

# MCMix: Anonymous Messaging via Secure Multiparty Computation

Nikolaos Alexopoulos<sup>1</sup>, Aggelos Kiayias<sup>2</sup>, Riivo Talviste<sup>3</sup>, and Thomas Zacharias<sup>2</sup>

<sup>1</sup>Technische Universität Darmstadt

<sup>2</sup>School of Informatics, University of Edinburgh, UK

<sup>3</sup>Cybernetica AS, Estonia

{alexopoulos@tk.tu-darmstadt.de, akiayias@inf.ed.ac.uk, riivo@cyber.ee,  
tzachari@inf.ed.ac.uk}

## Abstract

We present MCMix, an anonymous messaging system that completely hides communication metadata and can scale in the order of hundreds of thousands of users. Our approach is to isolate two suitable functionalities, called dialing and conversation, that when used in succession realize anonymous messaging. With this as a starting point, we apply secure multiparty computation (“MC” or MPC) and proceed to realize them. We present an implementation using a prevalent MPC system (Sharemind) that is competitive in terms of latency with previous messaging systems that only offer much weaker privacy guarantees. Our solution can be instantiated in a variety of different ways with different MPC implementations, overall illustrating how MPC is a viable and competitive alternative to mix-nets and DC-nets for anonymous communication.

## 1 Introduction

In an era in which privacy in communications is becoming increasingly important, it is often the case that two parties want to communicate anonymously, that is to exchange messages while hiding the very fact that they are in conversation. A major problem in this setting is hiding the communication metadata: while existing cryptographic techniques (e.g., secure point-to-point channels implemented with TLS) are sufficiently well developed to hide the communication content, they are not intended for hiding the metadata of the communication such as its length, its directionality, and the identities of the communicating end points. Metadata are particularly important, arguably some times as important to protect as the communication content itself. The importance of metadata is reflected in General Michael Hayden’s quote “We kill people based on metadata”<sup>1</sup> and in the persistence of security agencies with programs like PRISM (by the NSA) and TEMPORA (by the GCHQ) in collecting metadata for storage and mining.

Anonymous communication has been pioneered in the work of Chaum, with mix-nets [Cha81] and DC-nets [Cha88] providing the first solutions to the problem of sender-anonymous communication. In particular, a mix-net enables the delivery of a set of messages from  $n$  senders to a recipient so that the recipient is incapable of mapping outgoing messages to their respective senders. A DC-net on the other hand, allows  $n$  parties to implement an anonymous broadcast channel so that any one of them can use it

---

<sup>1</sup>Complete quote: “We kill people based on metadata. But that’s not what we do with this metadata.” General M. Hayden. The Johns Hopkins Foreign Affairs Symposium. 1/4/2014.

to broadcast a message to the set of parties without any participant being able to distinguish the source. While initially posed as theoretical constructs, these works have evolved to actual systems that have been implemented and tested, for instance in the case of Mixminion [DDM03], that applies the mix-net concept to e-mail, in the case of Vuvuzela [VDHLZZ15] that applies the mix-nets concept to messaging and in the case of Dissent [WCGFJ12] that implements DC-nets in a client-server model.

It is important to emphasize that the adversarial setting we wish to protect against is a model where the adversary has a *global view* of the network, akin say to what a global eavesdropper would have if they were passively observing the Internet backbone, rather than a localized view that a specific server or sub-network may have. Furthermore, the adversary may manipulate messages as they are transmitted and received from users as well as block users adaptively. Note that in a more “localized” adversary setting one may apply concepts like Onion routing [SGR97], e.g., as implemented in the Tor system [DMS04], or Freenet [CSWH01] to obtain a reasonable level of anonymity with very low latency. Unfortunately such systems are susceptible to traffic analysis, see e.g., [JWJ<sup>+</sup>13], and, in principal, they cannot withstand a global adversary.

Given the complexity of the anonymous communication problem in general, we focus our application objective to the important special case of *anonymous messaging*, i.e., bidirectional communication with *both* sender and receiver anonymity against a third party, that requires moderately low latency and has relatively small payloads (akin to SMS text messaging). The question we ask is whether it is possible to achieve it with *simulation-based security*<sup>2</sup> while scaling to *hundreds of thousands* of users. In particular, we consider two types of entities in our problem specification, clients and servers, and we ask how is it possible that the servers assist the clients that are online to communicate privately without leaking *any* type of metadata to a global adversary, apart from the fact that they are using the system. Furthermore, we seek a decentralized solution, specifically one where no single entity in the system can break the privacy of the clients even if it is compromised. We allow the adversary to completely control the network as well as a subset of the servers and adaptively drop clients’ messages or manipulate them as it wishes.

