

Fast and Unconditionally Secure Anonymous Channel

Juan Garay
Yahoo Labs
701 First Ave.
Sunnyvale, CA 94089, USA
garay@yahoo-inc.com

Rafail Ostrovsky*
Department of Computer
Science and Department of
Mathematics
UCLA, Los Angeles, CA, USA
rafail@cs.ucla.edu

Clint Givens
Maine School of Science and
Mathematics
Limestone, MN, USA
cgivens@gmail.com

Pavel Raykov
ETH Zurich
Zurich, Switzerland
pavel.raykov@inf.ethz.ch

ABSTRACT

In this paper we focus on sender-anonymous channels (a.k.a. *Dining Cryptographers networks*) and present a construction requiring a very low (constant) number of rounds of interaction while tolerating actively malicious behavior by some of the participants (up to less than half of them). Our construction is unconditionally secure (meaning that no bounds are placed on the computational power of the adversary), makes black-box use of a verifiable secret sharing (VSS) protocol, and is based on a special-purpose secure multiparty computation protocol implementing the method of “throwing darts;” its round complexity is essentially equal to that of the VSS protocol.

In addition, since broadcast *cannot* be simulated in a point-to-point network when a third or more of the participants are corrupt, it is impossible to construct VSS (and, more generally, any other basic multiparty protocol) in this setting without using a “physical broadcast channel,” and a recent line of research has sought to minimize the use of this expensive resource. Our anonymous channel protocol’s reduction to VSS is broadcast-round-preserving, thus making

*Supported in part by NSF grants 0830803, 09165174, 1065276, 1118126 and 1136174, US-Israel BSF grant 2008411, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, and Lockheed-Martin Corporation Research Award. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0392. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'14, July 15–18, 2014, Paris, France.

Copyright 2014 ACM 978-1-4503-2944-6/14/07 ...\$15.00.

<http://dx.doi.org/10.1145/2611462.2611494>.

the fewest (known to date) calls to the broadcast channel while running in an overall constant number of rounds.

Finally, anonymous channels play an important role in the setup phase of an authentication technique known as *pseudosignatures*, which then may be used to simulate *authenticated Byzantine agreement* protocols in the information-theoretic setting. Plugging in our anonymous channel translates into a fast (and broadcast-efficient) pseudosignature construction.

Categories and Subject Descriptors

F.0 [Theory of Computation]: General

Keywords

Anonymous message transmission; DC-nets; pseudosignatures; Byzantine agreement

1. INTRODUCTION

Anonymous (or untraceable) message transmission is a powerful privacy-preserving tool, both in theory and in practice (see, non-exhaustively, [A14]; see also [IKOS06] for a glance at its power). In this paper we consider the seminal technique introduced by Chaum known as *Dining Cryptographers networks* (a.k.a. DC-nets, or, plainly, “anonymous channel”) [Cha88], where the responsibility of providing “anonymization” lies on the same participants providing the input to the computation. This is in contrast to Chaum’s alternative anonymity technique known as *mixnets* [Cha81], which operates by “proxy.”

Now, while the latter has been extensively studied in the literature and serves as the basis of multiple deployed systems (e.g., [GRS99, RR99, BSG00, DDM03, DMS04]), this is considerably less the case for the former. Indeed, only a handful of anonymous-channel constructions exist, both in the information-theoretic (a.k.a. “unconditional,” where no bounds are placed on the computational power of an adversary) and cryptographic settings ([Cha88, BdB89, Wai89, PW96] and, more recently, [Zha11], and [vABH03, GJ04], respectively), all of which perform poorly in terms of running time (specifically, rounds of interaction) when having to cope with disruptive behavior by some of the participants.

This is unsatisfactory, not only because of scalability reasons, but also because anonymous channels play a powerful role in the pre-processing and setup stages of secure multiparty protocols (an application also treated in this paper), and where a speedy execution is of utmost importance.

1.1 Our results

In this work we present an unconditionally secure sender-anonymous channel construction requiring a very low (constant) number of rounds of interaction. Our protocol tolerates actively malicious parties, relies on the standard resources of secure pairwise channels and broadcast, and works in the case the number of such parties $t < n/2$, where n is the total number of participants. For a long time, the best unconditional anonymous channel construction, due to Pfitzmann and Waidner [PW96], required $\Omega(n^2)$ rounds¹. Recently, a constant-round construction was proposed by Zhang [Zha11] with a value however in the hundreds. In contrast, the round complexity of the protocol presented here is essentially equal to that of the used VSS protocol (e.g., 7 rounds in the case of [RB89]). Constant-round anonymous channel protocols were not considered in the cryptographic setting (see Related work below).

As mentioned above, our protocol makes black-box use of a (linear) *verifiable secret sharing* (VSS) [CGMA85] protocol, and is based on a special-purpose secure multiparty computation (MPC) [GMW87, BGW88, CCD88, RB89] protocol implementing the method of “throwing darts” (e.g., [Hag91]). (See next section for a description of the method and comparison to its previous use.) In addition, because broadcast *cannot* be simulated in the information-theoretic setting on a point-to-point network when a third or more of the parties are corrupt [LSP82], it is impossible to construct VSS, and more generally, any basic MPC protocol in this setting without using a “physical broadcast channel” (that is, a black box which securely implements broadcast), or some equivalent addition to the model. Consequently, a recent line of research [GGOR13] has sought to minimize the use of this expensive resource. Our anonymous channel protocol’s reduction to VSS is broadcast-round-preserving, thus making the fewest (known to-date) calls to the broadcast channel (namely, *two*) while running in an overall constant number of rounds. The protocol is presented in Section 3.

