number of validators: Since a validator does not validate more than one payment from a given fund, the seller should not get validations from many validators because that can affect the ability of the buyer to get other partial spending transactions validated. The solution to this is to have the buyer approve the validators, but that needs to be done without the buyer learning the identities of the validators. This can be achieved by having the buyer *blindly* sign a limited number of *validation approvals* such that one approval cannot be used with two different validators. This way, the buyer cannot contact more than a fixed number of validators for a given transaction and the buyer does not learn the identities of the validators.

The partial spending amount should account for an adaptive adversary and a corrupt seller: If the seller is corrupt, then the identities of the validators is known to the adversary, who can then corrupt validators in the selected quorum adaptively. The adversary is limited in the number of validations beyond k_2 by the *validation slack* which we have discussed earlier. We require that if a (k_1, k_2) -quorum system with validation slack s_v is used, then the partial spending fraction should not be $1/k_2$ but $1/k'_2$, where $k'_2 = k_2 + f/s_v$.

Corrupt buyer and seller: If both the buyer and the seller are corrupt, then it is possible to corrupt all the validators in the chosen quorum and erase the record of some payments. This should not create an issue because the amounts that can be erased are only those involving corrupt sellers. Since at settlement time everyone is contacted and a large quorum must validate the settlement, the buyer and the seller cannot simultaneously receive credit for the erased transaction.

In what follows, we present the solution and elaborate on each of its components.

7 PARTIAL SPENDING AND SETTLEMENT PROTOCOLS

We start by presenting the protocol for choosing quorums randomly, then we present the partial spending protocol and we finish with the settlement protocol. In what follows we assume that an asynchronous (k_1, k_2) -quorum system construction is known to all participants with the parameters $k_1, k_2, \alpha, \beta, \mu$ and m available as global constant values.

7.1 Random Quorum Selection

The (k_1, k_2) -quorum systems definition is just a mathematical construction that assumes that a quorum can be chosen randomly, so we need to specify how quorums can be chosen randomly by the seller and how to prevent corrupt sellers from fixing the membership of the chosen quorum.

The protocol (Algorithm 10 shown in the Appendix) allows the seller to choose a quorum of m validators and ensures that if the seller is correct, the quorum of validators is chosen uniformly at random and is not known to the adversary or to the buyer. The arguments to the algorithm are a buyer-provided transaction identifier $tx = \langle F, pkb, pks \rangle$ that ties the chosen quorum to the two parties and the specified fund, and a seller-generated random r -bit string N_s , where r is a security parameter. The transaction's identifier is concatenated (denoted with $||$) with N_s to obtain a seller transaction identifier T_s , which is used as a seed for quorum selection. This seed is guaranteed to be unique *w.h.p.* if the seller is not corrupt. The goal is to select m different servers. The seller uses the seed h concatenated with an index j to select the validator $Server(H(h||j))$, where $Server$ is a function that maps the output of H to server identifiers. Since the number of possible validators is n , it is possible that the same validator is chosen twice for different values of j , so the seller tries successive values of j until m distinct validators are chosen. The algorithm returns the identities of the selected validators. This information is all that is needed to verify later that a particular quorum of validators was properly chosen according to the protocol. In fact, the set of chosen validators is a deterministic function of the protocol arguments, but this requires knowledge of N_s without which an adversary cannot guess the identities of the validators.

It is important to point out that the randomness of the quorum is ensured by using $\langle F, pkb, pks \rangle || N_s$ as a seed quorum selection and that a corrupt seller cannot reuse an old seed to double spend because this would result in the same seller transaction identifier. One final consideration is ensuring random selection for a corrupt seller that attempts to run the algorithm multiple times in the hope of maximizing the number of previously corrupted validators in the chosen quorum or to present a different N_s at validation time to ensure that the quorum contains a large number of corrupt validators. Since the quorum system is asynchronous, *w.h.p.* there are no more than $(1 + \mu)p_fm$ previously corrupted validators in a randomly chosen quorum. If the seller wants to increase the number to $(1 + 2\mu)p_fm$, for example, the seller should make an exponential number of attempts (in

Algorithm 1 Partial spending: Buyer's Protocol

```

1: procedure BUYERPARTIALSPEND( $pks$ )
2:    $tx := \langle F, pkb, pks \rangle$  ▷ Transaction Id
3:   send (PAY,  $tx$ ) to  $pks$  ▷ send payment message with transaction Id
4:   wait ▷ Wait until commitments
5:   until (QUORUM,  $tx, Q_c = [c_i]$ ) received from  $pks$  ▷ received from seller
6:   if  $|Q_c| = m$  then ▷ If quorum has the correct size,
7:     send (SIGNED_QUORUM,  $tx, [\text{Sign}_{skb}(tx||h_s||c_i)]$ ) to  $pks$  ▷ sign and send response
8:   else ▷ otherwise
9:     return  $\perp$  ▷ abort

```

Algorithm 2 Partial spending: Seller's Protocol

```

1: procedure SELLERPARTIALSPENDING( $tx$ )
2:    $[v_i] = Q_{tx} \leftarrow \text{SELECTQUORUM}(tx, N_s \leftarrow \{0, 1\}^r)$  ▷ Seller selects  $m$  validators
3:    $[N_i] \leftarrow [(\{0, 1\}^r)^m]$  ▷  $N_i$  is seller's blinding nonce for  $i$ 'th validator
4:    $[c_i] \leftarrow H(v_i||N_i)$  ▷  $c_i$  is commitment to  $i$ 'th validator's identity
5:    $h_s = H(N_s)$  ▷ commitment to  $N_s$ 
6:   send (QUORUM,  $tx, h_s, [c_i]$ ) to  $pkb$  ▷ send commitments to buyer
7:   wait until (SIGNED_QUORUM,  $tx, [\sigma_i]$ ) received from  $pkb$  ▷ and wait for signatures
8:   for all  $v \in Q_{tx}$  do ▷ Send signatures with nonces
9:     send  $\langle tx, h_s, \sigma_i, N_i \rangle$  to  $v$  ▷ to validators
10:   $replies = witnesses = \emptyset$ 
11:  repeat ▷ Wait for replies from
12:    if received  $resp$  from  $v \in Q_{tx} \wedge v \notin replies$  then ▷ validators while keeping track of
13:       $replies = replies \cup \{v\}$  ▷ which validators replied and
14:      if  $resp = \langle \text{valid}, tx, h_s, \sigma = \text{Sign}_{skv}(tx||h_s) \rangle$  then ▷ which validators
15:         $witnesses = witnesses \cup (resp, v)$  ▷ are witnesses
16:  until  $|replies| \geq m - (1 + \mu)p_fm$ 
17:  if  $|witnesses| \geq (1 - \alpha) \times m$  then ▷ If enough validators validated,
18:    return  $\langle tx, N_s, witnesses \rangle$  ▷ return certificate for partial spending transaction
19:  else
20:    return  $\perp$ 

```

μp_fm) to choose a quorum. So, we assume that the quorum system is such that *w.h.p.* a randomly chosen quorum does not contain more than $(1 + \mu/2)p_fm$ previously corrupted validators, and that $(\mu/2)p_fm/2$ is large enough so that *w.h.p.* the computationally bounded corrupt client cannot choose a quorum with more than $(1 + \mu)p_fm$ previously corrupted validators.