**Our Contributions.** We present MCMix, the first anonymous messaging service that offers simulation-based security, under a well specified set of assumptions, and can scale to hundreds of thousands of users. In our solution, we adopt a different strategy compared to previous approaches to anonymous communication. Specifically, we provide a way to cast the problem of anonymous messaging natively in the setting of secure multiparty computation (MPC). MPC, since its initial inception [GMW87], is known to be able to distribute and compute securely any function, nevertheless, it is typically considered to be not particularly efficient for a large number of parties and thus inconsistent with problems like anonymous messaging. However, the commodity-based approach for MPC [Bea97] (client-server model), and more recent implementation efforts such as Fairplay [BDNP08], VIFF [DGKN09], Sharemind [Bog13], PICCO [ZSB13], OblivM [LWN<sup>+</sup>15], Araki et al. [AFL<sup>+</sup>16] and [FLNW17] increasingly suggest otherwise.

We first propose two ideal functionalities that correspond to the dialing operation and the conversation operation. The MCMix system proceeds in rounds, where in each round an invocation of either the dialing or the conversation ideal functionality is performed. The dialing functionality enables clients to either choose to dial another client or check whether anyone is trying to dial them (in practice in most dialing rounds the overwhelming majority of clients will be in dial-checking mode). If a matching pair is determined by the ideal functionality, then the caller will be notified that the other client has accepted their call and the callee will be notified about the caller. Moreover, the ideal functionality will deliver

---

<sup>2</sup>We use this term to refer to a level of metadata hiding that ensures, in a simulation based sense, that *no information* is leaked to an adversary. This is distinguished from weaker levels of privacy, such as e.g., a differential privacy setting where some controlled but non-trivial amount of information is leaked to the adversary.

to both clients a random tag that can be thought of as the equivalent of a “dead drop” or “rendezvous” point. Subsequently, the clients can access the conversation functionality using the established random tag. When two clients use the same random tag in the conversation functionality, their messages are swapped and thus they can send messages to each other (even concurrently).

The two ideal functionalities provide a useful abstraction of the anonymous messaging problem. We proceed now to describe how they can be implemented by an MPC system. It is easy to see that a straightforward implementation of the functionality programs results in a circuit of size  $\Theta(n^2)$ , where  $n$  is the number of online users accessing the functionalities. Such a solution would clearly be not scalable. We provide more efficient implementations that achieve  $O(n \log n)$  complexity in both cases with very efficient constants using state of the art oblivious sorting algorithms [HKI<sup>+</sup>12, BLT14].

Given our high level functionality realizations, we proceed to an explicit implementation in the Sharemind system [Bog13] using its SecreC programming language [BLR14]. We provide benchmarks for the Dialing and Conversation solutions. The Sharemind platform provides a 3-server implementation of information theoretically secure MPC. Our results showcase that our system can handle hundreds of thousands of users in a reasonable latency (little over a minute), that is consistent with messaging.

In order to provide theoretical evidence of further improving performance and scaling to even larger anonymity sets, we provide a parallelized version of the conversation functionality. Parallelization is a non-trivial problem in our setting since we would like to maintain anonymity across the whole user set; thus, a simplistic approach that breaks users into chunks solving dialing and conversation independently will isolate them to smaller “communication islands”; if two users have to be on the same island in order to communicate, this will lead to privacy loss that is non-simulatable and we would like to avoid. Our parallelized solution manages to make the interaction between islands, in a way that maintains strong privacy guarantees, at the cost of a correctness error that can become arbitrarily small. In this way, by utilizing a large number of servers, we provide evidence that the system can scale up to anonymity sets of up to half a million of users. To sum up, our contributions can be expressed by the following points:

1. A model for simulation-based anonymous messaging.
2. A realization of this model with a set of programs that are provably secure and expressed in a way so that they can be implemented in any MPC platform.
3. An implementation of our programs in Sharemind that can accommodate anonymity sets of hundreds of thousands of users.
4. A novel parallelization technique that allows our system to scale, in theory, even beyond the order of hundreds of thousands of users.

**Organization.** After shortly presenting some preliminary topics in section 2, we formalize the concept of anonymous messaging via an ideal MPC functionality and introduce the Dialing and Conversation programs in an abstract form that together solve the sender and receiver anonymous messaging problem (cf. Section 3). In Section 4, we present the general architecture of MCMix and in Sections 5 and 6, we propose a way to realize the Dialing and Conversation programs, using MPC. Then, in Section 7, we give more details regarding how the MCMix system implements anonymous messaging in a provably secure and privacy-preserving way. In Section 8, we present the results of benchmarking our prototype and in Section 9, we account for the client-side load of our system. In Section 11, we introduce a novel way to parallelize our conversation protocol in order to achieve even better scalability. Finally, in Section 10, we provide an overview of noticeable anonymous communication systems and when applicable, we compare their performance and security level to MCMix. Our concluding remarks are in Section 12.

## 2 Background

### 2.1 Secure Multiparty Computation and the Sharemind framework

Secure Multiparty Computation (MPC), is an area of cryptography concerned with methods and protocols that enable a set of users  $\mathcal{U} = u_1, \dots, u_n$  with private data  $d_1, \dots, d_n$  from a domain set  $D$ , to compute the result of a public function  $f(d_1, \dots, d_n)$  in a range set  $Y$ , without revealing their private inputs. For clarity, we assume that  $D$  consists only of actual messages, but  $f$  accepts also  $\perp$  as input, which denotes abstain behavior.