As already discussed above, anonymous message transmission is a fundamental privacy-preserving tool, both in theory and in practice. In particular, many-to-one anonymous channels play an important role in the construction of *pseudosignatures*, an information-theoretic authentication technique for multiparty protocols introduced by Pfitzmann and Waidner [PW96], as an extension of and improvement to a scheme by Chaum and Roijackers [CR90]. Suppose there is a setup phase during which the parties enjoy access to a physical broadcast channel (but need not know their future inputs), in a setting where $t \geq n/3$. The parties

¹The anonymous channel protocol in [PW96] in fact tolerates an arbitrary number of corruptions. The protocol uses fault detection and localization: any failed invocation leads to public identification of either a single corrupt player, or a pair of players at least one of whom is corrupt. Thus, even with an honest majority, there are $\Omega(n^2)$ pairs of players with one of them corrupt; therefore the adversary can force the protocol to run for $\Omega(n^2)$ rounds. (Potentially, this could be reduced to $\Omega(n)$ rounds by making use of more recent “player elimination” techniques [HMP00].)

may use such a setup phase to implement pseudosignatures; then, using the pseudosignatures for authentication, they may simulate future invocations of broadcast by running an *authenticated* Byzantine agreement protocol (e.g., [DS83, KK06]), thus avoiding any need for a physical broadcast channel during the main phase of the protocol.

The pseudosignatures of [PW96] rely on a subprotocol implementing a secure, many-to-one anonymous channel. Thus, “plugging in” our constant-round anonymous channel construction translates into a very fast and broadcast-efficient pseudosignature construction. We recall the [PW96] approach and further elaborate on this application in Section 4.

1.2 Related work

We already mentioned the introduction by Chaum of his seminal *Dining Cryptographers networks* (DC-nets) technique for anonymous message publication [Cha88]. DC-nets’ major obstacle, namely, being subject to “jamming” by actively malicious players, was subsequently overcome by Waidner and Pfitzmann and Waidner [Wai89, PW96], by means of a somewhat complicated procedure of setting “traps” during a “slot reservation” phase, resulting in protocols running in $\Omega(n^2)$ rounds, as pointed out above. These constructions (as well as other, cryptography-based constructions [GJ04]—see below) rely on a special algebraic technique, wherein after players have shared pair-wise secret keys (called “DC-keys” or “pads”) plus some other secret information which specifies, for example, a sender’s turn to send a message, the receiver is only able to reconstruct the *sum* of the real message with the shared pads; as the pads cancel each other out, the receiver retrieves the message, but is unable to establish the identity of the sender.

In [Zha11], the author presents a sorting algorithm that enjoys unconditional security in the setting of a malicious adversary corrupting $t < n/2$ players (the setting considered in the current paper). Based on this sorting algorithm, [Zha11] gives a protocol for the “obfuscated shuffle” problem that can be directly translated into an anonymous channel protocol. The resulting protocol employs four functionalities: VSS, comparison, equality testing and multiplication. Its total round complexity is $r_{\text{VSS-share}} + r_{\text{comp}} + r_{\text{eq}} + r_{\text{mult}}$, where $r_{\text{VSS-share}}$, r_{comp} , r_{eq} , r_{mult} are the round complexities of the corresponding functionalities. If all r ’s are constant (e.g., VSS is implemented as in [RB89, GGOR13] and comparison, equality testing and multiplication as in [DFK⁺06]), then the overall round complexity of the protocol is constant. In contrast, the round complexity of the anonymous channel protocol presented in this paper is essentially equal to $r_{\text{VSS-share}}$, which currently is significantly faster than the [Zha11] construction given that comparison and equality testing functionalities require bit decomposition of the shared values, whose round costs are significantly larger than that of a single VSS execution—i.e., 114 rounds for bit decomposition [DFK⁺06] vs. 7 rounds for VSS sharing [RB89].

While in the cryptographic setting the anonymous message transmission functionality can in principle be realized in constant rounds by a generic MPC protocol (e.g., [DI05]), that has not been the case by existing non-generic constructions [vABH03, GJ04], which we now briefly review. The work by Golle and Juels [GJ04], in particular, employs bilinear maps [BF03] and achieves what the authors call “non-interactivity,” in the sense that subsequently to key estab-

lishment, players may publish their messages in a single broadcast round (while detecting and identifying cheating players with high probability). “Collisions,” namely, a common problem which arises in DC-nets in the selection of message positions, even when all the players are honest, are not considered. The suggestion to plainly repeat the execution of the protocol until the messages get delivered, however, is a bit problematic, as it also allows the adversary to introduce additional spurious values; thus, in addition to being unreliable the construction becomes *malleable* (see security properties in Section 2).

While the work of von Ahn *et al.* [vABH03] focuses on the weaker goal of “ k -anonymity” (as opposed to full anonymity, for k a constant smaller than but otherwise not related to n), in the sense that the adversary is able to learn something about the origin (as well as destination) of a particular message, but cannot narrow down its guess to a set of less than k participants, technique-wise it is related to ours, as it also follows the “dart-throwing” method mentioned above. In some detail, in order to anonymously publish a message (or send it to a specific receiver) the method consists of each player randomly choosing a few indices in a long vector, and that player’s message appears at these indices. (Some collisions between different messages, of course, may occur.) Then, the long vector is revealed (to the receiver) with an outcome corresponding to the multiset of messages appearing sufficiently many times in the vector. Provided the players can correctly simulate the dart-throwing in such a way that each player’s choice of indices remains private from the receiver, and so that only the long vector with the messages is revealed, the channel will be anonymous. In our construction, a careful choice of parameters guarantees the channel’s security properties (cf. Section 2) except with negligible probability. In contrast, the construction in [vABH03] guarantees success (specifically, the Reliability property²) with probability $1/2$ only, bringing us back to the challenge pointed out above of achieving $(1 - \epsilon)$ -reliability, for negligible ϵ , just by repetition without sacrificing non-malleability.

An approach similar to the “dart-throwing” technique has also been recently used in [AIKW13] for a different problem, namely, the randomized encoding of the subset function. There the authors encrypt each of the sparse vectors under a separate key. Then, if an additive homomorphic encryption is employed, it is guaranteed that, for any set of indices I , the partial sum of vectors’ encryptions with indices in I is an encryption of the corresponding partial sum of plain vectors (with indices in I) under the corresponding partial sum of the keys (with indices in I).