The quorum selection protocol satisfies the following properties, which we prove in the appendix.

LEMMA 7.1. *If the seller is correct, the quorum of validators is chosen uniformly at random.*

LEMMA 7.2. *If the seller is correct, the identities of the correct validators in the chosen quorum are not known to the adversary or to the buyer.*

LEMMA 7.3. *For $0 < \mu < 1$, if a seller chooses K quorums, of size m each, at random, then with probability at most $Ke^{-\mu^2 \times p_fm / (2 + \mu)}$, every chosen quorum has no more than $(1 + \mu)p_fm$ previously corrupt validators.*

Algorithm 3 Validator v Payment Validation

```

1: if received  $\langle tx, h_s, \sigma, N \rangle$  from  $pks \wedge tx.F \notin settle \wedge$ 
    $(tx.pkb \in tx.F.owners) \wedge (tx.pks = pks) \wedge (tx.F \notin validated\_fund) \wedge$ 
    $VERIFYSIG(tx.pkb, tx||h_s||H(v||N)), \sigma)$  then
2:    $validated\_fund = validated\_fund \cup F$ 
3:    $validated\_transactions = validated\_transactions \cup \langle tx, h_s, \sigma, N \rangle$ 
4:   send  $\langle \text{valid}, tx, h_s, \text{Sign}_{sk_v}(tx||h_s) \rangle$  to  $pks$ 
5: else
6:   send  $\langle \text{invalid}, tx \rangle$  to  $pks$ 

```

7.2 Partial Spending Protocols

The partial spending protocols for the buyer and seller are shown in Algorithms 1 and 2 respectively. The validator code relating to partial spending is shown in Algorithm 3. The code closely follows the solution overview above. In the code we assume that the quorum parameters including the validation slack v_s are fixed and known by all parties and therefore the partial amount to be paid from fund F , which is equal to $F.bl/(k_2 + f/v_s)$, needs not be explicitly shown in the code.

7.2.1 Buyer's Code. The buyer sends a PAY message specifying the transaction identifier (recall that we only consider partial spending transactions in the protocols). The buyer then waits for a response from the seller which will be a vector of commitments $[c_i]$ to the identities of a quorum of validators (explained below in the overview of the seller's code). The buyer checks that the quorum size is m and sends a reply that includes for each validator a signature of $tx||c_i$ to link the commitment to the transaction. The buyer does not need to check that the chosen quorum is valid. That is done when the seller settles the payment.

7.2.2 Seller's Code. The seller's SELLERPARTIALSPENDING() function is executed in response to a PAY message from the buyer and takes the tx of the PAY message as argument. The seller starts by calling SELECTQUORUM() using the buyer's transaction identifier and a randomly generated nonce N_s . The call returns a quorum of validators Q_{tx} represented as a vector of validators $[v_i]$. The seller generates m r -bit nonces that it uses in generating a vector of commitments $[c_i] = [H(v_i||N_i)]$ to the validators identities, which it sends to the buyer, and waits to receive from the buyer the signed commitments. Then the seller sends to each validator a validation request that includes the transaction identifier tx , the signature σ_i (which should be equal to $\text{Sign}_{sk_b}(tx||c_i)$), and the nonce N_i used in generating the commitment. It then waits to receive replies from $m - (1 + \mu)p_f m$ validators and checks if $(1 - \alpha)m$ validators validated the transaction. If so, the transaction is validated and the replies from validators that validated the transaction constitute a *certificate* of validation for the transaction. In the code, we do not explicitly represent the resulting fund, but we note that the transaction, and therefore the fund that it creates, is specified by the transaction identifier $tx = \langle F, pkb, pks \rangle$ and h_s , the seller's commitment to the nonce N_s .

7.2.3 Validator's Code. The validator checks for the validity of the request and sends the result of the validation to the seller. A request is valid if: the fund from which the payment is made has not been settled; the seller specified by the transaction is the one making the request; the buyer specified in the transaction is the buyer specified in the fund of the transaction; the fund specified by the transaction has not been previously validated by the validator; the signature σ in the request is a valid signature for $tx||c$ by the buyer $tx.pkb$ specified in the request; and, c is a valid commitment to the validator's identity. If the request is valid, the validator adds the fund to the set of validated funds and the transaction to the list of validated transactions. The list of validated funds is needed

to ensure that a validator validates only one partial payment from a given fund. It is used during settlement of buyer's fund F to ensure that all partial spending from F will be accounted for. That is why σ , c and N which are needed to validate a transaction are stored along with tx .

7.3 Settlement Protocols

As we explained in the solution overview, during settlement there is a need to prevent the adversary from learning the identities of the validators prematurely and corrupting them. The settlement algorithms use a PROPAGATE() protocol that allows a party to propagate information to all but $2f$ correct servers without the adversary having the ability to suppress the information being propagated. We outline the PROPAGATE() protocol first (the details and code are in the appendix), then we present the settlement protocols for the seller, buyer and validators.

7.3.1 Propagating Information. As we discussed in the protocol overview, the adversary can corrupt a validator if it learns that the validator validated a particular transaction and the adversary wants to suppress that information. The identities of validators that are involved in validating a particular transaction are kept secret by the validators themselves (if honest) or the seller until settlement time. The PROPAGATE() protocol allows a party (seller or validator) to send a message to all validators so that the adversary would either have to corrupt the party to learn the message or, if the adversary learns the message without corrupting the party, then all but $2f$ honest are also guaranteed to learn the message. The PROPAGATE() protocol is implemented using secret sharing and reconstruction. To distinguish between messages propagated by different calls to PROPAGATE() by clients, a unique N_{prop} nonce is generated for the call and provided as a second argument to PROPAGATE(). When a message from client c is finally propagated to a server, the server will have $message[c, N_{prop}]$ equal to the propagated message. The details are given in the appendix.

7.3.2 Seller's Settlement and Corresponding Validation. The settlement algorithm for the seller is straightforward. The seller propagates a settlement request to validators. The fund is fully specified by $tx = \langle F, pks, pkb \rangle$ and N_s of the partial spending transactions that created the fund. The certificate of the fund consists of the set of *payment_witnesses* that validated the partial spending transaction. The information $tx, N_s, payment_witnesses$ enables a validator of the settlement transaction to recalculate the quorum used in validating the transaction and to verify that the witnesses are provided by members of that quorum. The validator checks that the quorum and validations provided by the seller are correct and that the validation for the transaction are received from a large enough subset of validators. In addition, the validator checks that either it has no record of SETTLE transaction for F ($F \notin settle$) or the transaction being settled, represented as $(tx, h_s = H(N_s))$ was added to *transactions*[F] which is the set of transactions calculated when the buyer settles F . This check is needed to avoid double spending. As shown in the proofs, if the buyer settles F , every payment to an honest seller will be added to *transactions*[F] which ensures that this condition will be satisfied for honest sellers. If all the information checks out, the validator sends a validation for the settlement. The seller waits until it receives $n - f$ validations, which is guaranteed to happen if the seller is correct. These $n - f$ validations guarantee that only one fund can result from settling a partially certified fund. Finally, note that the settlement of the seller's partially certified fund can go through even the validators that validated the original spending transaction are corrupted.

