

Cryptography from Anonymity

Yuval Ishai*

Eyal Kushilevitz[†]

Rafail Ostrovsky[‡]

Amit Sahai[§]

Abstract

There is a vast body of work on implementing anonymous communication. In this paper, we study the possibility of using anonymous communication as a building block, and show that one can leverage on anonymity in a variety of cryptographic contexts. Our results go in two directions.

- **Feasibility.** *We show that anonymous communication over insecure channels can be used to implement unconditionally secure point-to-point channels, broadcast, and general multi-party protocols that remain unconditionally secure as long as less than half of the players are maliciously corrupted.*
- **Efficiency.** *We show that anonymous channels can yield substantial efficiency improvements for several natural secure computation tasks. In particular, we present the first solution to the problem of private information retrieval (PIR) which can handle multiple users while being close to optimal with respect to both communication and computation.*

1 Introduction

There are many scenarios in which anonymous communication can be implemented at a low cost, either via physical means (e.g. in wireless networks, or small wired networks) or by means of special-purpose protocols. Indeed, a lot of systems work has been done on implementing anonymous communication (see [1, 12, 47, 7] and refer-

ences therein). Anonymizing web browsers and anonymous email accounts are already widely available. In this work, we ask the question: If anonymity is already out there, can we harness its power for other purposes? To what extent can anonymity be used as a *building block* for obtaining better solutions to other important cryptographic tasks? We elaborate on this question below.

Anonymity vs. privacy. Anonymous communication allows users to send messages to each other without revealing their identity. However, in contrast to popular misconception, anonymity is far from answering all concerns of “privacy”.¹ Conceptually, anonymity is aimed at hiding *who* performs some action, whereas full privacy requires additionally hiding *what* actions are being performed. In the context of distributed computation, anonymity allows hiding which users hold which local inputs, whereas privacy requires hiding all information about the inputs except what follows from the outputs. In a sense, the relation between anonymity and privacy is analogous to the relation between unpredictability and indistinguishability: while the former notions of security might be sufficient for some applications, they are generally considered inadequate; in particular, they are vulnerable to attacks that exploit a-priori information about the secrets.

The aim of the current work is to study the extent to which anonymity can be useful as a *primitive*. Can the gap between hiding the *Who* and the *What* be closed at a small additional cost? Can anonymity be exploited in other cryptographic contexts, beyond those that involve privacy?

A toy example. As a simple motivating example, consider the following scenario. Two players, A and B , wish to agree on an unconditionally secret random bit (a “key”). Their only means of communicating is by posting anonymous messages on a *public* internet bulletin board. (In this example, we assume that the board operator, as well as the users, are “honest but curious”.) The key agreement protocol proceeds as follows. Each player $P \in \{A, B\}$ independently picks a random 50-bit integer r_P , and posts the message (“AB”, r_P) on the board. The common bit is taken to be 0 if $r_A > r_B$ and 1 if $r_A < r_B$. (In the unlikely event of a tie, the protocol aborts.) Note that since each player P knows its integer r_P they can both compute the (same)

¹The term “privacy” has different interpretations. Our usage of this term below follows the common terminology in the literature on cryptographic protocols.

*Computer Science Department, Technion. Research supported by grant 36/03 from the Israel Science Foundation and grant 2004361 from the U.S.-Israel Binational Science Foundation. E-mail: yuval@cs.technion.ac.il

[†]Computer Science Department, Technion. Research supported by grant 36/03 from the Israel Science Foundation and grant 2002354 from the U.S.-Israel Binational Science Foundation. E-mail: eyalk@cs.technion.ac.il

[‡]Computer Science Department and Department of Mathematics, UCLA. Supported in part by IBM Faculty Award, Intel equipment grant, NSF Cybertrust grant No. 0430254, Xerox Award and grant 2002354 from the U.S.-Israel Binational Science Foundation. E-mail: rafail@cs.ucla.edu

[§]Computer Science Department, UCLA. Research supported in part by an Alfred P. Sloan Foundation Research Fellowship, an Intel equipment grant, NSF ITR/Cybertrust grants 0456717 and 0627781, and grant 2004361 from the U.S.-Israel Binational Science Foundation. Email: sahai@cs.ucla.edu

common bit, whereas other users (as well as the board operator) cannot distinguish r_A from r_B and thus learn nothing about the common bit. Of course, the 1-bit key can now be used by A and B to communicate a bit with unconditional secrecy using the public bulletin board.

1.1 Our Contribution

We demonstrate the usefulness of the “cryptography from anonymity” paradigm in two settings: (1) Establishing *feasibility* results for traditional cryptographic tasks unconditionally, based solely on the assumption of public anonymous channels. (2) Showing that anonymous channels can lead to much more *efficient* solutions to several cryptographic problems. We now provide a detailed account of both types of results.

1.1.1 Feasibility results using anonymity

We start by studying which tasks can be implemented with *unconditional* security based on anonymous communication. To this end, we consider the following weak model of anonymity over *public* channels. In each round each player can send a message to a chosen destination. The adversary can learn, for every player in the network (including uncorrupted players), the *multiset* of all messages received by that player. The adversary does *not* learn the identity of the senders, except when the sender itself is corrupted. In addition to such anonymous channels, we also assume that the players can communicate via authenticated but public point-to-point channels.

One of the challenges that need to be faced when attempting to exploit such a network is the fact that anonymity can also serve as a *shelter* for malicious players. For instance, if the protocol instructs only some strict subset of the players to (anonymously) send messages to the same receiver, malicious players outside this set can interfere by sending their own messages. Note that one cannot make a direct use of authentication to distinguish “legitimate” messages from illegitimate ones, as this would violate anonymity.

In the above model, we show how to realize the following primitives:

1. Secure point-to-point communication with unconditional security against an arbitrary number of malicious players. (This protocol strengthens the simple key agreement protocol described above in that it is resilient against “Denial of Service” attacks mounted by malicious players.)
2. General multi-party protocols with unconditional security against any minority of malicious players. (Note that such protocols cannot be based on secure point-to-point channels alone even for simple functionalities such as broadcast.)

As an intermediate step towards establishing the above results, we construct a stronger form of *private* anonymous channels out of the basic public anonymous channels defined above. Such private anonymous channels allow the adversary to learn only messages sent or received by corrupted players. The equivalence between private and public anonymity is of independent interest, since there is no obvious construction of private anonymity from public anonymity even if one is additionally given secure (but non-anonymous) point-to-point channels.