Regarding pseudosignatures, an interesting alternative construction to [PW96]’s was proposed by Shikata *et al.* [SHZI02], relying on the evaluation of random low-degree multivariate polynomials in the setup phase. In [BTHR07] it is shown how to compute these polynomials using a special type of MPC protocol, which can also be made in a constant number of rounds. A salient difference of the approach is that pseudosignatures obtained this way can only be applied to messages drawn from the underlying field where the MPC computation takes place, whereas the [PW96] approach is domain-independent, as the generated setup can be used to sign messages from domains that are not known during the setup phase. On the flip side, the [SHZI02]-[BTHR07] real-

ization is more communication-efficient than [PW96]’s with our anonymous channel, thus presenting a tradeoff between versatility and speed on one hand and communication efficiency on the other. (Further details in Section 4.)

In that sense, we finally remark that since the focus of this paper is on the feasibility of *fast* (i.e., constant-round) anonymous channels, we forgo explicit treatment of communication complexity. We do however note that the protocols described herein can be compiled via generic techniques (see, for example, [BFO12]) into more communication-efficient versions.

2. MODEL, DEFINITIONS AND TOOLS

We consider a complete, synchronous network of n players P_1, \dots, P_n who are pairwise connected by secure (private and authenticated) channels, and who additionally have access to a broadcast channel. Some of these players, up to $t < n/2$, are corrupted by a centralized adversary \mathcal{A} with *unbounded computing power*. The adversary is *active*, directing players under his control to deviate from the protocol in arbitrary ways, and may be *static* or *adaptive*, depending on whether he chooses which players to corrupt prior to the start of the protocol execution or on the fly, respectively; the type of adversary tolerated by our protocol will in turn depend on the security properties of the underlying VSS protocol³.

The network computation evolves as a series of rounds. In a given round, honest players’ messages depend only on information available to them from prior rounds; \mathcal{A} , however, is *rushing*, and receives all messages (and broadcasts) sent by honest players before deciding on the messages (and broadcasts) of corrupted players. Sometimes we refer to \mathcal{A} thus defined as a *t-adversary*. Messages are drawn from and computations are carried out in a fixed finite field \mathbb{F} , $|\mathbb{F}| > n$. We consider statistical security (i.e., protocols are subject to some [negligibly small] probability of error), as perfect security is possible only when $t < n/3$ [CCD88, DDWY93]⁴, and let $\kappa \in \mathbb{N}$ denote the error parameter, $\kappa \geq 2n$. As usual, we say a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *negligible* if for all polynomials p , $f \leq 1/p$.

For the sake of simplicity, we adopt in all our protocols the usual convention that whenever a party fails to send an expected message, or sends a syntactically incorrect message, it is replaced with some default message. Thus we do not deal separately with the case of missing or improper messages.

2.1 Many-to-one anonymous channel

As mentioned above, *Dining Cryptographers networks* (or DC-nets) are a privacy-preserving primitive introduced by Chaum [Cha81, Cha88] for anonymous message publication. Here we formally define a *secure, many-to-one anonymous channel* which allows n players to simultaneously transmit messages to any single player P^* . Let X denote the multiset of honest senders’ messages and Y denote the multiset output by P^* . Then the following conditions are satisfied in the presence of a t -adversary \mathcal{A} :

³Our constant-round, most broadcast-efficient construction, based on [GGOR13], will only be statically secure, for example.

⁴This result is mentioned also in [BGW88, RB89]; [CCD88] give an informal argument and [DDWY93] a formal proof.

²Called “Robustness” in [vABH03].

ANONYMITY: Even if P^* is corrupt, \mathcal{A} learns no more than X (and in particular gains no information on which messages came from which honest players).

PRIVACY: If P^* is honest, then \mathcal{A} learns no information on X .

RELIABILITY: If P^* is honest, then $X \subseteq Y$.

NON-MALLEABILITY: If P^* is honest, then $|Y| \leq n$ and the probability distribution of $Y \setminus X$ is independent from X .

Some of these properties (in particular, Reliability—cf. our comment on statistical security above) may hold except with a negligible error probability in the security parameter. In addition, these security properties are required to hold under parallel composition of the anonymous channel. We now turn to the definition of the main building block used by the anonymous channel construction.

2.2 Verifiable secret sharing

An (n, t) -verifiable secret sharing scheme [CGMA85] is a pair of protocols (VSS-Share, VSS-Rec) for a set of players $\mathcal{P} = \{P_1, \dots, P_n\}$, one of whom, the dealer D , holds input $s \in \mathbb{F}$. It must satisfy the following guarantees in the presence of a t -adversary \mathcal{A} :

COMMITMENT: With high probability (w.h.p.), at the end of VSS-Share there exists a fixed value $s^* \in \mathbb{F}$, defined by the joint view of the honest parties, such that all honest parties will output s^* in VSS-Rec. If D is honest, then $s^* = s$. On the other hand, if D is dishonest, then the value of s^* is efficiently computable given the sharing-phase views (messages sent and received) of all correct parties.⁵

PRIVACY: If D is honest, then w.h.p. prior to VSS-Rec the adversary gains no information on s (i.e., his view is statistically independent of s).

We say that the parties *verifiably share* a secret s if each (honest) party maintains some state such that, when the honest parties invoke VSS-Rec on that joint state, they will reconstruct the value s (w.h.p.). Clearly, if a dealer D has just completed VSS-Share with effective input s , then the parties verifiably share s .

LINEARITY: If the parties verifiably share secrets $\{s^{(k)}\}$, then they also (without further interaction) verifiably share any (public) linear combination of the secrets.

In this paper we are interested in *fast* (i.e., constant-round) anonymous channel constructions. Linear, constant-round VSS protocols for the $t < n/2$ regime are presented in [RB89, Rab94, CDD⁺01, GGOR13]⁶; the protocol presented in [GGOR13], in particular, only requires the use of the broadcast channel twice in the sharing phase and none in reconstruction, although is not as round-efficient⁷.