7.3.3 Buyer's Settlement and Corresponding Validation. To settle the buyer's fund F , we need to make sure that all partial spending from the buyer's fund are deducted from the settled balance. The buyer's starts by sending a settlement request to all validators. Every validator that receives the buyer's request *propagates* information about any payments from F to sellers that it is aware of (there can only be at most one such payment per validator). A validator will either provides proof

Algorithm 4 Seller Fund Settlement

```

1: procedure SELLERSETTLE( $tx = \langle F, pkb, pks \rangle, N_s, payment\_witnesses$ )
2:    $N_{settle} \leftarrow \{0, 1\}^r$ 
3:   PROPAGATE( $\langle tx, N_s, payment\_witnesses, SETTLE \rangle, N_{settle}$ )
4:   repeat
5:     if received (valid,  $\sigma, N_{settle}$ ) from  $v$  then
6:        $Id_1 = H(tx || N_s || PAY)$ ;    $Id_2 = H(Id_1 || SETTLE)$ 
7:        $F'.id = Id_2$ ;  $F'.bl = F.bl/k'_2$ ;  $F'.owners = \{pks\}$ 
8:       if  $\sigma = \text{Sign}_v(\langle F'.id, F'.bl, F'.owners \rangle)$  then
9:          $settle\_witnesses = settle\_witnesses \cup \{(v, \sigma)\}$ 
10:    until  $|settle\_witnesses| \geq n - f$ 
11:     $F'.certificate = settle\_witnesses$ 
12:    return ( $\langle F' \rangle$ )

```

Algorithm 5 Code of Validator v for Seller Fund Settlement

```

1: if  $message[pks, N_{settle}] = \langle tx, N_s, payment\_witnesses, SETTLE \rangle \wedge$ 
    $Q = \text{SELECTQUORUM}(tx, N_s) \wedge ((tx, H(N_s)) \in transactions[F] \vee F \notin settle) \wedge$ 
    $\text{VALIDATEDTRANSACTION}(tx, payment\_witnesses) \wedge$ 
    $Q' = \{q : (resp, q) \in payment\_witnesses\} \subseteq Q \wedge |Q'| \geq m - m\mu p_f$  then
2:    $transactions[F] = transactions[F] \cup \{(tx, H(N_s))\}$ 
3:    $Id_1 = H(tx || N_s || PAY)$ 
4:    $Id_2 = H(Id_1 || SETTLE)$ 
5:    $F'.id = Id_2$ ;  $F'.bl = F.bl/k'_2$ ;  $F'.owners = \{pks\}$ 
6:   send (valid,  $\text{Sign}_v(\langle F'.id, F'.bl, F'.owners \rangle, N_{settle})$ ) to  $pks$ 

```

that it witnessed a payment or states that it did not witness any payment. The proof of a witnessed payment (tx, σ, N_h) where $\sigma = \text{Sign}_{pkb}(H(tx || N_h))$ is the blind signature and N_h is the blinding factor. Validators wait until $n - f$ different messages about payments from F are propagated. Every validator then counts the number of different payment transactions from F that it heard about, calculates the resulting balance after subtracting the amounts for those payments and send the seller a signed balance for the settlement fund (see problem definition). The buyer waits until it receives $n - 2f$ signatures for the same balance which will form the set of witnesses for the settlement fund resulting from F . The code is given in the appendix.

8 CONCLUSION

We introduced (k_1, k_2) -quorum systems and shown how to use them to execute payment transactions with less than f validations per transaction. By carefully considering intersection and non-intersection properties, we are able to allow up to k_1 non-interfering payments in parallel and prevent double spending by limiting the number of such concurrent transactions. It might seem that the construction succeeds because when the seller is honest, the quorum is randomly selected and has, with high probability, a small fraction of corrupt validators. We observe that in the case of corrupt sellers, the whole quorum could be corrupted, and our definition of k'_2 was done explicitly to handle that possibility. The constructions we proposed work for large values of n . It is not clear how to modify the constructions to work for smaller values of n and larger values of f relative to n , or how to reduce the difference between k_1 and k_2 . This is a subject for future work.

REFERENCES

- [1] Amitanand Aiyer, Lorenzo Alvisi, and Rida A Bazzi. On the availability of non-strict quorum systems. In *International Symposium on Distributed Computing*, pages 48–62. Springer, 2005.
- [2] Amitanand S Aiyer, Lorenzo Alvisi, and Rida A Bazzi. Byzantine and multi-writer k-quorums. In *International Symposium on Distributed Computing*, pages 443–458. Springer, 2006.
- [3] Zeta Avarikioti, Eleftherios Kokoris-Kogias, Roger Wattenhofer, and Dionysis Zindros. Brick: Asynchronous incentive-compatible payment channels. In *International Conference on Financial Cryptography and Data Security*, pages 209–230. Springer, 2021.
- [4] Amos Beimel. Secret-sharing schemes: A survey. In *International conference on coding and cryptology*, pages 11–46. Springer, 2011.
- [5] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
- [6] Christian Cachin, Klaus Kursawe, Anna Lysyanskaya, and Reto Strobl. Asynchronous verifiable secret sharing and proactive cryptosystems. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 88–97, 2002.
- [7] Christian Decker and Roger Wattenhofer. A fast and scalable payment network with bitcoin duplex micropayment channels. In *Proceedings of the 17th International Symposium on Stabilization, Safety, and Security of Distributed Systems - Volume 9212*, page 3–18, 2015.
- [8] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE transactions on information theory*, 31(4):469–472, 1985.
- [9] Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, and Dragos-Adrian Seredinschi. Scalable byzantine reliable broadcast. In *33rd International Symposium on Distributed Computing, DISC 2019*, volume 146 of *LIPICs*, pages 22:1–22:16, 2019.
- [10] Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovič, and Dragos-Adrian Seredinschi. The consensus number of a cryptocurrency. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC '19, page 307–316, 2019.
- [11] Saurabh Gupta. A non-consensus based decentralized financial transaction processing model with support for efficient auditing. Master's thesis, Arizona State University, June 2016.
- [12] Dahlia Malkhi, Michael Reiter, and Rebecca Wright. Probabilistic quorum systems. In *Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing*, pages 267–273, 1997.
- [13] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, USA, 2005.
- [14] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008. URL: <https://bitcoin.org/bitcoin.pdf>.
- [15] Alejandro Ranchal Pedrosa, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. Scalable lightning factories for bitcoin. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, SAC '19, page 302–309, New York, NY, USA, 2019. Association for Computing Machinery. doi: 10.1145/3297280.3297312.
- [16] Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments, 2016.
- [17] Ronald L. Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 26(1):96–99, 1983.
- [18] Jakub Sliwinski and Roger Wattenhofer. Abc: Proof-of-stake without consensus, 2019. URL: <https://arxiv.org/abs/1909.10926>, doi: 10.48550/ARXIV.1909.10926.