**Sharemind.** Sharemind [Bog13] is an MPC framework that offers a higher level representation of the circuit being computed in the form of a program written in a C-like language, namely the SecreC language [BLR14]. It uses three-server protocols that offer security in the presence of an honest server majority. That is, we assume that no two servers will collude in order to break the systems privacy. Our implementation is designed over the Sharemind system, but the general approach that we introduce for anonymous messaging can also be deployed over other MPC protocols. The security of Sharemind has been analyzed several settings including semi-honest and active attacks (e.g., [Bog13, PL15]).

### 2.2 Oblivious sorting

Sorting is used as a vital part of many algorithms. In the context of secure multiparty computation, sorting an array of values without revealing their final position, is called oblivious sorting. The first approach to sorting obliviously is using a data-independent algorithm and performing each compare and exchange execution obliviously. This approach uses sorting networks to perform oblivious sorting. Sorting networks are circuits that solve the sorting problem on any set with an order relation. What sets sorting networks apart from general comparison sorts is that their sequence of comparisons is set in advance, regardless of the outcome of previous comparisons. Various algorithms exist to construct simple and efficient networks of depth  $\mathcal{O}(\log^2 n)$  and size  $\mathcal{O}(n \log^2 n)$ . The three more used ones are Batcher’s odd-even mergesort and bitonic sort [Bat68] and Shellsort [She59]. All three of these networks are simple in principle and efficient. Sorting networks that achieve the theoretically optimal  $\mathcal{O}(\log n)$  and  $\mathcal{O}(n \log n)$  complexity in depth and total number of comparisons, such as the AKS-network [AKS83] exist, but the constants involved are so large that make them impractical for use. Note that even for 1 billion values, i.e.,  $n = 10^9$ , it holds that  $\log n < 30$  so, in practice, the extra log factor is preferable to the large constants. A major drawback of all sorting network approaches is that sorting a matrix by one of its columns would require oblivious exchange operations of complete matrix rows, which would be very expensive.

In recent years techniques have been proposed from Hamada et. al [HKI<sup>+</sup>12] to use well known data-dependent algorithms such as quicksort in an oblivious manner to achieve very efficient implementations, especially when considering a small number of MPC servers, which is very often the case. This approach uses the “shuffling before sorting“ idea, which means that if a vector has already been randomly permuted, information leaked about the outcome of comparisons does not leak information about the initial and final position of any element of the vector. More specifically, the variant of quicksort proposed in [HKI<sup>+</sup>12], needs on average  $\mathcal{O}(\log n)$  rounds and a total of  $\mathcal{O}(n \log n)$  oblivious comparisons. Complete privacy is guaranteed when the input vector contains no equal sorting keys, and in the case of equal keys, their number leaks. Furthermore, performance of the algorithm is data-dependent and generally depends on the number of equal elements, with the optimal case being that no equal pairs exist. Practical results have shown [BLT14] that this quicksort variant is the most efficient oblivious sorting algorithm available, when the input keys are constructed in a way that makes them unique.

In our algorithms we use the Quicksort algorithm together with a secret-shared index vector as

described in [BLT14]. This way, each sortable element becomes a unique value-index pair, providing us the optimal Quicksort performance and complete privacy. It also has the added benefit of making the sorting algorithm stable.

### 2.3 Identity-Based Key Agreement Protocols

Like in [LZ16], we make use of identity-based cryptography [Sha84] to circumvent the need for a Public Key Infrastructure (PKI), here, for the computation of the dead drops<sup>3</sup>. In identity-based cryptography, a Key Generation Center (KGC) using a master secret key, generates the users' secret keys, while the users' public keys are a deterministic function of their identity. In an *identity-based key agreement (ID-KA) protocol* (e.g. [Gün89, SKO00, Sma01, CK03, YL05, FG10, Wan13]), after receiving their secret keys, the users can mutually agree on shared keys given their secret keys and the other user's identity.

In our setting, we will apply ID-KA for the computation of the dead drops, where now the users compute their secret keys by combining partial secret keys issued by the MPC servers. Therefore, we adjust ID-KA to a multiple KGC setting where each MPC server plays the role of a KGC. In general, we can manage distributed key generation in a fault tolerant manner, using threshold secret-sharing techniques. However, since our threat model considers a passive (semi-honest adversary), we consider an  $m$ -out-of- $m$  instantiation, keeping protocol description simple. In particular, we naturally extend a pairing-based single KGC ID-KA protocol to a setting with  $m$  KGCs denoted by  $\text{KGC}_1, \dots, \text{KGC}_m$ . A *cryptographic pairing*  $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ , where  $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$  are multiplicative cyclic groups of prime order  $q$ , is an efficiently computable function such that for every pair of generators  $g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2$  and every pair of exponents  $x, y \in \mathbb{Z}_q$  it holds that:

1.  $e(g_1^x, g_2^y) = e(g_1, g_2)^{xy}$  (*bilinearity*).
2.  $e(g_1, g_2)^{xy}$  is a generator of  $\mathbb{G}_T$  (*non-degeneracy*).