⁵Combined with Privacy, this latter property guarantees malicious dealers’ inputs are independent of honest dealers’ inputs when several copies of VSS are run in parallel.

⁶The VSS protocols presented in [KPC10] are also constant-round, but not (apparently) linear.

⁷21 rounds in [GGOR13] vs. 9 in [Rab94], for example.

3. FAST AND UNCONDITIONAL ANONYMOUS CHANNEL

We first give a high-level description of the protocol. We instantiate the anonymous channel using a special-purpose MPC protocol. As mentioned in Section 1, our special-purpose MPC protocol can be characterized by the method of “throwing darts” [Hag91]. To wit, each player randomly chooses a few indices in a long vector, and that player’s message appears at these indices. (Some collisions between different messages, of course, may occur.) Then, the long vector is revealed to the receiver who delivers the multiset of messages appearing sufficiently many times in the vector. Provided the players can correctly simulate the dart-throwing in such a way that each player’s choice of indices remains private from the receiver, and so that the long vector with the messages is revealed only to the receiver, the channel will be anonymous. A careful choice of parameters guarantees the channel’s security properties (cf. Section 2) except with negligible probability.

Security against maliciously active adversaries requires that dishonest players *commit* to and are not able to change their original inputs, and provide a proof of good behavior throughout the computation. For the former, our construction makes use of a (constant-round, broadcast-efficient) VSS protocol, the distributed analogue of a commitment functionality. The latter is achieved through the traditional “cut-and-choose” approach of preparing many instances and then collectively deciding on which ones to open while maintaining the privacy and anonymity of the construction. However, while in the cryptographic setting a variety of zero-knowledge proofs exists for this type of task, the design of such a mechanism in our (unconditional) setting requires some additional care.

We now present the protocol, AnonChan, which implements an anonymous channel for $t < n/2$ using secure point-to-point channels and black-box access to a linear VSS protocol. Although the VSS itself must rely on physical broadcast when $t \geq n/3$, our construction uses no additional broadcast rounds beyond those required by the calls to VSS.

First of all we ensure that the messages the honest players send are all unique and non-zero. This is achieved by players appending non-zero random tags to them, which P^* strips off upon receipt—this decreases reliability by at most a negligible amount (the probability that two honest tags collide).

We now explain how the “throwing darts” paradigm described above is implemented. Each player P_i commits, by invoking the VSS sharing phase in many parallel executions, to a sufficiently long and sparse vector $\mathbf{v}^{(i)}$ in which the few nonzero entries are set equal to P_i ’s input. If all players do so correctly, then the sum of the vectors will have relatively few collisions among nonzero entries—in particular, it will preserve enough instances of each input value with high probability. Then the players use the linearity of the VSS scheme to locally compute shares of the sum of the $\mathbf{v}^{(i)}$ ’s, and send the resulting shares to the receiver P^* (privately). P^* then simulates the VSS reconstruction function internally to recover the sum, which contains at least one copy of each player’s input (along with some additional garbage arising from collisions, which we can tolerate).

Of course, dishonest players may not prepare their $\mathbf{v}^{(i)}$ ’s properly—if a cheater were to commit to (say) a vector full

of random entries, then including it in the sum would destroy all information about honest players' inputs. Therefore we require that, after committing to his vector, each player must *prove* that it is indeed sparse, and that nonzero entries are equal. A simple cut-and-choose proof suffices: The prover P_i prepares many additional vectors $\mathbf{w}_j^{(i)}$, $1 \leq j \leq \kappa$, in the same manner as the original $\mathbf{v}^{(i)}$, except that he re-randomizes the locations of nonzero entries. He also commits, for each $\mathbf{w}_j^{(i)}$, to both the list of locations of nonzero entries, and to a permutation on vector entries which sends nonzero locations of $\mathbf{v}^{(i)}$ to nonzero locations of $\mathbf{w}_j^{(i)}$.

The parties then jointly generate a random challenge r : For each $\mathbf{w}_j^{(i)}$, P_i is challenged to reveal either (a) the permutation, in which case the parties reconstruct the *difference* of the permuted $\mathbf{v}^{(i)}$ and $\mathbf{w}_j^{(i)}$, to verify it is the zero vector; or (b) the nonzero entries of $\mathbf{w}_j^{(i)}$, in which case the parties reconstruct the alleged *zero* entries to verify they are all zero, and reconstruct the *differences* of nonzero entries, which should be zero to verify that all nonzero entries are equal. Vectors which fail this check are excluded, and with overwhelming probability all accepted vectors are indeed sparse, with nonzero entries equal. Finally, the players randomly permute each of the accepted vectors in order to ensure that their nonzero entries appear at random indices. (If a random permutation is not applied, then the non-zero entries of accepted vectors belonging to malicious players appear at the indices chosen by the adversary and hence are not random.)

The protocol parameters are the player set $\mathcal{P} = \{P_1, \dots, P_n\}$, a designated receiver $P^* \in \mathcal{P}$, and error parameter κ . All computations are done over a finite field $\mathbb{F} = GF(2^\kappa)$, where we assume that $\kappa \geq 2n$. The length of vectors ℓ and their sparseness d used in the protocol are polynomial in n and κ (both ℓ and d are given explicitly in the proof). Protocol **AnonChan** is shown in Figure 1. Constant number of rounds can easily be verified by inspection, and the fact that the invoked VSS protocol is also constant-round.

THEOREM 1. *Let (VSS-Share, VSS-Rec) be a linear verifiable secret sharing scheme which fails with probability at most $2^{-\Omega(\kappa)}$ in the presence of an unbounded adversary corrupting up to $t < n/2$ parties. Then protocol **AnonChan** implements a many-to-one anonymous channel, except with probability $2^{-\Omega(\kappa)}$.*

PROOF. We now show each of the security properties of a many-to-one anonymous channel (cf. Section 2):

ANONYMITY. Without loss of generality consider an adversary corrupting P^* . By the privacy of VSS, \mathcal{A} 's view of the VSS sharing phases at the end of step 1 reveals only negligible information on the honest players' inputs. The random value r , reconstructed in step 2, is independent of the honest players' inputs.