APPENDIX

This appendix contains the following:

- (1) **Partial spending from partially validated funds** (Section I): This section includes a description of how partially validated funds can be spent without full validation. It turns out that this is possible if the buyer is correct, but a corrupt buyer can prevent some such payments, but it will be detected.
- (2) **Impossibility Proofs** (Section II): shows that no solution to the partial spending problem that tolerates an adaptive adversary can be deterministic or allow the spending of the whole balance.
- (3) **Propagating Information** (Section III): This section includes the protocols for propagating information and its correctness proof.
- (4) **Partial Spending and Settlement** (Section IV): This section provides code for the buyer settlement protocol and the proofs that our solution for the payment and settlement satisfy the problem requirements.
- (5) $(k_1 - k_2)$ -**Quorum Systems Properties** (Section V): This section provides the proofs for the lemmas relating to $(k_1 - k_2)$ -quorum systems that are listed in the main text.
- (6) **Quorum Selection Protocol Proofs** (Section VI): This section contains proofs of the lemmas relating to the quorum selection protocol.

I PARTIAL SPENDING FROM PARTIALLY VALIDATED FUNDS

If a buyer pays a seller with a partial spending transaction, how can the seller turn around and spend some of the received amount without first settling the fund? It turns out, maybe not surprisingly, that in our asynchronous system model with an adaptive adversary, this does not seem to be achievable without some restrictions. One issue is the inability of the seller to prove to a potential buyer that it has a properly validated level-1 payment because that would require divulging the identities of the validators, which is problematic as we described. Also, it is not clear that a zero-knowledge proof of validity is possible. A seller who is the recipient of level-1 payment has only the validation from the buyer (the signatures of the hashes) that it can share as part of a level-2 payment. This would at least prevent a seller who has not received a level-1 payment from attempting to make a level-2 payment, but this is a weak guarantee because both the buyer and the seller can be corrupt.

We outline the main idea for supporting level- i payments, but do not present any protocols. The idea for supporting level- i payments, $i > 1$, is the following. We start with a buyer who has an fully certified fund. The buyer can make up to k'_2 level-1 partial spending transactions. If we want to spend from the resulting k'_2 funds (level-2 payments), we can do k'_2 partial spendings from each one of them resulting in k'^2_2 level-2 partial spending transactions. So, in general, our approach will allow a total of k'^i_2 level- i spending transactions, $i \geq 1$.

Of course all of this should be done so that double spending is not possible. For level-1 payments, the total spending can be as high as $k'_2 * F.bl/k'_2 = F.bl$, and potentially there might be no funds left for level-2 payments! One way to deal with this is to reduce the partial spending fraction from $1/k'_2$ to $1/2k'_2$. This way, the total of level-1 payments is $k'_2 * F.bl/2k'_2 = F.bl/2$. The total for level-2 payments is $k'^2_2 * F.bl/4k'^2_2 = F.bl/4$. So the total spending at all levels is $F.bl/2 + F.bl/4 + F.bl/8 + \dots \leq F.bl$. So, essentially, multi-level spending from partially certified funds can be achieved by increasing the amount that cannot be spent from a given balance. Also, it would require that the fund *level* be part of the fund description, which could be maintained as a chain of spending transactions starting from level-1 up to the level of the fund.

One limitation of this approach is that it does not guarantee for every recipient of a level-1 payment the ability to make k'_2 level-2 payments. Instead, the guarantee is on the total number of level-2 payments. With this solution, a corrupt buyer, can issue many level-1 payments that are not properly validated to sellers under the control of the adversary. These sellers can then make level-2 payments that *consume* the total quota of level-2 payments. So, if the buyer is corrupt, a seller can be denied the opportunity to make level-2 payments, but the buyer's corruption will be detected. The situation can be worse as we go up the levels. Nonetheless, the fact that level-2 and higher payments can be made at all is surprising.

II IMPOSSIBILITY PROOFS

We show that a deterministic solution cannot satisfy all the problem's requirements (Section 4.2) and that no solution that uses less than full quorums can spend the whole amount and avoid double spending.

LEMMA II.1. *If all spending transactions are partially certified (certified by quorums of size less than $f + 1$), then a deterministic solution to the spending problem defined in Section 4.2 cannot guarantee progress with w.h.p. for correct buyers and sellers.*

PROOF. In a deterministic solution, the identities of the validators for a given transaction is solely determined by the buyer and seller identifiers pkb and pks , the history of transactions at the buyer and seller and the initial states of the buyers and sellers. Consider an execution in which all payments are made by one buyer and assume that the buyer pays the seller with a partial spending transaction. The identity of the validators that will be used by the seller can be calculated by the buyer because it has all the information needed to do that. If the buyer is corrupt, the adversary can learn the identities of the validators for the transaction and corrupts all of them because they number less than $f + 1$. The seller can then be prevented from successfully settling the fund resulting from the transaction. \square

LEMMA II.2. *If all quorums used in validation have size less than $f + 1$, then it is not possible to prevent double spending and allow the spending of the whole balance of a fund.*

PROOF. We consider an execution in which a buyer spends the whole amount in a fund F by executing k partial spending transactions, T_1, T_2, \dots, T_k , for some k , to honest sellers who validate and accept the payments. No validators are corrupted during the executions of these k partial spending transactions. After these partial spending transactions are validated, the adversary corrupts the buyer who issues a partial spending transaction T to a corrupt seller. To validate the transaction, less than $f + 1$ validators are contacted to validate the transaction. In a general solution, some of the validators for a transaction are contacted directly by the seller and buyer and some of them could be contacted indirectly by other validators. Since the adversary controls the buyer and seller, it knows the identities of the validators that are directly contacted and can corrupt them because there are at most f validators involved in validating the transaction. In turn, and for the same reason, this allows the adversary to learn the identities and corrupt any validators that are used to validate the transaction whether they are directly or indirectly contacted by the seller and buyer. Since all the validators for transaction T can be corrupted by the adversary, the transaction can be successfully validated and successfully settled. Since the transactions T_1, T_2, \dots, T_i are to honest sellers, according to the problem requirements, the sellers for these transactions should be able to successfully settle the funds resulting from them, resulting in a total spending that exceeds the original balance F , a double spending. \square