The pairing  $e$  is called *symmetric* if  $\mathbb{G}_1 = \mathbb{G}_2 = \mathbb{G}$ , and *asymmetric* otherwise.

We provide two secure constructions of multiple KGC ID-KA protocol. The second construction additionally achieves forward secrecy, i.e. if the users' secret keys are compromised then past session keys are not leaked.

**Construction 1: Multiple KGC ID-KA.** We build upon the SOK ID-KA protocol introduced in [SKO00] proven secure in [PS09]. Our multiple KGC ID-KA protocol consists of the following algorithms:

– *Setup*: On common input  $1^\lambda$ , where  $\lambda$  is the security parameter,  $\text{KGC}_1, \dots, \text{KGC}_m$  agree on a symmetric cryptographic pairing  $e$  with parameters  $(e, \mathbb{G}, \mathbb{G}_T, q, g)$ , where  $g$  is a generator of  $\mathbb{G}$ , and two cryptographic hash functions  $H_1 : \{0, 1\}^* \rightarrow \mathbb{G}$  and  $H_2 : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ .

Next, each  $\text{KGC}_j, j \in [m]$  randomly chooses a partial master secret key  $\text{msk}_j = x_j \xleftarrow{\$} \mathbb{Z}_q$  and publishes its partial public key  $\text{mpk}_j = g^{x_j}$  that are combined in the protocol's public key  $\text{pk} := \prod_{j \in [m]} \text{mpk}_j$ . The public parameters of the protocol are  $\text{params} := (e, \mathbb{G}, \mathbb{G}_T, q, g, H_1, H_2, \text{pk})$ .

– *Secret Key Derivation*: For every user  $u_i$  with identity  $\text{ID}_i$ , each  $\text{KGC}_j, j \in [m]$ , on input  $\text{msk}_j$  generates the partial secret key  $\text{sk}_{i,j} := H_1(\text{ID}_i)^{x_j}$  and sends it to  $u_i$ . Upon receiving  $\text{sk}_{i,1}, \dots, \text{sk}_{i,m}$ , the user  $u_i$  obtains its secret key  $\text{sk}_i$  by setting

$$\text{sk}_i := \prod_{j \in [m]} \text{sk}_{i,j} = H_1(\text{ID}_i)^{\sum_{j \in [m]} x_j} .$$

– *Key Agreement*: Using their secret keys  $\text{sk}_a, \text{sk}_b$ , two users  $u_a, u_b$  agree on a key value  $K$  as follows:

<sup>3</sup>If preexisting PKI has already resolved the issue of users' public key distribution, then we can turn to the easier solution of classic Diffie-Hellman key exchange for dead drop computation (cf. Remark 6).

- $u_a$  computes the value  $K_{a,b} = e(\text{sk}_a, H_1(\text{ID}_b))$ .
- $u_b$  computes the value  $K_{b,a} = e(\text{sk}_b, H_1(\text{ID}_a))$ .
- $u_a$  and  $u_b$  agree on the key  $K = H_2(K_{a,b}) = H_2(K_{b,a})$ .

The correctness of the protocol follows from the bilinearity property of  $e$  as shown below:

$$\begin{aligned} K_{a,b} &= e(\text{sk}_a, H_1(\text{ID}_b)) = e(H_1(\text{ID}_a)^{\sum_{j \in [m]} x_j}, H_1(\text{ID}_b)) = e(H_1(\text{ID}_a), H_1(\text{ID}_b))^{\sum_{j \in [m]} x_j} = \\ &= e(H_1(\text{ID}_b), H_1(\text{ID}_a))^{\sum_{j \in [m]} x_j} = e(H_1(\text{ID}_b)^{\sum_{j \in [m]} x_j}, H_1(\text{ID}_a)) = e(\text{sk}_b, H_1(\text{ID}_a)) = K_{b,a}. \end{aligned}$$

The security of the original single SOK ID-KA protocol proven in [PS09], which is a special case of the multiple KGC protocol described above for  $m = 1$ , holds under the assumptions that  $H_1$  and  $H_2$  are modeled as random oracles and that the *computational bilinear Diffie-Hellman problem (CBDH)* is hard for the group  $\mathbb{G}$  of pairing  $e$ . Briefly, CBDH hardness assumption for  $\mathbb{G}$  states that for a randomly chosen triple of exponents  $x, y, z \in \mathbb{Z}_q$  and on input  $(g^x, g^y, g^z)$  it is hard to compute the value  $e(g, g)^{xyz}$ .

Given the security of the original ID-KA protocol for  $m = 1$ , it is straightforward that the multiple KGC ID-KA protocol described above is secure against any polynomially bounded semi-honest adversary that corrupts all-but-one of the  $m$  KGCs.