Each honest player P_i will have correctly shared sparse vectors $\mathbf{v}^{(i)}, \mathbf{w}_j^{(i)}$. Consequently, in step 3 \mathcal{A} can predict beforehand, the outcome of all reconstructions associated with honest players' vectors (namely, they will be zero), and so \mathcal{A} learns negligible information on honest inputs. Also, w.h.p. all honest players will be included in **PASS**.

In step 4, P^* receives shares of $\mathbf{v} = \sum_{P_i \in \text{PASS}} g_i(\mathbf{v}^{(i)})$. Because \mathcal{A} knows the value of $\mathbf{v}^{(j)}$ for each dishonest player P_j in **PASS**, he may subtract these from \mathbf{v} and get $\mathbf{v}_{\text{honest}} =$

$\sum_{P_i \in \text{honest}} g_i(\mathbf{v}^{(i)})$. Which is to say, \mathcal{A} learns no more information from \mathbf{v} than he would from learning $\mathbf{v}_{\text{honest}}$ directly. But since each honest P_i chose his d indices of $\mathbf{v}^{(i)}$ independently at random, the distribution of $\mathbf{v}_{\text{honest}}$ is a function only of the *set* of honest players' (unique) values. Hence \mathcal{A} gains no information from seeing $\mathbf{v}_{\text{honest}}$ other than what he could learn from seeing the set of honest players' values.

PRIVACY. Assume P^* is honest. As above, by the privacy of VSS, \mathcal{A} learns at most a negligible amount of information about honest players' inputs from the VSS sharing phases in step 1; no information from the reconstruction of r in step 2, which is independent of the x_i 's; and a negligible amount from the step 3 proof of sparseness, the outcome of which he can predict with all but negligible probability. Then statistical privacy is ensured since all communication in step 4 is on the private channels to P^* .

RELIABILITY. We must show that Y output by an honest P^* at step 4 contains a copy of each honest input x_i , with overwhelming probability. As noted in the proof of Anonymity, each honest player is included in **PASS** with high probability. Call a vector *proper* if it is d -sparse, and all its non-zero entries are equal (if not, we call it *improper*). First we prove the following claim showing that all vectors used in step 4 are proper.

CLAIM 1. *If a dishonest P_i shares in step 1 an improper vector $\mathbf{v}^{(i)}$, he will be disqualified w.h.p. at step 3.*

PROOF. P_i must commit to $\mathbf{v}^{(i)}$ and to all $\mathbf{w}_j^{(i)}$'s in step 1, before \mathcal{A} has any information on the value of the challenge r , which will be uniformly distributed over \mathbb{F} .

For each j : If $\mathbf{w}_j^{(i)}$ is proper, then no permutation will map the entries of an improper $\mathbf{v}^{(i)}$ onto matching entries of $\mathbf{w}_j^{(i)}$; hence the vector $\mathbf{u} := \pi_j^{(i)}(\mathbf{v}^{(i)}) - \mathbf{w}_j^{(i)}$ will be nonzero and P_i will be disqualified w.h.p. provided $b_j = 0$. Alternatively, if $\mathbf{w}_j^{(i)}$ is improper, then either the alleged zero entries will actually contain some non-zero value, or the alleged *non-zero* entries will contain unequal values (or both). In either case, opening the alleged zero entries, and the differences of alleged non-zero entries, will reveal an error and P_i will be disqualified w.h.p. provided $b_j = 1$.⁸

Since each b_j is uniformly random, P_i will be disqualified with probability $\geq 1 - 2^{-\Omega(\kappa)}$, and with probability $\geq 1 - t \cdot 2^{-\Omega(\kappa)} = 1 - 2^{-\Omega(\kappa)}$, every dishonest P_i who commits to an improper $\mathbf{v}^{(i)}$ will be disqualified. \square

Now we know that w.h.p. all the vectors $\mathbf{v}^{(i)}$ used in step 4 are proper. Since P^* is honest we have that each $\mathbf{v}^{(i)}$ used in step 4 is randomly permuted with g_i . Let I_i denote the non-zero indices of permuted $\mathbf{v}^{(i)}$. In the following claim we estimate how many indices in I_1, \dots, I_n may coincide.

CLAIM 2. *Let I_1, \dots, I_n be random subsets of size d chosen from $[\ell]$. Let $X_{ij} = |I_i \cap I_j|$. Then for any $C \geq 0$ holds:*

$$\Pr \left[\sum_{i \neq j} X_{ij} \geq n^2(d^2/\ell + C \cdot d) \right] \leq n^2 \exp(-C^2 \cdot d)$$

⁸The reason P_i will only be disqualified "with high probability" is that, with negligible probability, the adversary may succeed in causing incorrect values to be reconstructed for one or more of the VSS sharings.

Protocol AnonChan(\mathcal{P}, P^*, κ)

Input: Each $P_i \in \mathcal{P}$ has as input a message $x_i \in \mathbb{F}$.

Output: P^* outputs a multiset Y .