1.1.2 Efficiency improvements based on anonymity

We now turn the attention to the question of *efficiency*, attempting to identify natural cryptographic tasks for which anonymity can give rise to substantial efficiency gains. In contrast to the feasibility results discussed above, here we do not restrict ourselves to unconditional results. In particular, we would like to improve over the best known solutions under *any* cryptographic assumption. A key observation, that underlies our protocols in this setting, is that *local randomization* of inputs, via secret-sharing, when combined with the *global mixing* of the shares, provided by anonymity, allows us to keep the inputs private and, at the same time, allows us to carry out some useful computations on the inputs. We elaborate below.

“Split and Mix” approach. Consider a scenario in which several clients want to access or query a central server without revealing their sensitive data to the server. For instance, the clients may want the server to compute some global function of their joint queries (e.g., average salary of employees), or alternatively to respond to each query separately (as in the case of retrieving data from a central database). As discussed above, anonymity alone does not provide a good solution to this problem. While the mixing effect achieved by anonymity eliminates some information about the query made by each client, most information remains. Our key idea is to boost the effect of anonymity by using local randomization as a *catalyst*. (Indeed, in our analysis it will be useful to view the local randomization as a seed to a randomness extractor, and the partial randomization provided by anonymity as an imperfect source.) More concretely, our approach is to first have each client locally *split* its query into few randomized sub-queries, e.g. via the use of secret-sharing, and then let the clients *mix* all their sub-queries by anonymously sending them to the server. (Here we assume that all sub-queries are mixed together, so that the server cannot tell whether two sub-queries were sent by the same client.) The hope is that the mixed sub-queries can *totally* eliminate unnecessary information about the queries, either statistically or computationally. Moreover, the splitting should be done in a way that allows carrying out the desired computation on the original queries based on the mixed sub-queries. We stress that neither mixing nor splitting alone can provide an adequate level of privacy; but as it will turn out, their combination is surprisingly

powerful. We demonstrate the usefulness of this “split and mix” approach in several contexts, described below.

Non-interactive private statistics. We consider the case where two or more clients hold m -bit integers and wish to reveal to a server the *sum* of their inputs (and nothing else) by simultaneously sending anonymous messages to the server. (Note that without anonymity, it is impossible to solve this problem in a completely non-interactive way as we require, regardless of efficiency.) A simple but inefficient solution is to let each client i , holding an integer x_i , anonymously send x_i distinct dummy messages to the server. The server can now compute the sum of all inputs by simply counting the number of messages it received. This simple solution provides perfect privacy, but does not scale well with the bit length m . Towards a more efficient solution, we use the split and mix approach in the following natural way. Each client locally splits its input into $O(m)$ shares via the use of *additive* secret-sharing, and anonymously sends its shares to the server. The server can now recover the sum of the inputs by adding up all the shares it received. We show that with this choice of parameters, the mixed messages reveal only a negligible amount of information about the inputs (other than their sum). This basic integer summation protocol can be used as a building block for computing other kinds of useful statistics on distributed data. We also show an application of this protocol in a general context of two-party computation, where each client wants to privately compute a function of its own input and the server’s input.

Optimal amortized PIR. In the problem of Private Information Retrieval (PIR) [15, 37] the server holds a (large) database of size n , and each client wants to retrieve a specific item from this database while hiding what it is after. This can be trivially done by having the server communicate the entire database to each client, but this solution is prohibitively expensive. In recent years, there has been a significant body of work on improving the communication complexity of PIR, either in the above single-server scenario [37, 10, 38, 23] or using multiple servers [15, 3]. Given the low (and essentially optimal) communication complexity of the best known PIR protocols, the efficiency bottleneck shifts to the *local computation* performed by the server. Indeed, it is not hard to see that even in the case of a single query, the server must read every bit of the database in order for full privacy to be maintained. Thus, the best one could hope for is to amortize this cost over multiple queries.

The question of amortizing the computational cost of PIR has been previously considered in [4, 32]. However, all previous solutions to this problem either require multiple servers (and fail to protect against colluding servers) or only allow to amortize the cost of several simultaneous queries that *originate from the same client*.² A remaining open problem in this area is to obtain solutions to PIR that

²In [4] it was demonstrated that the computational cost of PIR can be *slightly* amortized even in the case of queries that originate from different clients.

are close to optimal with respect to *both* communication and computation even in the case where queries originate from different clients. We suggest a solution to this problem using (two-way) anonymous communication, and assuming that the queries are made simultaneously by the clients. Our solution applies the “split and mix” technique as follows. First, each client randomizes its query into a small number of sub-queries by simulating an appropriate *multi-server* PIR protocol (with privacy threshold equal to the security parameter). Then, all sub-queries are anonymously sent to the server, who responds to each sub-query separately without knowing which client it originates from. Each client can recover the answer to its query from the server’s answers to its sub-queries as in the underlying multi-server PIR protocol. The computation in this protocol can be amortized by choosing the parameters so that the space of all possible sub-queries is of polynomial size. This enables precomputing the answers to all possible sub-queries.

The security of our PIR protocol relies on an intractability assumption related to the hardness of reconstructing *noisy* low-degree curves in a low-dimensional space. A similar assumption for the two-dimensional case was introduced by Naor and Pinkas [41]. Roughly speaking, the original assumption from [41] asserts that noisy two-dimensional curves cannot be reconstructed in a much better way than using the Guruswami-Sudan list decoding algorithm [27]. Our generalized assumption asserts that, in the low-dimensional case, one cannot do much better than the Coppersmith-Sudan algorithm [16] (which, in a sense, extends [27] to the multi-dimensional case). We note that this assumption does not seem to be affected by the recent progress in the field of list-decoding [43, 26, 25].

On relaxing the anonymity assumption. While we assume for simplicity that the network provides *perfect* anonymity, this assumption can be relaxed. Most of our protocols only require “sufficient uncertainty” about the origins of messages to maintain their full security, at a modest additional cost. Thus, our approach is quite insensitive to imperfections of the underlying network, and can be applied even in more realistic scenarios where only some crude form of anonymity is available. See [33] for further details.

1.2 Related Work

A variant of the toy example presented above (for key agreement using anonymity) was suggested by Alpern and Schneider [2]. A similar idea was previously used by Winkler [51] for establishing secure channels in the game of Bridge. Our work, in contrast, achieves key agreement in the presence of *malicious* parties, a problem that was posed and left open in [2]. The related problem of obtaining key agreement from *recipient anonymity*, which hides the identity of the receiver rather than that of the sender, was considered in [44]. Anonymous communication has also been exploited in the context of practically oriented applications such as voting [13] and electronic cash [49].