Algorithm 6 Propagating Information: Client c 's code

```

1: procedure PROPAGATECLIENT( $message, N_{prop}$ )
2:    $[s_i] = \text{SECRETSHARE}(message, n, f)$ 
3:   for  $i := 1 \rightarrow n$  do
4:     send (SHARE,  $N_{prop}, s_i, \text{Sign}_c(s_i || v_i || N_{prop})$ ) to  $v_i$ 

5:    $acks[N_{prop}] = \emptyset$ 
6:   repeat
7:     if received (SHARE_ACK,  $N_{prop}$ ) from  $v$  then
8:        $acks[N_{prop}] = acks[N_{prop}] \cup \{v\}$ 
9:     until  $|acks[N_{prop}]| \geq n - f$ 
10:    send (RECONSTRUCT,  $N_{prop}$ ) to validators
11:     $reconstructed[N_{prop}] = \emptyset$ 
12:    repeat
13:      if received (RECONSTRUCTED,  $message, N_{prop}$ ) from  $v$  then
14:         $reconstructed[N_{prop}] = reconstructed[N_{prop}] \cup \{v\}$ 
15:      until  $|reconstructed[N_{prop}]| \geq n - f$ 

```

III PROPAGATING INFORMATION: PROTOCOLS AND PROPERTIES

As we discussed in the protocol overview, the adversary can corrupt a validator if it learns that the validator validated a particular transaction. The identities of validators that are involved in validation are kept secret by the validators themselves (if honest) or the seller until settlement time, at which time they need to be divulged. The PROPAGATE() protocol allows a party (seller or validator) to send secret information to all validators so that the adversary would either have to corrupt the party to learn the information or, if the adversary learns the information without corrupting the party, then all but f honest are also guaranteed to learn the information. The PROPAGATE() protocol is implemented using secret sharing and reconstruction. The code is given in Algorithms 6 and 7.

Propagating information is relatively simple in our setting and is easier than the asynchronous verifiable secret sharing problem [6] (a solution to which can also be used, but would be an overkill). The client protocol takes as input two arguments: the $message$ being propagated and a unique nonce N_{prop} to distinguish between different calls to the propagate protocol. The client c who propagates the information creates shares of the information that needs to be propagated so that any $f + 1$ out of n shares can be used to reconstruct the $message$. Depending on the settlement protocol, c is either a seller or a validator. Each share s_i is addressed to a particular validator v_i . In addition to the share s_i , v_i receives from the client a signature $\sigma_i = \text{Sign}_c(s_i || v_i || N_{prop})$ that v_i can later use to prove that s_i is a properly received share. After the shares are distributed and $n - f > f + 1$ validators acknowledge receiving the shares, the client sends a RECONSTRUCT message so that the validators can reconstruct the $message$ from the shares. Validators broadcast their shares and collect shares that they authenticate until $f + 1$ authenticated shares are collected. At that point, the $message$ is reconstructed using the collected shares. It is important to note here that all the shares used in the reconstruction are validated to be from the client (line 9 of validator code), so if the client is honest, the $f + 1$ shares will all be valid even if they are received from corrupt validators. If the client is not honest, we don't care about the value that is reconstructed. A validator stops participating in the propagation protocol for a particular N_{prop} when $n - f$ different validators announce that they have reconstructed the messages. Of those validators, $n - 2f$ must be honest.

Algorithm 7 Propagating Information: validator v 's code. The client c in the code can be either a seller invoking the client side of information propagation to settle a fund resulting from partial payment or a validator invoking propagation information to handle a buyer's settlement transaction

```

1: procedure PROPAGATESERVER( $N_{prop}$ )
2:   repeat
3:     upon receipt of (SHARE,  $N_{prop}$ ,  $s$ ,  $\sigma = \text{Sign}_{pk_s}(s||v||N_{prop})$ ) from  $c$ :
4:       send (SHARE_ACK,  $N_{prop}$ ) to  $c$ 
5:        $share[c, N_{prop}] = s, \sigma$ 
6:     upon receipt of (RECONSTRUCT,  $N_{prop}$ ) from  $c$ :
7:       send (FORWARD,  $c$ ,  $share[c, N_{prop}]$ ,  $N_{prop}$ ) to validators
8:        $shares[c, N_{prop}] = \emptyset, Forwarded[c, N_{prop}] = \emptyset$ 
9:     upon receipt of (FORWARD,  $c$ ,  $s$ ,  $\sigma$ ,  $N_{prop}$ ) from  $v$ :
10:      if  $\sigma = \text{Sign}_c(s||v||N_{prop}) \wedge v \notin Forwarded[c, N_{prop}]$  then
11:         $shares[c, N_{prop}] = shares[c, N_{prop}] \cup \{(s, \sigma)\}$ 
12:         $Forwarded[c, N_{prop}] = Forwarded[c, N_{prop}] \cup \{v\}$ 
13:      if  $|shares[c, N_{prop}]| \geq f + 1$  then
14:         $message[c, N_{prop}] = \text{RECONSTRUCTSECRET}(shares[c, N_{prop}])$ 
15:        send (RECONSTRUCTED,  $message[c, N_{prop}]$ ,  $N_{prop}$ ) to  $c$  and validators
16:   until  $(n - f)$  RECONSTRUCT messages received

```

The algorithms guarantees that if the client calling $\text{PROPAGATE}(message)$ is honest, all but $2f$ honest validators will receive $message$. If the client propagating the shares is not honest, there are no guarantees as to what is reconstructed and the reconstruction might even fail, but that does not affect the algorithms that use the $\text{PROPAGATE}()$ function.

LEMMA III.1. *If an honest client calls $\text{PROPAGATE}(m)$, all but $2f$ honest validators will be guaranteed to receive m .*

PROOF. (sketch) If the client is honest, when a validator receives a RECONSTRUCT message, $n - f$ validators must have already received the SHARE message. Of these validators $n - 2f \geq f + 1$ validators are honest and will each send FORWARD message containing its share to every other honest validator who will be able to reconstruct the message. Validators stop their execution when they receive $n - f$ RECONSTRUCT messages from other validators. Of these $n - 2f$ are from honest validators. \square

LEMMA III.2. *Let N be a random value that is generated by a client c such that no value that is directly or indirectly dependent on N is sent to any validator with the exception of $H(N)$. If the adversary learns N w.h.p., then w.h.p. the adversary must have corrupted the client c .*

PROOF. (sketch) The proof is by induction on the number of messages sent by c . If c sends no messages, it should be clear that the adversary cannot learn N without corrupting c . Assume that the lemma holds for the first i messages sent by c and consider the $i + 1$ message sent by c . The $i + 1$ message either contains no value dependent on N or contains $H(N)$ and other values that are not dependent on N . In the first case, the adversary clearly doesn't learn anything about N . In the second case, given the adversary's bounded computational power, it cannot learn the value of N except with negligible probability. It follows that the chain must be a 0-length chain and the adversary learns N by corrupting c . \square

Algorithm 8 Buyer Fund Settlement

```

1: procedure BUYERSETTLE( $F$ )
2:   send (SETTLE,  $F$ ) to validators
3:   repeat
4:     if received  $resp$  from  $q_i$  then                                ▶ Wait for replies from
5:        $replies = replies \cup resp$                                        ▶ validators until enough
6:   until  $\exists F' : |\{resp : resp \in replies \wedge resp = F'\}| = n - 2f$  ▶ identical replies are received
7:   return  $F'$ 

```

LEMMA III.3. Let N be a random value that is generated by a client c such that no value that is directly or indirectly dependent on N is sent to any validator with the exception of $H(N)$ or by executing $PROPAGATE(m)$ for a message m that depends on N . If the adversary learns N w.h.p., then w.h.p. the adversary must have corrupted the client c or all but $2f$ honest validators are guaranteed to learn m .