**Construction 2: Multiple KGC ID-KA with forward secrecy.** We build upon the pairing-based ID-KA protocol introduced in [Sma01] as modified in [CK03] that achieves security and forward secrecy as proven in [CCS07]. Our multiple KGC ID-KA protocol with forward secrecy consists of the following algorithms:

– *Setup*: On common input  $1^\lambda$ ,  $\text{KGC}_1, \dots, \text{KGC}_m$  agree on an asymmetric cryptographic pairing  $e$  with parameters  $(e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2)$  and two cryptographic hash functions  $H_1 : \{0, 1\}^* \rightarrow \mathbb{G}_1$  and  $H_2 : \{0, 1\}^* \times \{0, 1\}^* \times \mathbb{G}_2 \times \mathbb{G}_2 \times \mathbb{G}_T \rightarrow \{0, 1\}^*$ .

Next, each  $\text{KGC}_j$ ,  $j \in [m]$  randomly chooses a partial master secret key  $\text{msk}_j = x_j \stackrel{\$}{\leftarrow} \mathbb{Z}_q$  and publishes its partial public key  $\text{mpk}_j = g_2^{x_j}$  that are combined in the protocol's public key  $\text{pk} := \prod_{j \in [m]} \text{mpk}_j$ . The public parameters of the protocol are  $\text{params} := (e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, H_1, H_2, \text{pk})$ .

– *Secret Key Derivation*: For every user  $u_i$  with identity  $\text{ID}_i$ , each  $\text{KGC}_j$ ,  $j \in [m]$ , on input  $\text{msk}_j$  generates the partial secret key  $\text{sk}_{i,j} := H_1(\text{ID}_i)^{x_j}$  and sends it to  $u_i$ . Upon receiving  $\text{sk}_{i,1}, \dots, \text{sk}_{i,m}$ , the user  $u_i$  obtains its secret key  $\text{sk}_i$  by setting

$$\text{sk}_i := \prod_{j \in [m]} \text{sk}_{i,j} = H_1(\text{ID}_i)^{\sum_{j \in [m]} x_j}.$$

– *Key Agreement*: Using their secret keys  $\text{sk}_a := (s_a, r_a)$ ,  $\text{sk}_b := (s_b, r_b)$ , two users  $u_a, u_b$  agree on a key value  $K_{a,b} = K_{b,a}$  as follows:

- $u_a$  picks a random value  $t_a \stackrel{\$}{\leftarrow} \mathbb{Z}_q$  and sends  $g_2^{t_a}$  to  $u_b$ ;
- $u_b$  picks a random value  $t_b \stackrel{\$}{\leftarrow} \mathbb{Z}_q$  and sends  $g_2^{t_b}$  to  $u_a$ ;
- $u_a$  computes the values  $K_{a,1} = e(H_1(\text{ID}_b)^{t_a}, \text{pk}) \cdot e(\text{sk}_a, g_2^{t_b})$  and  $K_{a,2} = (g_2^{t_b})^{t_a}$ ;
- $u_b$  computes the values  $K_{b,1} = e(H_1(\text{ID}_a)^{t_b}, \text{pk}) \cdot e(\text{sk}_b, g_2^{t_a})$  and  $K_{b,2} = (g_2^{t_a})^{t_b}$ ;
- $u_a$  and  $u_b$  agree on the key

$$K_{a,b} := H_2(\text{ID}_a, \text{ID}_b, g_2^{t_a}, g_2^{t_b}, K_{a,2}, K_{a,1}) = H_2(\text{ID}_a, \text{ID}_b, g_2^{t_a}, g_2^{t_b}, K_{b,2}, K_{b,1}).$$

The correctness of the protocol follows from the bilinearity property of  $e$  as shown below:

$$\begin{aligned}
K_{a,1} &= e(H_1(\text{ID}_b)^{t_a}, \text{pk}) \cdot e(\text{sk}_a, g_2^{t_b}) = e(H_1(\text{ID}_b), g_2)^{t_a \sum_{j \in [m]} x_j} \cdot e(H_1(\text{ID}_a), g_2)^{t_b \sum_{j \in [m]} x_j} = \\
&= e(H_1(\text{ID}_b)H_1(\text{ID}_a), g_2)^{(t_a+t_b) \sum_{j \in [m]} x_j} = e(H_1(\text{ID}_a)H_1(\text{ID}_b), g_2)^{(t_b+t_a) \sum_{j \in [m]} x_j} = \\
&= e(H_1(\text{ID}_a), g_2)^{t_b \sum_{j \in [m]} x_j} \cdot e(H_1(\text{ID}_b), g_2)^{t_a \sum_{j \in [m]} x_j} = e(H_1(\text{ID}_a)^{t_b}, \text{pk}) \cdot e(\text{sk}_b, g_2^{t_a}) = K_{b,1} \\
K_{a,2} &= (g_2^{t_a})^{t_b} = (g_2^{t_b})^{t_a} = K_{b,2}.
\end{aligned}$$

The security and forward secrecy of the original single KGC ID-KA protocol proven in [CK03], which is a special case of the multiple KGC protocol described above for  $m = 1$ , holds under the assumptions that  $H_1$  and  $H_2$  are modeled as random oracles and that CBDH is hard for the group pair  $(\mathbb{G}_2, \mathbb{G}_1)$  of pairing  $e$ . Briefly, CBDH hardness assumption for  $(\mathbb{G}_2, \mathbb{G}_1)$  states that for a randomly chosen triple of exponents  $x, y, z \in \mathbb{Z}_q$  and on input  $(g_2^x, g_1^y, g_2^z)$  it is hard to compute the value  $e(g_1, g_2)^{xyz}$ .