Our work can be cast in the setting of investigating secure reductions between different multiparty functionalities: An anonymous channel can be modelled by such a functionality, and we are investigating what other functionalities can be realized using this functionality (and how efficiently). This area has a rich history (see [35, 36, 21, 28, 5, 22, 40] and references therein). However, most of the work in this area has been restricted to the two-party setting, and known results for the multi-party setting are not general enough to apply to the case of the anonymity functionality. Moreover, relatively little attention has been paid to the *efficiency* of such reductions.

Pfitzmann and Waidner [45] use a variant of private anonymous communication as an intermediate step for obtaining highly resilient broadcast protocols (in a model that allows broadcast in a precomputation phase). We rely on their protocol as a step in obtaining our feasibility result for MPC with honest majority. We stress that the protocol of [45] requires the use of *private* anonymity, and thus does not directly imply an implementation of broadcast in our basic model of non-private anonymity. For self-containment, we also present a simpler alternative to the use of the protocol of [45] in our context.

Finally, our approach is also reminiscent of work on data privacy through data *perturbation* (cf. [17, 14, 50]). These works examine privacy through “blending in with the crowd” [14], obtained via the perturbation of data revealed to the adversary. In our work, we also examine how privacy can be achieved by blending in with the crowd, but via *mixing* rather than perturbation.

Organization. In Section 2, we provide some necessary background and definitions. In Section 3, we apply anonymity for obtaining new feasibility results and, in Section 4, we apply anonymity for improving the efficiency of cryptographic protocols. For lack of space, some of the material was omitted from this extended abstract and can be found in the full version [33].

2 Preliminaries

We denote by $[n]$ the set $\{1, 2, \dots, n\}$ and by $\binom{[n]}{k}$ the collection of all subsets of n of size k . We use \log to denote \log_2 , the logarithm to the base 2. We denote by $SD(X, Y)$ the statistical distance between probability distributions X, Y , and by $H_\infty(X)$ the min-entropy of X defined by $H_\infty(X) = \min_x (-\log \Pr[X = x])$, where the minimum is taken over all x in the support of X .

Network model. We consider a network of N players that are connected via authenticated but *non-private* point-to-point channels. (The availability of authenticated channels is a standard assumption in the context of multi-party computation, and is necessary to prevent *denial of service* attacks where even a single malicious player can prevent the entire system from functioning.) In addition, we allow the players to communicate via anonymous but non-

private channels: at each round, each player P_i sends a single anonymous message to some player P_j of its choice, including itself.³ The adversary learns the contents of *all* messages exchanged between the players, including the destination of each message but not its source. This basic type of anonymous communication can be naturally captured by an N -party functionality that we denote by *Anon*. Alternatively, one could assume that the adversary only learns messages received by corrupted players; we refer to such a primitive as a *private-anonymous channel* and denote the corresponding functionality by *PrivAnon*. We will show, in Section 3.1, that the two variants are in fact equivalent. We stress that, in both cases, the adversary cannot learn the sources of messages that originate from uncorrupted players. See [33] for formal definitions of the functionalities *Anon* and *PrivAnon*. Our definitions allow the adversary to be *rushing*; namely, to choose the messages it sends depending on the messages received by corrupted players at the same round.

Some of our applications will rely on *two-way anonymous communication*. An invocation of this primitive consists of two rounds, allowing each player P_i to anonymously send a message to any player P_j and to receive a reply to its message. For each such message-reply pair sent from player P_i to P_j and back, the adversary learns the identity of the destination player j , but not the identity of i ; it can also tell that the second message is a reply to the first. Most physical or algorithmic implementations of anonymity achieve this type of two-way communication almost for free (cf. [12, 7, 48]). Finally, in some of our applications, we will consider scenarios where the players are partitioned into two or more *clients* and a single *server*, so that each client only needs to interact with the server and not with other clients. The above definitions apply to such a setting as well.

Secure reductions. Our protocols can be viewed as reductions between cryptographic primitives; namely, they show how to implement a certain primitive (or functionality) g given a black-box access to a functionality f , where f is typically an anonymity functionality. By a t -secure reduction from g to f , we refer by default to a *statistically t -secure* protocol for g in the so-called f -hybrid model (i.e., in a model where players have access to a trusted oracle computing f). For a formal definition of “statistically t -secure protocols”, see [11, 24]. For simplicity, we consider *non-adaptive* adversaries, who choose the set of corrupted players in advance.

3 Feasibility Results Using Anonymity

In this section, we present unconditionally secure implementations of several cryptographic primitives based on anonymous communication over public channels.

³This can be extended by allowing each player to send up to λ messages at each round, for some parameter λ . Note that without such a bound the adversary may “flood” the network with anonymous messages.

3.1 From Public to Private Anonymity

We start by showing how to realize the private anonymity functionality PrivAnon using the basic (non-private) anonymity functionality Anon . This will serve as a stepping stone towards implementing other primitives. The above goal is nontrivial even if we were additionally given private (non-anonymous) point-to-point channels. Indeed, there is no obvious way to combine the advantages of non-private but anonymous channels and private but non-anonymous channels. Instead, we suggest the following direct implementation of PrivAnon based on Anon .

Assume for now that there are only 3 players, A , B (senders) and R (receiver); see Remark 3.1 below for the generalization to the N -party case. We wish to construct a protocol allowing A and B to send messages to R with the following properties: (1) If A and B are honest then their anonymity is preserved. (2) If the adversary corrupts only one sender, then it cannot violate the privacy or the correct delivery of the message sent by the other sender, nor can it correlate its own message with the other message.

Below, we write $AE_K(m)$ to denote a (statistically secure) one-time authenticated encryption of the message m using the key K . Such an encryption can be decrypted and authenticated using the secret key K . It can be implemented by using a one-time pad for encrypting and an unconditionally-secure MAC for authenticating. We let k denote a statistical security parameter. The protocol proceeds as follows:

1. Repeat $3k$ times sequentially (each is referred to as a “session”):
 - (a) Each of A , B and R sends a random k -bit number to R , using anonymous (non-private) channels.
 - (b) R considers the numbers received in the previous step, ignoring repetitions, and chooses 2 out of these numbers, including its own. R sends these 2 numbers in lexicographic order to both players A and B via authenticated channels. The order of these 2 numbers defines a (secret) bit between R and either A or B according to who’s number is chosen. (With overwhelming probability, the honest players choose distinct numbers. The adversary, being rushing, may duplicate numbers selected by honest players, but does not know who’s numbers it duplicate.)
2. Each of A and B sends to R , via non-private anonymous channels, a list of the first k session numbers in which they obtained shared bits. (Note that, with overwhelming probability, each honest sender obtained at least k shared bits.) This results in the definition of two k -bit secret keys K_A and K_B . The receiver, R , knows both keys, but does not know which of them belongs to A and which belongs to B (i.e., from R ’s perspective they are two keys $\{K_1, K_2\} = \{K_A, K_B\}$).