PROOF. (sketch) Since each share generated by the secret sharing of the $PROPAGATE(m)$ is independent of N , by an argument similar to that in the previous lemma, the adversary cannot learn N w.h.p. if it does not learn $f + 1$ shares. If the adversary learns N , then it must have corrupted a validator that received $f + 1$ valid shares. One of these shares must be from an honest validator that received a RECONSTRUCT message. It follows that $n - f$ validators must have received a reconstruct message and all of these validators send FORWARD messages to all other validators who will be able to reconstruct m . Validators participate in the protocol until $n - f$ validators announce that they reconstructed the message at which point $n - 2f$ honest validators must have reconstructed the message. \square

It might seem a little strange that even though every honest validator sends messages to every other honest validator, we only guarantee that $n - 2f$ validators will receive the message if the adversary doesn't corrupt the client. The reason is that even though honest validators might be guaranteed to receive messages sent by other honest validators eventually, we need to make a statement about what holds when the execution of $PROPAGATE(m)$ ends.

IV PARTIAL SPENDING AND SETTLEMENT PROOFS

The protocol for buyer fund settlement is shown in Algorithms 8 and 9. The rest of this section presents the correctness proof of the solution for partial spending and settlement protocols.

We show that the protocol satisfies the problem requirements.

LEMMA IV.1 (TRANSACTIONS TO HONEST SELLERS ARE ACCOUNTED FOR). Let $T_s = (tx, N_s)$ be a partial spending transactions for which the seller is honest. If the buyer executes (SETTLE, F), then $n - 2f$ honest validators will have $(tx, h_s = H(N_s)) \in transactions[F]$.

PROOF. (sketch) If the seller already executed a settlement transaction for the fund resulting from T_s , then it should have received validations from $n - f$ validators each of which adds $(tx, h_s = H(N_s))$ to $transactions[F]$. So, we assume that the seller did not execute a settlement transaction for the fund resulting from T_s and did not divulge N_s . It follows that the identities of the validators of T_s are not known to the adversary which means, by the model assumptions, that some of the honest validators of T_s will succeed in propagating their information to $n - 2f$ honest validators that add (tx, h_s) to $transactions[F]$. \square

LEMMA IV.2. All partially certified funds with honest owners can be settled successfully.

Algorithm 9 Code of Validator v for Buyer Fund Settlement

```

1: procedure VALIDATORBUYERSETTLE( $F, pkb$ )
2:   if received (SETTLE,  $F$ ) from  $pkb \wedge pkb \in F.owners$  then
3:      $N_{prop} \leftarrow \{0, 1\}^n$  ▷ different validators choose different  $N_{prop}$  values
4:      $value[F, N_{prop}] = \perp$  ▷  $F$  will have multiple entries, one for each such value
▷ If  $v$  is witness to payment from  $F$ , propagate payment info
5:     if  $\exists (tx, h_s, \sigma, N) \in validated\_transactions \wedge tx.F = F \wedge$  and  $pkb \in tx.owners$  then
6:       PROPAGATE( $\langle\langle F, SETTLE \rangle\rangle, v, (tx, h_s, \sigma, N), N_{prop}$ )
7:     else ▷ otherwise propagate that  $v$  is not witness to any payment from  $F$ 
8:       PROPAGATE( $\langle\langle F, SETTLE \rangle\rangle, v, Sign_v(F||none), N_{prop}$ )
▷ Update known transactions[ $F$ ] based on propagated information
9:   if  $message[F, N_{prop}] = \langle\langle F, SETTLE \rangle\rangle, v, (tx, h_s, \sigma, N), N_{prop}$  then
10:     $transactions[F] = transactions[F] \cup \{(tx, h_s)\}$ 
▷ Forward any witness or non-witness information received
11:  if  $message[F, N_{prop}] = (\langle F, SETTLE \rangle, v', *)$  then
12:    send ( $F, SETTLE, value[F, N_{prop}]$ ) to all validators
13:  if received ( $\langle F, SETTLE \rangle, v', *) \wedge v' \notin settle\_validators[F]$  then
14:     $settle\_validators[F] = settle\_validators[F] \cup v'$ 
15:  if received ( $F, SETTLE, v', (tx, \sigma, N)$ )  $\wedge v' \notin settle\_validators[F] \wedge$ 
    VALIDPAYMENT( $F, (tx, h_s, \sigma, N)$ ) then
16:     $payments[F] = payments[F] \cup (tx, h_s, \sigma, N)$ 
17:  if  $|settle\_validators[F]| = n - f$  then
18:    for all  $(tx, h_s, \sigma, N) \in payments[F] : VALIDPAYMENT(tx, h_s, \sigma, N)$  do
19:       $transactions[F] = transactions[F] \cup \{(tx, h_s)\}$ 
20:       $settle = settle \cup \{F\}$ 
21:      if  $|transactions[F]| > k_1$  then ▷ if many transactions were issued from  $F$ 
22:        abort ▷ abort
▷ Add  $F$  to funds with SETTLE transactions
23:     $new\_balance = F.bl - |transactions[F]| \times 1/k'_2$ 
24:     $F'.tx = H(F.tx) ; F'.bl = new\_balance ; F'.owners = F.owners$ 
25:    send ( $F', Sign_v(F')$ ) to  $F'.owners$ 

```

PROOF. (sketch) Consider a partially certified fund F_s resulting from a partial spending transaction tx from a fund F . If the owner (seller) of F_s is honest, then when the payment was made, the seller selected a quorum according to the quorum selection protocol and then got the payment validated by a sufficient number of honest validators from the selected quorum. The identities of these validators are not known to the adversary before the owner settles the fund. Also, each of these honest validator received the validation request from the seller before receiving a settlement request for F because one of the conditions for validating a payment request is for the validator not to have F is the set of funds for which there is a SETTLE transaction ($F.tx \notin settle$). Every honest validator that receives the settlement request for F_s will either have $F \notin settle$ or $F \in settle$. If $F \in settle$, by the previous lemma, the validator must have added F_s to $transactions[F]$. The validation of the settlement of F_s will go through in both cases because the other conditions for validating a seller's

settlement transaction will hold because the honest seller provides a correctly selected quorum and a large enough set of witnesses that validates the partial spending transactions that created F_s . \square

LEMMA IV.3. *Settlement for a partially certified funds equals payment amount: If F'' is a fully certified fund resulting from executing (F', SETTLE) for partially validated fund F' , then $F''.bl = F'.bl$.*

PROOF. (sketch) This follows immediately from the code. When a fund F' of a partial spending transaction is settled, the balance of the resulting fund F'' is equal to $F.bl/k'_2$ which is also the spending amount (which we do not explicitly represent in the protocols). \square