0. Each P_i creates a random non-zero κ -bit tag $a_i \in \mathbb{F}$.
1. Each P_i constructs a random d -sparse vector $\mathbf{v}^{(i)} \in (\mathbb{F} \times \mathbb{F})^\ell$ whose d nonzero entries are (x_i, a_i) . At the same time, P_i constructs κ random permutations $\pi_j^{(i)} : [\ell] \rightarrow [\ell]$. For each $\pi_j^{(i)}$ let $\mathbf{w}_j^{(i)}$ be a vector $\mathbf{v}^{(i)}$ permuted with $\pi_j^{(i)}$, i.e., $\mathbf{w}_j^{(i)}[k] = \mathbf{v}^{(i)}[\pi_j^{(i)}(k)]$ for all $k \in [\ell]$.
 P_i invokes **VSS-Share** $O(\ell\kappa)$ times in parallel to share each of the following values:
 - Each coordinate of $\mathbf{v}^{(i)}$ and of the $\mathbf{w}_j^{(i)}$'s;
 - each of the permutations $\pi_j^{(i)}$;
 - for each $\mathbf{w}_j^{(i)}$, its list of nonzero indices;
 - a random element $r^{(i)} \in \mathbb{F}$.
 At the same time P^* generates n random permutations $g_i : [\ell] \rightarrow [\ell]$ and shares them using **VSS-Share**.
2. The players invoke **VSS-Rec** to reconstruct the sum $r := \sum r^{(i)}$, and interpret it as a random bit-string of length κ .
3. For each bit $b_j \in r$ (in parallel):
 - $b_j = 0$:
 1. The players invoke n instances of **VSS-Rec** to reconstruct the permutations $\pi_j^{(i)}$, for all P_i . (If the result is not a valid permutation, P_i is disqualified.)
 2. For each $\mathbf{w}_j^{(i)}$, the players invoke ℓ instances of **VSS-Rec** to reconstruct the vector $\mathbf{u} := \pi_j^{(i)}(\mathbf{v}^{(i)}) - \mathbf{w}_j^{(i)}$. If $\mathbf{u} \neq (0, 0)^\ell$, P_i is disqualified. ($\pi_j^{(i)}$ acts on a vector by permuting its components.)
 - $b_j = 1$:
 1. The players invoke n instances of **VSS-Rec** to reconstruct the list of nonzero indices for $\mathbf{w}_j^{(i)}$, for all P_i . (If the result is not a valid list of d distinct indices in $[\ell]$, P_i is disqualified.)
 2. For each $\mathbf{w}_j^{(i)}$, invoke $\ell - d$ instances of **VSS-Rec** to reconstruct the values at (alleged) zero-indices of $\mathbf{w}_j^{(i)}$. If any are actually nonzero, P_i is disqualified. Also invoke $d - 1$ instances of **VSS-Rec** to reconstruct consecutive *differences* of (alleged) nonzero entries of $\mathbf{w}_j^{(i)}$. If any such differences are nonzero, P_i is disqualified.
4. Let **PASS** denote the set of players who are not disqualified. Players use **VSS-Rec** to reconstruct P^* 's permutations g_1, \dots, g_n globally. Each player locally computes his shares of the vector $\mathbf{v} := \sum_{P_i \in \text{PASS}} g_i(\mathbf{v}^{(i)})$. He sends the shares, over private channel, to the receiver P^* . P^* uses the received shares to internally simulate **VSS-Rec** and recover the vector \mathbf{v} . Let T denote the set of nonzero pairs which appear $\geq d/2$ times in \mathbf{v} . Let Y be an empty multiset; then, for each $(x_i, a_i) \in T$, include x_i in Y . P^* outputs Y .

Figure 1: A secure anonymous channel protocol using black-box VSS.

PROOF. Consider X_{ij} for any fixed i and j . There are d fixed indices in I_i and ℓ indices in total. The intersection of I_j and these fixed indices follows a hypergeometric distribution. We have that $\mathbb{E}[X_{ij}] = d^2/\ell$. We use a tail bound on the hypergeometric distribution appearing in [Chv79, Ska13] to conclude that, for any $C \geq 0$:

$$\Pr[X_{ij} \geq d^2/\ell + C \cdot d] \leq \exp(-C^2 \cdot d).$$

We trivially have that

$$\Pr\left[\sum_{i \neq j} X_{ij} \geq n^2(d^2/\ell + C \cdot d)\right] \leq \Pr\left[\bigvee_{i \neq j} (X_{ij} \geq d^2/\ell + C \cdot d)\right].$$

Using the union bound we upper bound the latter probability with $n^2 \exp(-C^2 \cdot d)$. \square

Now we choose ℓ, C and d such that the following holds: $n^2(d^2/\ell + C \cdot d) = d/2$ and $C^2 \cdot d \in \Omega(\kappa)$. This will ensure that w.h.p. in total $< d/2$ indices coincide in I_1, \dots, I_n . This requirement is satisfied for $C = 1/(4n^2)$, $d = n^4\kappa$ and $\ell = 4n^6\kappa$.

Hence, among d indices I_i of any honest P_i at least $d/2$ are untouched in \mathbf{v} implying that $(x_i, a_i) \in T$ at Step 4.

Finally, since a_i are chosen uniformly at random from $\mathbb{F} \setminus \{0\}$ the probability that a_i collides with any other a_j is at most $2^{-\Omega(\kappa)}$. So P^* will include a copy of x_i sent by P_i in Y .

NON-MALLEABILITY. By independence of inputs for the VSS, all the values which dishonest players commit to in step 1 are independent of those which honest players commit to in step 1.

Consider any value $x \in Y$. Since $x \in Y$ this implies that for some $a \in \mathbb{F}$ it holds that $(x, a) \in T$ at step 4. We know that (x, a) is included in T only if it appears at least $d/2$ times in \mathbf{v} . Due to Claim 1 and Claim 2 this implies that (x, a) must have been a non-zero value in some vector $\mathbf{v}^{(i)}$ with $P_i \in \text{PASS}$. The vector $\mathbf{v}^{(i)}$ originates either from an honest player or from a dishonest. In the latter case since adversary learns no information about inputs of the honest players, the value (x, a) is generated by \mathcal{A} independently of all honest players' messages. Hence, Y can be split into two subsets: X —multiset with honest players' messages, and $Y \setminus X$ —multiset with the messages independent of X . Finally, since each value in Y originates from some $\mathbf{v}^{(i)}$ we have that $|Y| \leq n$. \square

4. APPLICATIONS: PSEUDOSIGNATURES

As mentioned in Section 1, *pseudosignatures* [PW96] are an information-theoretic authentication technique for multiparty protocols. They require a setup phase during which the parties enjoy access to a physical broadcast channel (recall that we are in the $t \geq n/3$, information-theoretic regime, where broadcast cannot be simulated on a point-to-point network); during this phase the parties need not know their future inputs. After the setup phase, using the pseudosignatures for authentication, the parties are able to simulate future invocations of broadcast by running an authenticated Byzantine agreement protocol [DS83, KK06], thus avoiding any need for a physical broadcast channel during the main phase of the protocol.