3. To send private anonymous messages m_A and m_B to R , the two players A and B send $AE_{K_A}(m_A)$ and $AE_{K_B}(m_B)$, respectively, via the non-private anonymous channel to R . The receiver R authenticates and decrypts each message using the keys K_1 and K_2 (this allows identifying the key corresponding to the message).

Remark 3.1 The following changes should be applied to the above protocol, when dealing with an N -player network (a receiver R and $N - 1$ potential senders). We increase the number of “sessions” to $2Nk$ which guarantees that the numbers chosen by each sender appear in the pairs selected by R at least k times (also, k is sufficiently large so that in each session all the numbers chosen by honest players are distinct). This allows each sender to send a list of k session numbers for which it knows the corresponding secret bit.

Theorem 3.2 *The above protocol defines a statistically N -secure reduction from the private anonymity functionality PrivAnon to the non-private anonymity functionality Anon .*

3.2 Private Point-to-Point Channels

In this section, we show how to use anonymous channels to realize private point-to-point channels. In light of the previous subsection, it suffices to implement key agreement using *private* anonymous channels. Recall that the simple key agreement protocol described in the Introduction (as well as a similar protocol from [2]) allows A and B to agree on a secret random key r by posting messages on an anonymous bulletin board. This protocol assumes that the bulletin board operator as well as the other users in the system are honest-but-curious.

We now describe a key-agreement protocol that works in our standard model, namely where the players A , B are just two players in a network of N , possibly malicious, players. The main difficulty in utilizing anonymity in this case is that when one of the players needs to send an anonymous message, corrupted players may attack the protocol by also sending messages over the anonymous channel. Our protocol prevents this attack, even if all other $N - 2$ players are malicious, via a simple reduction to private anonymity (i.e., we assume the availability of private anonymous channels, as constructed above).

1. A sends a random k -bit key, K_A , to B via a private anonymous channel. [Note that malicious players may also send to B anonymous k -bit messages.]
2. B sends to A the list of all keys received in step (1), via a private anonymous channel. [Note that malicious players may also send to A such lists.]
3. A finds the single list L which includes her original key K_A from all lists received in step (2) (if there is no such list then A aborts), and sends over a public authenticated channel the position of K_A in this list.

Theorem 3.3 *The above protocol defines a statistically N -secure reduction from the key agreement functionality to the private anonymity functionality PrivAnon.*

In the full version [33] we discuss an alternative means for obtaining key agreement in our setting via a reduction to the problem of key agreement using a deck of cards [20].

3.3 General Multiparty Computation

Using a result of Pfitzmann and Waidner [45], it is possible to use *private* anonymity for implementing broadcast, which together with private point-to-point channels can be used to securely compute any N -party functionality with resilience $t < N/2$ [46]. In the full version we provide more details [33] as well as an alternative to the Pfitzmann-Waidner construction. We also discuss limitations on what can be achieved using anonymity, and in particular conclude that the bound $t < N/2$ is tight. Thus, the anonymity functionality is *nontrivial* in the sense that it allows key agreement, but on the other hand it is not *complete* for all N -party functionalities with respect to N -secure reductions.

4 Efficiency Improvements Using Anonymity

In this section we consider different scenarios in which anonymous communication can yield *efficiency* improvements over the best known solutions (even ones that rely on cryptographic assumptions).

4.1 Non-interactive Private Statistics

We show how $n \geq 2$ clients can privately compute statistics (such as mean, standard deviation, correlations) on their combined inputs by each sending few anonymous messages to a central server. Our protocols only require one-way anonymous communication and are private with respect to an adversary corrupting the server along with an arbitrary number of clients.⁴ Note that it is impossible to obtain such non-interactive protocols in the standard model, even if one settles for computational privacy. (It is possible to solve this problem in the non-interactive model of [18]; however, such a solution requires setup assumptions and provides a weaker security guarantee.)

Our basic building block is a protocol for integer summation. We assume that each client P_i holds an integer x_i , where $0 \leq x_i < M$. We want to design a protocol in which each client sends a small number of anonymous messages to the server, from which the server can recover the sum of all inputs without learning additional information about the inputs. This basic building block for privately computing the sum immediately allows privacy-preserving computation of

⁴We provide no guarantee of correctness in the presence of malicious clients. However, in most applications of the kind considered here malicious clients can cause nearly as much damage also in an idealized implementation involving a trusted party.

the mean of a distributed set of data, and can also be applied to privately compute “suites” of statistics such as: (1) both the mean and variance of a set of numbers, and (2) the means of and covariance between two or more sets of numbers (where each player holds corresponding elements from the sets).⁵ The sum protocol can also be used to efficiently compute randomized linear *sketches* of the data that reveal approximate statistics (e.g., an approximate histogram).⁶

Our goal is to obtain a (statistically) private protocol in which the communication complexity is essentially optimal: the total number of bits sent by each player depends only logarithmically on M .

The protocol SUM. We present a protocol for adding n inputs in a finite group G of size L . (The above integer summation problem reduces to addition over $G = \mathbb{Z}_L$, where $L = nM$.) To compute the sum of the inputs, each player *additively shares* its input into k shares in G (where k will be specified later) and sends each share to \mathcal{S} in a separate anonymous message. The server can recover $\sum x_i$ by adding up (in G) the kn messages it received.

Analysis. We now analyze the parameters for which mixing additive shares hides the values of the shared secrets.⁷ We start with the case of $n = 2$ players and consider the experiment of running the above protocol with uniformly chosen inputs in G . Let (X, Y) denote the players’ random inputs and V denote the mixed shares received by \mathcal{S} . Let $V(x, y)$ denote the distribution of V conditioned on $X = x, Y = y$ and $V(x)$ denote the distribution of V conditioned on $X = x$. Finally, let U be a random variable uniformly distributed in G , independently of V .