LEMMA IV.4. *Settlement amounts for payments from F are subtracted from settlement for F : If executing transaction (F, SETTLE) results in a fully certified fund F_R :*

$$F_R.bl \leq F.bl - \sum_{F' \in \text{settled}_F} F'.bl$$

PROOF. (sketch) Consider $F' \in \text{settled}_F$ that results from settling a fund F_s identified by transaction (tx, N_s) when F_s is settled, resulting in F' , either F is not yet settled and $(tx, H(N_s))$ gets added to $\text{transactions}[F]$ by the validators of the settlement transaction or $(tx, H(N_s))$ is already in $\text{transactions}[F]$. In either case, when \square

LEMMA IV.5 (NO MORE THAN $k_2 + f/v_s$ PARTIAL SPENDING TRANSACTIONS). *With high probability, no more than $k_2 + f/v_s$ partial spending transaction can be validated, where v_s is the validation slack.*

PROOF. With high probability, k_2 is an upper bound on the number of partial spendings that can be validated without any failures. If the adversary can only randomly chose the validators to corrupt, which is the case when the seller is honest, then k_2 will also be the number of partial spendings that are possible. So, we assume that k_2 partial spending transactions are executed with no validator corruption. Any additional transactions would have quorums of validators that intersect the previously selected quorums in βm validators and the adversary would need to corrupt $(\beta - \alpha)m$ for each additional partial spending transactions so that the number of validators that validate the transaction is raised to αm . A total of $f/((\beta - \alpha)m)$ additional spending transactions. The total number of partial spending transactions possible is therefore $k_2 + f/((\beta - \alpha)m) = k_2 + f/v_s$. \square

LEMMA IV.6. *If the owner of a fund F is honest, settlement amount is no less than the the initial balance of F minus payments made from F : If executing transaction (F, SETTLE) results in a fully certified fund F_R and F .owners is honest:*

$$F_R.bl \geq F.bl - \sum_{F' \in \text{funds}_F} F'.bl$$

PROOF. (sketch) This follows directly from how the settlement amount is calculated in which only transactions originating from the owner of F are deducted from the resulting balance. \square

LEMMA IV.7 (NON-INTERFERENCE). *If a total of $k \leq k_1$ payment transactions are initiated by $pkb \in F$.owners from fully certified fund F and no additional payment or settlement transactions are initiated by F .owners and pkb is honest, then, every one of the k transactions whose seller (payee) is honest will be validated.*

PROOF. (sketch) We consider one of these transactions with an honest seller. The seller will select a quorum of validators according to the protocol and gets validations from the buyer. Then the seller will contact the validators to get the transaction validated. By the properties of (k_1, k_2) -quorum systems, the seller will get replies from $m - (1 + \mu)p_f m$ validators, $(1 - \alpha) \times m$ of which have not previously validated another transaction from F and the transaction will be validated. \square

LEMMA IV.8. *Successful settlement for fully certified funds: If the owner of fully certified fund F is honest and executes an (F, SETTLE) transaction, the settlement transaction for F will terminate.*

PROOF. lemma Since the owner is honest, at most k_1 partial spending transactions from F can be executed. When the (F, SETTLE) transaction is executed, the validators will only have no more than k_1 transactions in $\text{transaction}[F]$ because every transaction in $\text{transaction}[F]$ must be checked for validity with $\text{VALIDPAYMENT}(tx, h_s, \sigma, N)$ which checks that σ is a valid signature by the buyer for $tx || h_s || N$. Since the only potentially blocking condition in the validator's buyer settlement code is when there are more than k_1 transactions in $\text{transaction}[F]$, every honest validator that handles the settlement will validate (F, SETTLE) and eventually the buyer will get $n - 2f$ identical replies from validators.

PROOF. lemma

V (k_1, k_2) -QUORUMS PROPERTIES

LEMMA 5.4. *An (m, n) uniform balanced quorum system is a (k_1, k_2) -quorum system with $\alpha = 1/3$ and $\beta = 2/3$ has $\epsilon = \delta = \text{negl}(m)$ and validation slack $= m/3$, if $p_f + \alpha_1 < 1/3$, where $\alpha_1 = k_1 m/n$ and $p_f = f/n$.*

PROOF. We need to show that:

- (1) Lower bound: $\Pr_{Q, Q_j \leftarrow Q; j \in J_1; |J_1| \leq k_1} [|Q \cap (\mathcal{F}_{pr} \cup \bigcup_{j \in J_1} Q_j) | > m/3] \geq 1 - \text{negl}(m)$
- (2) Upper bound: $\Pr_{Q, Q_j \leftarrow Q; j \in J_2; |J_2| \geq k_2} [|(Q \cap (\bigcup_{j \in J_2} Q_j)) - \mathcal{F}_{pr}| \leq 2m/3] \geq 1 - \text{negl}(m)$

For the lower bound, consider the intersection of a quorum Q with the union of up to k_1 previously selected quorums. Since $k_1 m/n = \alpha_1$, $\alpha_1 n$ is an upper bound on the size of the union of the k_1 previously selected quorums because each of them is of size m . The expected size of the intersection of Q with these k_1 previously selected quorums is at most $\alpha_1 |Q| = \alpha_1 m$ because the probability that a given randomly chosen server in Q is in the union is at most $\alpha_1 n/n = \alpha_1$. Similarly, the probability that a server in Q is previously corrupted is at most p_f and the expected number of servers in Q that are previously corrupted is at most $p_f |Q| = p_f m$. So, the expected size of the intersection of Q with the union of the previously corrupted servers together with the union of k_1 previously selected quorums is at most $(\alpha_1 + p_f)m$. Since $\alpha_1 + p_f < 1/3$, $(\alpha_1 + p_f)m < m/3$. Let $r = (1/3(\alpha_1 + p_f)) - 1$. This value of r is positive and independent of m . For this value of r , we have $(1+r) \times (\alpha_1 + p_f)m = m/3$. So, the probability that the size of the intersection exceeds by a factor $1+r$ the expected size of the intersection is the same as the probability that the size of the intersection exceeds $m/3$

By Chernoff's upper tail bounds [13], the probability that the size of the intersection exceeds by a factor of $1+r$, $r > 0$, the upper bound $(\alpha_1 + p_f)m$ on the expected size is:

$$\Pr_{Q \leftarrow Q} [|Q \cap (\mathcal{F}_{pr} \cup \bigcup_{j \in J_1} Q_j) | \geq (1+r) \times (\alpha_1 + p_f)m = m/3] \leq e^{-r^2 \times (\alpha_1 + p_f)m/3}$$

which is a negligible function of m . Now, it remains to show that the upper bound requirement holds:

$$\Pr_{Q, Q_j \leftarrow Q; j \in J; |J| \geq k_2} [|(Q \cap (\bigcup_{j \in J} Q_j)) - \mathcal{F}_{pr}| < 2m/3] \geq 1 - \text{negl}(m)$$

The expected number of servers in a quorum Q that are in the union of k_2 previously selected quorums is $(k_2 m/n)|Q| = (1 - \alpha_1)m$. The expected number of servers in Q that are previously corrupted is less than $p_f m$. So, the expected number of non-corrupted servers in Q that are in one of the last k_2 chosen quorums is at least $(1 - \alpha_1)m - p_f m = (1 - (\alpha_1 + p_f))m > 2m/3$ because