One does pay a price for not relying on cryptographic assumptions—in the case of pseudosignatures, the price is *limited transferability*: The integrity of a pseudosignature “degrades” each time it is passed from one party to another, so that it only remains valid for an *a priori* bounded number of transfers⁹. In the case of deterministic Byzantine agreement, for example, $O(t)$ -transferability would suffice, given the number of rounds lower bound [LSP82] and protocols matching it (e.g., [DS83]).

The pseudosignatures of [PW96] rely on a subprotocol implementing a many-to-one anonymous channel, as defined above. Assuming such a channel, the pseudosignature scheme is roughly as follows. (For simplicity, let us consider a one-time signature.) Each player chooses a large number of random keys, which will function as information-theoretic message authentication codes. The players invoke the anonymous channel many times in parallel, each sending one key per invocation to the signer P^* . Hence, for each invocation of the channel, P^* receives an associated *signature block* containing $n - 1$ anonymous keys. Now, in order to (pseudo-)sign a message M , P^* simply signs it—i.e., computes the message authentication code on the message—using every authentication key, in every block; these individual signatures are referred to as *minisignatures*, and the collection of all of them, arranged in blocks, is P^* 's *pseudosignature* on M .

Verification is more involved. The *first* verifier V_1 (i.e., the player to whom P^* originally sends M with pseudosignature) accepts the signature provided that, in *every* signature block, one of the minisignatures matches the key V_1 sent for that block. The second verifier V_2 accepts the signature provided that, in *most* of the signature blocks, one of the minisignatures matches the key V_2 sent for that block. The third verifier accepts the signature provided a *fair number* of minisignatures check out, and so on, where each new verifier has a lower acceptance threshold than the previous one.

The rationale for the increasingly tolerant verifiers is the fear that a cheating signer P^* , even though he does not know whose keys are whose in any given block, could (for example) correctly sign M with every key except for half the keys in a certain block. There is a good chance then that V_1 will find a correct minisignature in every block, but will pass it on to V_2 who will *not* find a correct minisignature in the half-signed block. If V_1 accepts the signature while V_2 rejects, P^* has successfully broken the signature scheme.

By having the parties simply invoke protocol AnonChan (Figure 1) for each P_i , $1 \leq i \leq n$, acting as receiver for many sessions in parallel, the duration of the setup phase is reduced from $\Omega(n^2)$ to constant—the number of rounds that AnonChan takes. Further, by using the broadcast-efficient VSS protocol in [GGOR13], the use of the physical broadcast channel is reduced to 2 (also from $\Omega(n^2)$ in [PW96]). On the other hand, we should note that the pseudosignature setup in [PW96] works for any $t < n$, while ours is limited to $t < n/2$. This would not impose additional limitations, as we consider an important application of pseudosignatures to simulate broadcast in unconditional honest-majority MPC protocols.

To conclude, we mentioned in Section 1 that an alternative pseudosignature construction has been proposed by Shikata *et al.* [SHZI02], relying on the evaluation of random low-degree multivariate polynomials in the setup phase. In [BTHR07] it is shown how to compute these polynomials using a special type of MPC protocol. This MPC computation uses a generic VSS scheme, and thus it can also be made in a constant number of rounds. (In [BTHR07] this is not the case as the goal is communication efficiency.) In addition to pseudosignatures being confined to messages drawn from the underlying field, as pointed out in Section 1, the MPC protocol, in contrast to ours, requires to compute multiplication of the shared values, which translates into additional interaction. On the flip side, the type of computation enjoys lower communication costs, thus presenting a versatility and speed versus communication efficiency tradeoff.

5. REFERENCES

- [A14] Selected papers in anonymity. The Free Haven Project. <http://freehaven.net/anonbib/full/date.html>, 2014.
- [AIKW13] B. Applebaum, Y. Ishai, E. Kushilevitz, and B. Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. In R. Canetti and J. A. Garay, editors, *CRYPTO (2)*, volume 8043 of *Lecture Notes in Computer Science*, pages 166–184. Springer, 2013.
- [BdB89] J. N. Bos and B. den Boer. Detection of disrupters in the dc protocol. In J.-J. Quisquater and J. Vandewalle, editors, *EUROCRYPT*, volume 434 of *Lecture Notes in Computer Science*, pages 320–327. Springer, 1989.
- [BF03] D. Boneh and M. Franklin. Identity-based encryption from the weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003.
- [BFO12] E. Ben-Sasson, S. Fehr, and R. Ostrovsky. Near-linear unconditionally-secure multiparty computation with a dishonest minority. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 663–680. Springer, 2012.
- [BGW88] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proc. 20th Annual ACM Symposium of the Theory of Computation*, pages 1–10, May 1988.

⁹Another price is finite, fixed-in-advance player set!