Lemma 4.1 *Suppose $\log \binom{2k}{k} > \ell + \sigma$. Then, $SD((V, X), (V, U)) \leq 2^{-\Omega(\sigma)}$.*

Proof: For $a \in G^{2k}$ and $\pi \in \binom{[2k]}{k}$ let $h_a(\pi) = \sum_{i \in \pi} a_i$. Note that h_a defines a family of pairwise independent hash functions from $\binom{[2k]}{k}$ to G . (Pairwise independence follows from the fact that each a_i is an independent element of the group.) Also note that (V, X) is distributed identically to $(V, h_V(\Pi))$, where Π is the uniform distribution over all sets in $\binom{[2k]}{k}$ independently of V . This follows from the fact that, by symmetry, every possible k -subset of shares is equally likely to coincide with the shares of X . Finally, the Leftover Hash Lemma [30] guarantees that $SD((V, h_V(\Pi)), (V, U)) \leq 2^{-\Omega(H_\infty(\Pi) - \ell)} = 2^{-\Omega(\log \binom{2k}{k} - \ell)}$ from which the lemma follows. ■

⁵It is important that a suite of statistics are being computed – for instance, in example (1) above, we cannot use the sum protocol to privately compute *only* the variance, without revealing the mean. However, it is most often desirable to compute both the mean and variance together.

⁶In general, the approximate output together with the randomness used to generate the sketch may reveal a few bits of additional information that do not follow from the exact output (see [19]). However, in most applications of sketching, this privacy loss is either insignificant or non-existent.

⁷A somewhat simpler variant of the problem we consider here was implicitly considered in the context of constructing pseudorandom generators based on subset sum [29].

Lemma 4.2 Suppose $SD((V, X), (V, U)) \leq \epsilon$. Then, for all $x, y, x', y' \in G$ such that $x + y = x' + y'$, we have

$$SD(V(x, y), V(x', y')) \leq 2|G|^2 \cdot \epsilon.$$

Proof: If $SD((V, X), (V, U)) \leq \epsilon$ then, by Markov's inequality, for every $x \in G$ we have $SD(V(x), V) \leq |G| \cdot \epsilon$. By the triangle inequality, for every x, x' we have $SD(V(x), V(x')) \leq 2\epsilon|G|$.

To complete the proof we show that a δ -distinguisher D between $V(x, y)$ and $V(x', y')$, where $x + y = x' + y'$, can be turned into a $\delta/|G|$ -distinguisher D' between $V(x)$ and $V(x')$. Such a distinguisher can be implemented as follows. Let $z = x + y (= x' + y')$. Given a challenge v (a vector of $2k$ mixed shares), D' checks whether the shares add up to z and if so invokes D on v ; otherwise it outputs 0. ■

From these two lemmas, we immediately conclude that the protocol privately computes the sum for $n = 2$ players, with the appropriate setting of k :

Theorem 4.3 Let $k = 1.5\ell + \sigma$. Then, protocol SUM privately computes the sum of $n = 2$ inputs in a group G , where $|G| < 2^\ell$, with statistical error $2^{-\Omega(\sigma)}$.

The following theorem extends the analysis to the case of $n > 2$ clients. Its proof, appearing in [33], applies a reduction to the case $n = 2$ using the fact any n -tuple of elements from G that add up to 0 can be written as the sum of at most $n - 1$ such n -tuples of Hamming weight 2.

Theorem 4.4 Let $k = 1.5\ell + \sigma + \log n$. Then, protocol SUM privately computes the sum of n inputs in a group G , where $|G| < 2^\ell$, with statistical error $2^{-\Omega(\sigma)}$.

Secure Two-Party Computation. In [33] we describe an application of the summation protocol for realizing general secure two-party computation between the server and each client, albeit in a rather weak security model.

4.2 Private Information Retrieval

In this section, we use anonymous channels to obtain a PIR protocol which allows a server to handle queries that may originate from many different clients using a nearly optimal amount of communication and computation.

The model. We consider a system with a single server, holding a database $x \in \{0, 1\}^m$, and several (typically many) clients. Each client holds a selection index $i \in [m]$ and wishes to learn x_i without revealing i to the server. The protocol requires only a single round of queries and answers. Each client can send several (simultaneous) anonymous queries to the server and receive a separate answer for each query, via the two-way anonymity functionality. We stress that in our protocol the clients do not need to interact with each other. Our protocol will provide the following, somewhat unconventional, security guarantee. An adversary corrupting the server and a subset of the clients

will be unable to learn the inputs of the remaining clients, in the sense that different choices for these inputs induce computationally indistinguishable views, provided that the number of *uncorrupted* clients exceeds some given threshold. (More precisely, it will suffice that the total number of queries originating from uncorrupted clients exceeds this threshold.) The value of the threshold will depend on the database size and the security parameter, but not on the number of clients. Thus, the fraction of corrupted parties that can be securely tolerated by the protocol tends to 1 as the number of clients grows.

Overview of construction. We take a t -server information-theoretic PIR protocol in which the client's privacy is protected against collusions of k servers. (In our typical choice of parameters, we let k serve as the security parameter and $t = O(k \cdot m^\epsilon)$.) The t queries sent by the client in this protocol can be viewed as points on a degree- k curve in a low-dimensional space. Any k of these t points jointly reveal nothing about the client's selection i , whereas any $k + 1$ of them completely determine i . A natural approach that comes to mind is to (computationally) hide the curve encoding i by adding random noise. As it turns out, the required amount of noise is very large – it has to be at least of the order of magnitude of m , the database size, in order to defeat an attack by Coppersmith and Sudan [16].⁸ Thus, the approach is entirely useless in case of a single client accessing the database. The key observation is that the same amount of noise would suffice to hide an arbitrarily large number of curves, possibly originating from different clients. Thus, the use of anonymity allows to amortize the required noise over multiple clients. When the number of uncorrupted clients is sufficiently large, the amount of noise each client needs to contribute is small.

The original polynomial reconstruction (PR) intractability assumption, introduced by Naor and Pinkas [41], asserts roughly the following. For an appropriate choice of parameters, the output of the following experiment keeps a secret field element $s \in F$ semantically secure with respect to a security parameter k : (1) pick a random polynomial $p(\cdot)$ of degree $\leq k$ such that $p(0) = s$; (2) pick t distinct evaluation points $a_1, \dots, a_t \in F$ and n random noise coordinates $r_1, \dots, r_n \in F$; (3) output the good points $(a_j, p(a_j))$ along with noise points (r_j, b_j) in a random order (where each b_j is random and independent of all r_i).