$\alpha_1 + p_f < 1/3$ as we have noted above. Again, if we define $r = 1 - ((1 - (\alpha_1 + p_f))/(2/3))$, we have $0 < r < 1$, and $(1 - r) \times (1 - (\alpha_1 + p_f))m = 2m/3$. Using Chernoff's lower tail bound, we have

$$\Pr_{Q, Q_j \leftarrow Q; j \in J; |J| \geq k_2} [|(Q \cap (\bigcup_{j \in J} Q_j)) - \mathcal{F}_{pr}| \leq (1 - r) \times (1 - (\alpha_1 + p_f))m = 2m/3] \leq e^{-r^2(1 - (\alpha_1 + p_f))m/2}$$

which is a negligible function of m . It is important to note that r does not depend on m but on the value of $\alpha_1 + p_f$ which is a constant for a given p_f and k_1 and k_2 . So, the function above decrease exponentially with m for fixed α_1 and p_f satisfying $\alpha_1 + p_f < 1/3$. Also, for a fixed k_1 and α_1 , m grows linearly with n : $m = \alpha_1 n / k_1$. In other words, the probabilities are negligible functions of n for fixed k_1 , α_1 and p_f . \square

Even though the proof shows that the probability of not satisfying the (k_1, k_2) -quorum requirements is negligible, values of α_1 and p_1 for which the sum $\alpha_1 + p_f$ is very close to $1/3$ result in probability bounds that are not useful in practice. For the parameter choices we made, the *validation slack* of the system is $m/3$. Different values for α_1 and p_f could result in lower validation slack but better overall performance. Determining the optimal combination of parameters is subject of future work.

LEMMA 5.6. *An (m, n) uniform balanced quorum system is a (k_1, k_2) -quorum system is a $(k_1, k_2, \epsilon, \delta, \alpha = 1/3, \beta = 2/3, \mu)$ -asynchronous quorum system with $\epsilon = \delta = \text{negl}(m)$ and validation slack = $m/3$, if $n > 8f$ and $k_1 m / n < 1/24$.*

PROOF. The proof is almost identical to that of Lemma 5.4. Recall that for asynchronous (k_1, k_2) -quorum systems, the properties intersection and non-intersection properties should hold even if we exclude from a quorum Q a set Q_s of size at most $(1 + \mu)p_f m$. For the non-intersection property, not including some replies can only make the intersection smaller, so the proof that with negligible probability, the intersection size is less than $m/3$ carries over to this setting. For the upper bound, the only difference is that we are excluding an additional $(1 + \mu)p_f n$ servers in addition to the already excluded $p_f n$ servers. The expected size of non-corrupt servers in Q that are also in one of the previous k_2 chosen quorums is at least $(1 - \alpha_1 - p_f)m - (1 + \mu)p_f m = (1 - \alpha_1 - (2 + \mu)p_f)m$. We show that the probability that the size is short by a factor $(1 - r)$ of the expected size is negligible, where $r = 1 - ((1 - (\alpha_1 + p_f + (1 + \mu)p_f))/(2/3))$. Similarly to what we did in Lemma 5.4 for the synchronous case, we have $0 < r < 1$ and using Chernoff's lower tail bound, we have for any Q_s such that $|Q_s| \leq mp_f(1 + \mu)$:

$$\Pr_{Q, Q_j \leftarrow Q; j \in J; |J| \geq k_2} [|(Q - Q_s) \cap ((\bigcup_{j \in J} Q_j) - \mathcal{F}_{pr})| < 2m/3] \quad (1)$$

$$= \Pr_{Q, Q_j \leftarrow Q; j \in J; |J| \geq k_2} [|(Q \cap ((\bigcup_{j \in J} Q_j) - \mathcal{F}_{pr} - Q_s))| < (1 - r) \times (1 - (\alpha_1 + p_f))m] \quad (2)$$

$$\leq e^{-r^2/2 \times (1 - (\alpha_1 + p_f))m} \quad (3)$$

\square

VI RANDOM QUORUM SELECTION

The algorithm for random quorum selection is shown in Algorithm 10. The following lemmas establish the relevant properties of the algorithm.

LEMMA 7.1. *If the seller is correct, the quorum of validators is chosen uniformly at random.*

PROOF. If the seller is correct, the validators are chosen using the hash function on different input values, seeded with a high entropy N_s , that w.h.p. have not been previously selected, until m

Algorithm 10 Random selection of a quorum of size m : Seller's's Code (pks)

```

1: function SELECTQUORUM( $\langle F, pkb, pks \rangle, N_s$ )
2:    $T_s = \langle F, pkb, pks \rangle || N_s$  ▷ Seller's transaction identifier
3:    $h = H(T_s)$ 
4:    $Q = \{ \} ; j = 1 ;$ 
5:   while  $|Q| < m$  do
6:     if  $Server(H(h||j)) \notin Q$  then ▷ If server not previously selected,
7:        $Q = Q \cup \{Server(H(h||j))\}$  ▷ add server to quorum
8:        $j = j + 1$ 
9:   return  $Q$ 

```

different validators are chosen. Given that the hash function is modeled as a random function, each of these validators will be chosen randomly. \square

LEMMA 7.2. *If the seller is correct, the identities of the correct validators in the chosen quorum are not known to the adversary or to the buyer.*

PROOF. Under the assumption that the adversary cannot intercept the messages of the seller relating to one transaction (Section 3), the adversary can only learn about the transaction from corrupt servers that are selected as part of the quorum or by guessing the seed. Given that the seed has high entropy, it cannot be guessed other than with negligible probability. The rest of the proof follows directly from the fact that the validators are randomly chosen. All validators that are not under the control of the adversary are equally likely to be chosen in a particular quorum. \square

LEMMA VI.1. *The probability that a randomly selected quorum of size m contains $(1 + \mu)p_f m$ previously corrupt validators, $0 < \mu < 1$ is at most $e^{-\mu^2 \times p_f m / (2 + \mu)}$.*

PROOF. Let p'_f be the fraction of validators that have been previously corrupted. $p'_f < p_f$. Let Q_m be the set of subsets of validators of size m . The probability that a randomly chosen validator is corrupt is at most p'_f . The expected number of validators that are corrupt in a quorum of size m is at most $p'_f m$. By Chernoff's bounds, we have

$$\Pr_{Q \leftarrow Q_m} [|Q \cap \mathcal{F}| \geq (1 + \mu)p'_f m] \leq e^{-\mu^2 \times p'_f m / (2 + \mu)} \leq e^{-\mu^2 \times p_f m / (2 + \mu)}.$$

\square

The following Lemma follows directly from Lemma VI.1.

LEMMA 7.3. *For $0 < \mu < 1$, if a seller choses K quorums, of size m each, at random, then with probability at most $K e^{-\mu^2 \times p_f m / (2 + \mu)}$, every chosen quorums has no more than $(1 + \mu)p_f m$ previously corrupt validators.*