Given the security and forward secrecy of the original ID-KA protocol for  $m = 1$ , it is straightforward that the multiple KGC ID-KA protocol described above preserves security and forward secrecy against any polynomially bounded semi-honest adversary that corrupts all-but-one of the  $m$  KGCs.

### 3 Ideal Anonymous Messaging

We formalize the concept of anonymous messaging in line with standard MPC security modeling. In particular, we capture the notion of an *ideal MPC functionality*  $\mathcal{F}$  that in presence of an ideal adversary  $\mathcal{S}$  receives inputs from a number of  $n$  users and computes the desired result w.r.t. some program  $f$ . An MPC protocol is said to be secure w.r.t. a class of programs, if its execution running in the presence of a real-world adversary results in input/output transcripts that are indistinguishable from the ideal setting that  $\mathcal{F}$  specifies for program  $f$ .

Subsequently, inspired by Tor, Vuvuzela and other related systems, we make use of the “rendezvous points” idea. Specifically, we instantiate  $\mathcal{F}$  w.r.t. two distinct “abstract” programs  $\text{DLN}_{\text{abs}}$  and  $\text{CNV}_{\text{abs}}$  that reflect the Dialing and Conversation functionalities respectively; the two programs are abstract in the sense that, in this section, they will be described at a high level algorithmic way that we will make concrete in the coming sections. The use of a random rendezvous point in the establishment of a communication channel between two users averts any denial of service attacks targeting specific users by other users at the conversation phase.

**Notation.** We write  $x \stackrel{\$}{\leftarrow} X$  to denote that  $x$  is sampled uniformly at random from set  $X$ . For a positive integer  $n$ , the set  $\{1, \dots, n\}$  is denoted by  $[n]$ . The  $j$ -th component of  $n$ -length tuple  $a$  is denoted by  $a[j]$ , i.e.  $a := (a[1], \dots, a[n])$ . We use  $\stackrel{c}{\approx}$  to express indistinguishability between transcripts, seen as random variables. By  $\text{negl}(\cdot)$  we denote that a function is negligible, i.e. asymptotically smaller than the inverse of any polynomial. We use  $\lambda$  as the security parameter.

Let  $\mathbf{x} = \langle x_1, \dots, x_n \rangle$  be a vector of users’ inputs. We denote by  $\text{EXEC}_{\mathcal{S}, \mathbf{x}}^{\mathcal{F}, f}(\lambda)$  the transcript of input/outputs in an ideal MPC execution of  $\mathcal{F}$  interacting with the ideal adversary  $\mathcal{S}$ , and by  $\text{EXEC}_{\mathcal{A}, \mathbf{x}}^{\mathbb{P}, f}(\lambda)$  the transcript of inputs/outputs in a real-world execution of MPC protocol  $\mathbb{P}$  w.r.t.  $f$  under the presence of adversary  $\mathcal{A}$ . By PPT, we mean that  $\mathcal{A}$  runs in probabilistic polynomial time.

#### 3.1 Entities and threat model

We consider a client-server MPC setting. Namely, the entities involved in an MPC protocol  $\mathbb{P}$  are (i) a number of  $n$  users  $u_1, \dots, u_n$  that provide their inputs  $\langle x_1, \dots, x_n \rangle$  and (ii) a number of  $m$  servers  $\text{Ser}_1, \dots, \text{Ser}_m$  that collectively compute an evaluation on the users’ inputs w.r.t. a program  $f$ . The users

engaged in a specific MPC execution form an active set  $\mathcal{U}_{\text{act}}$ . We consider an *ad-hoc* setting [BGIK16] of secure computation, where the program  $f$  is known in advance, but *not* the active user set  $\mathcal{U}_{\text{act}}$ .

An adversary against  $\mathbb{P}$  is allowed to have a *global view* of the protocol network. In addition, it may corrupt up to a fixed subset of  $\theta$  servers and has limited computational resources preventing it from breaking the security of the underlying cryptographic primitives.