The Guruswami-Sudan list decoding algorithm [27] implies that the above assumption does not hold when $t > \sqrt{(n+t)k}$. Thus, the assumption becomes plausible only when the amount of noise is higher, say when $n \gg t^2/k$. We rely on the following multi-dimensional variant of the above assumption: the secret s is replaced by a vector of c field elements $s = (s_1, \dots, s_c)$ and the polynomial p by a

⁸An attempt to base PIR on a stronger version of our intractability assumption was made in [34]. This assumption was broken by the Coppersmith-Sudan algorithm (see also [8]). Our protocol relies on a much more conservative choice of parameters, that is not known to imply PIR in the standard model.

$(c + 1)$ -dimensional curve, namely by a vector of c polynomials $p = (p_1, \dots, p_c)$. The above experiment can then be generalized in a natural way to the multi-dimensional case. Formally, the assumption is defined as follows.

Definition 4.5 (Noisy Curve Reconstruction (CR) Assumption) Let k denote a degree parameter, which will also serve as a security parameter. Given functions $F(k)$ (field), $c(k)$ (dimension), $t(k)$ (points on curve), and $n(k)$ (noise), we say that the CR assumption holds with parameters (F, c, t, n) if the output of the following experiment keeps a secret $s \in F(k)^{c(k)}$ semantically secure (with respect to security parameter k):

- Pick random polynomials $p_1(\cdot), \dots, p_c(\cdot)$, s.t. each p_h is of degree $\leq k$ and $p(0) \stackrel{\text{def}}{=} (p_1(0), \dots, p_c(0)) = s$;
- Pick t distinct evaluation points $a_1, \dots, a_t \in F \setminus \{0\}$ and n random noise coordinates $r_1, \dots, r_n \in F \setminus \{0\}$;
- Output the good points $p(a_j)$ along with random noise points $(b_j^1, \dots, b_j^c) \in_R F^c$, in a random order.

Towards relating the CR assumption to known attacks, it is convenient to consider an augmented (and “more adversarial”) experiment which outputs the evaluation point a_j along with each c -tuple $p(a_j)$ and a random element of F along with each noise point b_j . Clearly, if the augmented CR assumption (i.e., the CR assumption with respect to the augmented experiment) holds then it also holds with respect to the original experiment. The algorithm from [16] breaks the augmented CR assumption when $t > ((n + t)k^c)^{1/(c+1)} + k + 1$. Thus, when $t = o(nk^c)^{1/(c+1)}$ or equivalently $n = \omega(t \cdot (t/k)^c)$ the augmented assumption (let alone the original one) remains plausible. We stress that the assumption does not seem to be affected by the recent progress in the field of list-decoding [43, 26, 25].

Our protocol uses the following choice of parameters. Let $c > 1$ be a constant. (The amortized complexity per client will be of the order of $n^{1/c}$.) We view the entries of a database $x \in \{0, 1\}^m$ as the coefficients of a c -variate polynomial q_x of total degree at most $d = O(m^{1/c})$ over a field F , where the size of F will be specified later. This allows to associate with each selection index $i \in [m]$ a point $z_i \in F^c$ such that $q_x(z_i) = x_i$ (see, e.g., [15]).

The protocol. Each client, holding selection index i , picks a random degree- k curve $p = (p_1, \dots, p_c)$ such that $p(0) = z_i$, as well as $t = kd + 1$ random distinct evaluation points $a_j \in F \setminus \{0\}$. It anonymously sends to the server the t queries v_j where $v_j = p(a_j) \in F^c$. In addition, the client anonymously sends a number of random noise points of the form $b_j \in_R F^c$, so that the total number of noise points sent by all clients is at least n . (The security of the protocol will be guaranteed as long as the total number of noise points sent by *uncorrupted* clients is at least n .) The server replies to each query with an answer $s_j = q_x(v_j)$. (If all values of q_x were precomputed, this is done via a table lookup.) Finally, the client can recover x_i by interpolating the degree- k univariate polynomial defined by the points (a_j, s_j) . For

this interpolation to be possible, we need $|F| > t + 1$, though a larger F is desirable for enhancing the security.⁹

Privacy. The following lemma guarantees that if n points of noise are sufficient to (computationally) hide the selection of a single client, then this is also the case for an arbitrary polynomial number of clients.

Lemma 4.6 Let k denote a security parameter let $u(k)$ be a polynomial. Let $A(k) = (A_1(k), \dots, A_{u(k)}(k))$ be a distribution ensemble, where $A(k)$ is a sequence of $u(k)$ independent distributions over multisets of elements from a domain $D(k)$. Let $B(k) = (B_1(k), \dots, B_{u(k)}(k))$ be another distribution ensemble as above, and let $R(k)$ be a random multiset of $n(k)$ elements from $D(k)$. Moreover, suppose that for every index sequence $j(k)$, $1 \leq j(k) \leq u(k)$, we have $A_j \cup R \stackrel{\approx}{\approx} B_j \cup R$. (Here $\stackrel{\approx}{\approx}$ denotes computational indistinguishability with respect to polynomial-size circuits, and the dependence of all parameters on k is implicit in the notation.) Then, $A_1 \cup \dots \cup A_u \cup R \stackrel{\approx}{\approx} B_1 \cup \dots \cup B_u \cup R$.

Proof: Suppose the contrary. By a hybrid argument, there is a sequence $j(k)$ such that $A_1 \cup \dots \cup A_{j-1} \cup B_j \cup \dots \cup B_u \cup R$ can be distinguished from $A_1 \cup \dots \cup A_j \cup B_{j+1} \cup \dots \cup B_u \cup R$ with non-negligible advantage. The corresponding distinguisher T can be used to get a distinguisher between $A_j \cup R$ and $B_j \cup R$: given a multiset S , take the union of S with a sample from $A_1 \cup \dots \cup A_{j-1} \cup B_{j+1} \cup \dots \cup B_u$, and invoke T on the result. ■

Note that the CR assumption guarantees that the queries of a *single* client, when combined with n points of noise, keep the client’s selection computationally private. Thus, Lemma 4.6 establishes the privacy of the protocol for an arbitrary number of clients, with the same amount of noise as that required for the privacy of a single client:

Theorem 4.7 If the CR assumption (Definition 4.5) holds with parameters $(F(k), c(k), t(k), n(k))$, then the above anonymous PIR protocol remains computationally private for an arbitrary (polynomial) number of clients, as long as the total amount of noise contributed by uncorrupted clients is at least $n(k)$.