- [BIB89] J. Bar-Ilan and D. Beaver. Non-cryptographic fault-tolerant computing in constant number of rounds of interaction. In P. Rudnicki, editor, *PODC*, pages 201–209. ACM, 1989.
- [BSG00] P. Boucher, A. Shostack, and I. Goldberg. Freedom systems 2.0 architecture. White paper, Zero Knowledge Systems, Inc., December 2000.
- [BTHR07] Z. Beerliová-Trubíniová, M. Hirt, and M. Riser. Efficient byzantine agreement with faulty minority. In K. Kurosawa, editor, *ASIACRYPT*, volume 4833 of *Lecture Notes in Computer Science*, pages 393–409. Springer, 2007.
- [CCD88] D. Chaum, C. Crépeau, and I. Damgård. Multiparty unconditionally secure protocols. In *Proceedings 20th Annual Symposium on Theory of Computing, STOC*. Association for Computing Machinery, May 1988.
- [CDD⁺01] R. Cramer, I. Damgård, S. Dziembowski, M. Hirt, and T. Rabin. Efficient multiparty computations secure against an adaptive adversary. In *Advances in Cryptology—EUROCRYPT’01*, Lecture Notes in Computer Science. Springer Verlag, May 2001.
- [CGMA85] B. Chor, S. Goldwasser, S. Micali, and B. Awerbuch. Verifiable secret sharing and achieving simultaneity in the presence of faults. In *Proceedings of Twenty Sixth IEEE Symposium in Foundations of Computer Science*, pages 383–395, 1985.
- [Cha81] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–88, 1981.
- [Cha88] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *J. Cryptology*, 1(1):65–75, 1988.
- [Chv79] V. Chvátal. The tail of the hypergeometric distribution. *Discrete Mathematics*, 25(3):285 – 287, 1979.
- [CR90] D. Chaum and S. Roijakkers. Unconditionally secure digital signatures. In A. Menezes and S. Vanstone, editors, *CRYPTO*, volume 537 of *Lecture Notes in Computer Science*, pages 206–214. Springer, 1990.
- [DDM03] G. Danezis, R. Dingledine, and N. Mathewson. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, pages 2–15, May 2003.
- [DDWY93] D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *Journal of ACM*, 1(40):17–47, 1993.
- [DFK⁺06] I. Damgård, M. Fitzi, E. Kiltz, J. B. Nielsen, and T. Toft. Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In S. Halevi and T. Rabin, editors, *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 285–304. Springer, 2006.
- [DI05] I. Damgård and Y. Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In V. Shoup, editor, *CRYPTO*, volume 3621 of *Lecture Notes in Computer Science*, pages 378–394. Springer, 2005.
- [DMS04] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13*, SSYM’04, pages 21–21, Berkeley, CA, USA, 2004. USENIX Association.
- [DS83] D. Dolev and H. Strong. Authenticated algorithms for Byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983.
- [GGOR13] J. Garay, C. Givens, R. Ostrovsky, and P. Raykov. Broadcast (and round) efficient verifiable secret sharing. In C. Padro, editor, *ICITS*, volume 8317 of *Lecture Notes in Computer Science*, pages 200–219. Springer, 2013.
- [GJ04] P. Golle and A. Juels. Dining cryptographers revisited. In C. Cachin and J. Camenisch, editors, *EUROCRYPT*, volume 3027 of *Lecture Notes in Computer Science*, pages 456–473. Springer, 2004.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *Proc. 19th Annual ACM Symposium on Theory of Computation*, pages 218–229, May 1987.
- [GRS99] D. M. Goldschlag, M. G. Reed, and P. F. Syverson. Onion routing. *Commun. ACM*, 42(2):39–41, 1999.
- [Hag91] T. Hagerup. Fast parallel generation of random permutations. In J. Albert, B. Monien, and M. Rodríguez-Artalejo, editors, *ICALP*, volume 510 of *Lecture Notes in Computer Science*, pages 405–416. Springer, 1991.
- [HMP00] M. Hirt, U. Maurer, and B. Przydatek. Efficient secure multi-party computation. In *Advances in Cryptology, ASIACRYPT*, Lecture Notes in Computer Science, 2000.
- [IKOS06] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. Cryptography from anonymity. In *FOCS*, pages 239–248. IEEE Computer Society, 2006.
- [KK06] J. Katz and C. Koo. On expected constant-round protocols for Byzantine agreement. In *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 445–462. Springer, 2006.
- [KPC10] R. Kumaresan, A. Patra, and C. Pandu Rangan. The round complexity of verifiable secret sharing: The statistical case. In M. Abe, editor, *ASIACRYPT*, volume 6477 of *Lecture Notes in Computer Science*, pages 431–447. Springer, 2010.
- [LSP82] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, pages 382–401, July 1982.

- [PW96] B. Pfitzmann and M. Waidner. Information-theoretic pseudosignatures and byzantine agreement for $t \geq n/3$. Technical Report RZ 2882 (#90830), IBM Research, 1996.
- [Rab94] T. Rabin. Robust sharing of secrets when the dealer is honest or cheating. *J. ACM*, 41(6):1089–1109, 1994.
- [RB89] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *Proc. 21st ACM Symposium on the Theory of Computing*, pages 73–85, 1989.
- [RR99] M. K. Reiter and A. D. Rubin. Anonymous web transactions with crowds. *Commun. ACM*, 42(2):32–38, 1999.
- [SHZI02] J. Shikata, G. Hanaoka, Y. Zheng, and H. Imai. Security notions for unconditionally secure signature schemes. In L. R. Knudsen, editor, *EUROCRYPT*, volume 2332 of *Lecture Notes in Computer Science*, pages 434–449. Springer, 2002.
- [Ska13] M. Skala. Hypergeometric tail inequalities: ending the insanity. *ArXiv e-prints*, November 2013.
- [vABH03] L. von Ahn, A. Bortz, and N. Hopper. K-anonymous message transmission. In *Proceedings of the 10th ACM Conference on Computer and Communications Security, CCS '03*, pages 122–130, New York, NY, USA, 2003. ACM.
- [Wai89] M. Waidner. Unconditional sender and recipient untraceability in spite of active attacks. In J.-J. Quisquater and J. Vandewalle, editors, *EUROCRYPT*, volume 434 of *Lecture Notes in Computer Science*, pages 302–319. Springer, 1989.
- [Zha11] B. Zhang. Generic constant-round oblivious sorting algorithm for mpc. In X. Boyen and X. Chen, editors, *ProvSec*, volume 6980 of *Lecture Notes in Computer Science*, pages 240–256. Springer, 2011.