In standard MPC cryptographic modeling, the security of  $\mathbb{P}$  is argued w.r.t. the functionality  $\mathcal{F}$  that specifies an “ideal” evaluation of  $f$ , where the privacy leakage is the minimum possible for the honest users. Thus, indistinguishability between the ideal and the real world setting implies that an adversary against  $\mathbb{P}$  obtains essentially no more information than this minimum leakage. In our description,  $\mathcal{F}$  merely leaks whether an honest user is online or not. This information is impossible to hide against a network adversary and hence it is a minimum level of leakage. On the other hand, information that can be typically inferred by traffic analysis, is totally protected by  $\mathcal{F}$ . This level of anonymity, sometimes referred to as unobservability, requires the participation of all online parties and the generation of “dummy traffic” independently of whether or not they wish to send a message in a particular round. As a result, any protocol  $\mathbb{P}$  that securely realizes  $\mathcal{F}$  where  $f$  represents a dialing or conversation program, should incorporate such a methodology. As we demonstrate, using MPC to realize  $\mathbb{P}$  is a natural way to determine the appropriate level and form of “dummy traffic” needed to realize this level of anonymity.

### 3.2 An ideal MPC functionality with adversarial influence for a family of programs

In a messaging system, dialing and conversation among users are operations where conflicts are likely to appear, e.g. two users may dial the same person, or conversation may be accidentally established on colluding communication channels (three equal rendezvous points are computed). One can think several other examples of operations where conflicts are possible, such as election tally where exactly one out of multiple ballots per voter must be counted, or deciding on the valid sequence of transactions on a blockchain ledger when forking occurs. Any program implementing this type of an operation must be able to resolve these conflicts. The way that conflict resolution is achieved, may depend on parameters like computation efficiency, communication complexity or user priority, yet in any case, a set of programs that implement the same operation are in some sense equivalent and may be clustered under the same family. A plausible requirement is that the choice of the family member that will be utilized should not affect the security standards of the operation implementation.

Consequently, in an MPC setting that supports the realization of any program in the family, it is desirable that security is preserved w.r.t. to the entire family, so that one can choose the family member that suits their custom requirements. To express this formally, we introduce a relaxation of the usual MPC functionality. Namely, the relaxed ideal MPC functionality  $\mathcal{F}$  is for a family of programs  $\{f_z\}_z$  in the presence of an ideal adversary  $\mathcal{S}$  that chooses the index  $z$  (this is the relaxation), where  $z$  can be parsed as the “code” that determines the family member  $f_z$ . We call this MPC with *adversarial influence*. The program  $f_z$  accepts as input a vector  $\mathbf{x} = \langle x_1, \dots, x_n \rangle$  of (i) valid messages from some domain  $D$  or (ii)  $\perp$ , if the user is inactive, i.e. not in  $\mathcal{U}_{\text{act}}$ . In our description, computation takes place even when a subset of users abstain from the specific execution by not providing inputs. To formalize the abstain behavior of user  $u_i$ , for every  $i \in [n]$  we define an ‘abstain $_i(\cdot)$ ’ predicate over  $D \cup \{\perp\}$  as follows:

$$\text{abstain}_i(x_i) := \begin{cases} 1, & \text{if } x_i = \perp \\ 0, & \text{if } x_i \in D \end{cases} \quad (1)$$

The ideal MPC functionality  $\mathcal{F}$  is presented in Fig. 1. Note that the relaxation suggests that the users will receive output from a program  $f_z$  for  $z$  that will be the ideal adversary’s choosing.

The security of a real-world MPC protocol  $\mathbb{P}$  is defined w.r.t. a class of programs  $\mathbf{F}$  as well as a family selected from  $\mathbf{F}$  as follows:

*Ideal MPC functionality  $\mathcal{F}$  with adversarial influence for programs  $\{f_z\}_z$*

- Upon receiving ‘start’ from  $\mathcal{S}$ , it sets the status to ‘input’ and initializes two lists  $L_{\text{input}}$  and  $L_{\text{corr}}$  as empty.
- Upon receiving (corrupt,  $u_i$ ) from  $\mathcal{S}$ , it adds  $u_i$  to  $L_{\text{corr}}$ .
- Upon receiving (send\_input,  $x_i$ ) from  $u_i$ , if  $u_i \in L_{\text{corr}}$ , then it sends (send\_input,  $u_i, x_i$ ) to  $\mathcal{S}$ . If  $u_i \notin L_{\text{corr}}$ , then it sends (i) (send\_input,  $u_i, \text{abstain}_i(x_i)$ ) to  $\mathcal{S}$ , where  $\text{abstain}_i(\cdot)$  is defined in Eq. (1).
- Upon receiving (receive\_input,  $u_i, \hat{x}_i$ ) from  $\mathcal{S}$ , if (i) the status is ‘input’ and (ii)  $(u_i, \cdot) \notin L_{\text{input}}$ , then if  $u_i \notin L_{\text{corr}}$ , it sets  $\hat{x}_i := x_i$ , else it sets  $\hat{x}_i := \hat{x}_i$ . Next, it adds  $(u_i, \hat{x}_i)$  to  $L_{\text{input}}$ .
- Upon receiving (compute,  $z$ ) from  $\mathcal{S}$ , if  $L_{\text{input}}$  contains records for all users in  $\mathcal{U}$ , it executes the following steps: first, then it computes the value vector

$$\mathbf{y} = \langle y_1, \dots, y_n \rangle \leftarrow f_z(\tilde{x}_1, \dots, \tilde{x}_n).$$

Then, it sends  $y_i$  to  $u_i$  for  $i, \dots, n$ , (hence,  $\mathcal{S}$  obtains  $\{y_i\}_{u_i \in L_{\text{corr}}}$ ).