Parameters. Recall that we set c to be a constant and $t = O(km^{1/c})$. As noted above, a good choice of the noise parameter for the CR assumption is $n = \omega(t \cdot (t/k)^c) = \omega(k \cdot m^{1+1/c})$. Thus, the total amount of noise is comparable to the database size. Finally, we argue that the query space is polynomial, so that the answers to all queries can be precomputed by the server. Recall that we require that $|F| = \Omega(t) = \Omega(km^{1/c})$ but, as discussed above, it is safer to avoid many collisions and thus let $|F|$ be larger than n . Either way, the query space $|F|^c$ is polynomial in m .

Using the above choice of parameters, we get:

⁹The problem with letting $|F| \approx t$ is that the good points are likely to share the same X -coordinates with many noise points. In such a case, PR-type assumptions are susceptible to lattice-based attacks [9].

Corollary 4.8 Let $m(k)$ be the size of the database as a function of the security parameter. Let c be a positive integer and $\epsilon > 0$ a constant such that the CR assumption holds with parameters $(F(k), c, t(k), n(k))$, where $t = O(k \cdot m^{1/c})$, $n = O(k \cdot m^{1+1/c+\epsilon})$, and $|F(k)| = O(k \cdot m^{1/c+\epsilon})$. (A larger value of ϵ represents a more conservative assumption.) Then, assuming two-way anonymous communication, there is a one-round PIR protocol involving a single server and multiple clients in which the amortized communication and computation per query are $\tilde{O}(t) = \tilde{O}(km^{1/c})$. The protocol is computationally private as long as uncorrupted clients make together at least $n(k)$ random noise queries.

The feasible query domain of the above protocol allows to distribute the role of the server between many users without compromising efficiency or security. This gives rise to a conceptually attractive type of distributed storage systems, described in [33].

Achieving sublinear communication. While the PIR protocol given by Corollary 4.8 has a low complexity per client when the number of clients is large, it requires the total communication with all clients to be bigger than the database size. Indeed, a protocol with a sublinear total communication would imply a PIR protocol and hence key agreement in the standard model, which is not known to be implied by the CR assumption. We now briefly sketch a way for combining the protocol \mathcal{P} of Corollary 4.8 with any standard (single-server) PIR protocol \mathcal{P}' in order to reduce the communication complexity when the number of clients $u(k)$ is smaller than the database size $m(k)$.

Let $m'(k)$ be a database size for which \mathcal{P} remains secure if each of the u clients contributes a single noise query. (Note that $m'(k)$ should always be smaller than $u(k)$, to an extent that depends on the strength of the CR assumption; when c is big and ϵ is small, $m'(k)$ is close to $u(k)$.) We parse the m bits of the original database x as an $m' \times \ell$ matrix X where $\ell = \lceil m/m' \rceil$. Each client, who wishes to retrieve entry (i, j) of X , invokes the protocol \mathcal{P} as if it is retrieving the i -th bit from a database of size m' . Along with each of the sub-queries in \mathcal{P} , it sends a (standard) PIR query pointing to the j -th entry of a database of ℓ entries, generated according to \mathcal{P}' . (An independent invocation of \mathcal{P}' is used for each sub-query.) For each sub-query received from a client, the server obtains ℓ answers, each resulting from applying \mathcal{P} to the corresponding column of X , and then computes a single response by applying \mathcal{P}' to the database of ℓ answers. The resulting protocol has a low communication complexity if so do \mathcal{P} and \mathcal{P}' . Excluding the cost of pre-processing, the amount of server computation per client is typically of the order of ℓ , resulting in a total amount of computation that is close to $\ell \cdot m' = m$ when m' is close to u . Thus, a good choice of parameters yields a protocol which is close to optimal with respect to both communication and computation, regardless of the number of clients.

Acknowledgements. We thank Andreas Pfitzmann for pointing out the relevance of [2, 44] to implementing key

agreement based on anonymity, and Matthias Fitz for pointing out the relevance of [45] to implementing broadcast based on anonymity. We also thank David Chaum, Juan Garay, Venkat Guruswami, Tatsuaki Okamoto, Farzad Parvaresh and the anonymous referees for helpful discussions and pointers.

References

- [1] Anonymity bibliography. <http://www.freehaven.net/anonbib/>
- [2] B. Alpern and F. B. Schneider. Key exchange Using ‘Keyless Cryptography’. *Information Processing Letters* Vol. 16, pages 79–81, 1983.
- [3] A. Beimel, Y. Ishai, E. Kushilevitz, and J. F. Raymond. Breaking the $O(n^{1/(2k-1)})$ Barrier for Information-Theoretic Private Information Retrieval. In *Proc. 43rd FOCS*, pages 261–270, 2002.
- [4] A. Beimel, Y. Ishai, and T. Malkin. Reducing the servers’ computation in private information retrieval: PIR with pre-processing. *Journal of Cryptology*, 17(2), pages 125–151, 2004. Earlier version in CRYPTO 2000.
- [5] A. Beimel and T. Malkin. A Quantitative Approach to Reductions in Secure Computation. In *Proc. of 1st TCC*, pages 238–257, 2004.
- [6] C. H. Bennett, G. Brassard, and J. M. Robert. Privacy Amplification by Public Discussion. *SIAM J. Comput.* 17(2): 210–229 (1988).
- [7] R. Berman, A. Fiat, and A. Ta-Shma. Provable Unlinkability against Traffic Analysis. In *Proc. of 8th Financial Cryptography*, pages 266–280, 2004.
- [8] D. Bleichenbacher, A. Kiayias, and M. Yung. Decoding of Interleaved Reed Solomon Codes over Noisy Data. In *Proc. of ICALP 2003*, pages 97–108.
- [9] D. Bleichenbacher and P. Q. Nguyen. Noisy Polynomial Interpolation and Noisy Chinese Remaindering. In *Proc. of EUROCRYPT 2000*, pages 53–69.
- [10] C. Cachin, S. Micali, and M. Stadler. Computationally private information retrieval with polylogarithmic communication. In *Proc. of EUROCRYPT ’99*, pages 402–414.
- [11] R. Canetti. Security and composition of multiparty cryptographic protocols. In *J. of Cryptology*, 13(1), 2000.
- [12] D. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Commun. ACM*, Vol. 24(2), pages 84–88, 1981. Also: UC Berkeley M.Sc. Thesis, 1979.
- [13] David Chaum. Elections with Unconditionally-Secret Ballots and Disruption Equivalent to Breaking RSA. In *Proc. EUROCRYPT 1988*, pages 177–182.
- [14] S. Chawla, C. Dwork, F. McSherry, A. Smith, and H. Wee. Toward Privacy in Public Databases. In *Proc. of 2nd TCC*, pages 363–385, 2005.
- [15] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. *J. of the ACM*, 45:965–981, 1998. Earlier version in FOCS ’95.

- [16] D. Coppersmith and M. Sudan. Reconstructing curves in three (and higher) dimensional space from noisy data. In *Proc. of 35th STOC*, pages 136-142, 2003.
- [17] I. Dinur and K. Nissim. Revealing information while preserving privacy. In *Proc. of 22nd PODS*, pp. 202-210, 2003.
- [18] U. Feige, J. Kilian, and M. Naor. A minimal model for secure computation. In *Proc. of 26th STOC*, pages 554-563, 1994.
- [19] J. Feigenbaum, Y. Ishai, T. Malkin, K. Nissim, M. Strauss, and R. Wright. Secure Multiparty Computation of Approximations. In *Proc. 28th ICALP*, pages 927-938, 2001.
- [20] M. J. Fischer and R. N. Wright. Multiparty Secret Key Exchange Using a Random Deal of Cards. In *Proc. CRYPTO 1991*, pages 141-155.
- [21] M. Fitzi, J. Garay, U. Maurer, and R. Ostrovsky. Minimal Complete Primitives for Secure Multi-party Computation. In *Proc. Crypto 2001*, pages 80-100.
- [22] M. Fitzi, S. Wolf, and J. Wullschlegler. Pseudo-signatures, Broadcast, and Multi-party Computation from Correlated Randomness. In *Proc. CRYPTO 2004*, pages 562-578.
- [23] C. Gentry and Z. Ramzan. Single-Database Private Information Retrieval with Constant Communication Rate. In *Proc. 32nd ICALP*, pages 803-815, 2005.
- [24] O. Goldreich. *Foundations of Cryptography: Basic Applications*. Cambridge University Press, 2004.
- [25] V. Guruswami and F. Parvaresh. Personal communication.
- [26] V. Guruswami and A. Rudra. Explicit Capacity-Achieving List-Decodable Codes. In *Proc. 38th STOC*, pp. 1-10, 2006.
- [27] V. Guruswami, and M. Sudan. Improved decoding of Reed-Solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, Vol. 45(6), pages 1757-1767, 1999. Earlier version in FOCS '98.
- [28] D. Harnik, M. Naor, O. Reingold, and A. Rosen. Completeness in two-party secure computation: a computational view. In *Proc. 36th STOC*, pages 252-261, 2004.
- [29] R. Impagliazzo and M. Naor. Efficient Cryptographic Schemes Provably as Secure as Subset Sum. *J. Cryptology* 9(4), pages 199-216, 1996. Earlier version in FOCS '89.
- [30] R. Impagliazzo and D. Zuckerman. How to Recycle Random Bits. In *Proc. 30th FOCS*, pages 248-253, 1989.
- [31] Y. Ishai and E. Kushilevitz. Perfect Constant-Round Secure Computation via Perfect Randomizing Polynomials. In *Proc. 29th ICALP*, pages 244-256, 2002.
- [32] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. Batch codes and their applications. In *Proc. 36th STOC*, pages 373-382, 2004.
- [33] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. Cryptography from Anonymity. *Cryptology ePrint Archive*, Report 2006/084, 2006.
- [34] A. Kiayias and M. Yung. Secure Games with Polynomial Expressions. In *Proc. 28th ICALP*, pages 939-950, 2001.
- [35] J. Kilian. Founding cryptography on oblivious transfer. In *Proc. 20th STOC*, pages 20-31, 1988.
- [36] J. Kilian, E. Kushilevitz, S. Micali, and R. Ostrovsky: Reducibility and Completeness in Private Computations. *SIAM J. Comput.* 29(4): 1189-1208 (2000).
- [37] E. Kushilevitz and R. Ostrovsky. Replication is not needed: single database, computationally-private information retrieval. In *Proc. 38th FOCS*, pages 364-373, 1997.
- [38] H. Lipmaa. An Oblivious Transfer Protocol with Log-Squared Communication. In *Proc. ISC 2005*, pages 314-328.
- [39] U. M. Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory* 39(3): 733-742, 1993.
- [40] T. Moran and M. Naor. Basing cryptographic protocols on tamper-evident Seals. In *Proc. of 32nd ICALP*, pages 285-297, 2005.
- [41] M. Naor and B. Pinkas. Oblivious polynomial evaluation. *SIAM J. Comput.* 35(5), pages 1254-1281, 2006. Earlier version in STOC '99.
- [42] N. Nisan and D. Zuckerman. Randomness is Linear in Space. *J. Comput. Syst. Sci.*, Vol. 52(1), pages 43-52, 1996. Earlier version in STOC '93.
- [43] F. Parvaresh and A. Vardy. Correcting Errors Beyond the Guruswami-Sudan Radius in Polynomial Time. In *Proc. 46th FOCS*, pages, 285-294, 2005.
- [44] A. Pfitzmann and M. Waidner. Networks without user observability – design options. In *Proc. Eurocrypt '85*, pages 245-253, 1986. Revision in: *Computers and Security* 6/2 (1987) 158-166.
- [45] B. Pfitzmann and M. Waidner. Information-Theoretic Pseudosignatures and Byzantine Agreement for $t \geq n/3$. IBM Research Report RZ 2882 (#90830), 1996.
- [46] T. Rabin and M. Ben-Or. Verifiable Secret Sharing and Multiparty Protocols with Honest Majority. In *Proc. 21st STOC*, pages 73-85, 1989.
- [47] J. F. Raymond. Traffic Analysis: Protocols, Attacks, Design Issues, and Open Problems. In *Workshop on Design Issues in Anonymity and Unobservability*, pages 10-29, 2001.
- [48] M. K. Reiter and A. D. Rubin. Crowds: Anonymity for Web Transactions. *ACM Trans. Inf. Syst. Secur.*, 1(1), pages 66-92, 1998.
- [49] D. R. Simon. Anonymous Communication and Anonymous Cash. In *Proc. CRYPTO 1996*, pages 61-73.
- [50] P. L. Vora. Information Theory and the Security of Binary Data Perturbation. In *Proc. Indocrypt 2004*, pages 136-147.
- [51] P. Winkler. Cryptologic techniques in bidding and defense: Parts I, II, III, and IV. *Bridge Magazine*, April-July 1981.
- [52] A. C. Yao. How to generate and exchange secrets. In *Proc. 27th FOCS*, pages 162-167, 1986.