Figure 1: The ideal MPC functionality  $\mathcal{F}$  with adversarial influence for a family of programs  $\{f_z : (D \cup \{\perp\})^n \rightarrow Y\}_z$  on input  $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ , interacting with the ideal adversary  $\mathcal{S}$ .

**Definition 1.** Let  $\mathbb{P}$  be an MPC protocol with  $n$  users and  $m$  servers and let  $\mathbf{F}$  be a class of programs. We say that  $\mathbb{P}$  is a  $(\theta, m)$ -secure MPC protocol w.r.t.  $\{f_z\}_z \subseteq \mathbf{F}$ , if for every active user set  $\mathcal{U}_{\text{act}} \in \mathcal{U}$  and every PPT adversary  $\mathcal{A}$  corrupting up to  $\theta$  out of  $m$  servers, there is an ideal adversary  $\mathcal{S}$  s.t. for every input vector  $\mathbf{x} = \langle x, \dots, x_n \rangle$ ,

$$\text{EXEC}_{\mathcal{S}, \mathbf{x}}^{\mathcal{F}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{A}, \mathbf{x}}^{\mathbb{P}}(\lambda).$$

### 3.3 The families of programs $\text{DLN}_{\text{abs}}$ and $\text{CNV}_{\text{abs}}$

An anonymous messaging scheme comprises the following two functionalities: (i) the *Dialing functionality*, which consists of the computation of a rendezvous point for a given pair of users who want to communicate, and (ii) the *Conversation functionality*, which represents the actual exchange of messages. For the families  $\text{DLN}_{\text{abs}}$  and  $\text{CNV}_{\text{abs}}$ , the parameter  $z$ , enables the adversary to choose (i) *how to handle collisions between multiple dialers* in the case of  $\text{DLN}_{\text{abs}}$ , and (ii) *how to handle the presence of three or more equal dead drops* in the case  $\text{CNV}_{\text{abs}}$  (which happens only in the case of malicious users). We note that this minimum level of adversarial manipulation does not affect the security features of the anonymity system, yet it allows for substantial performance gains in terms of the implementation.

We formally express the above functionalities by instantiating the generic MPC functionality  $\mathcal{F}$  w.r.t. the *Dialing program family*  $\text{DLN}_{\text{abs}}$  and the *Conversation program family*  $\text{CNV}_{\text{abs}}$  (i.e. we set  $f$  as  $\text{DLN}_{\text{abs}}$  and  $\text{CNV}_{\text{abs}}$ ). We note that for both the dialing and conversation program families, the verification that the parameter  $z$  has the proper structure can be suitably restricted so that it is tested efficiently by the program. For brevity, we omit further details.

#### 3.3.1 The Dialing program family $\text{DLN}_{\text{abs}}$

In the Dialing functionality, a rendezvous point for users  $u_i$  and  $u_j$  is set when two requests of the form  $(\text{DIAL}, u_i, u_j)$  and  $(\text{DIALCHECK}, u_j)$  have been produced. Thus, the Dialing program family  $\text{DLN}_{\text{abs}}$  receives inputs that are vectors of  $(\text{DIAL}, \cdot, \cdot)$  or  $(\text{DIALCHECK}, \cdot)$  requests, as well as  $\perp$  to denote user inactivity. That is,  $\mathcal{U}_{\text{act}}$  is the set of users that do not provide a  $\perp$  input. The program  $\text{DLN}_{\text{abs}}$  is

– **Domain:**  $(D_{DLN_{abs}} \cup \{\perp\})^n$ , where

$$D_{DLN_{abs}} := \left\{ \left\{ (DIAL, u_i, u_j), (DIALCHECK, u_i) \right\}_{u_i \neq u_j \in \mathcal{U}} \right\}$$

Namely, let  $\mathcal{U}_{act} := \{u_i \in \mathcal{U} \mid x_i \neq \perp\}$ ; a valid input  $x_i$  for user  $u_i \in \mathcal{U}_{act}$  consists of either (i) a  $(DIAL, u_i, u_j)$  request for some user  $u_j$  that  $u_i$  wants to dial, or (ii) a  $(DIALCHECK, u_i)$  request. For a vector of inputs  $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ , if  $x_i = (DIALCHECK, u_i)$  then  $M_i(\mathbf{x}) = \{j \mid x_j = (DIAL, u_j, u_i)\}$ , else is  $\emptyset$ . Parse  $z$  as a deterministic program  $R_{DLN}^z$ , such that for any  $\mathbf{x}$  if  $M_i(\mathbf{x}) \neq \emptyset$ , then  $R_{DLN}^z(i, \mathbf{x}) \in M_i(\mathbf{x})$ , else it is equal to  $\perp$ .

– **Range:**  $Y_{DLN_{abs}} := \langle \{y_i \mid y_i \in [a, b]\} \rangle_{u_i \in \mathcal{U}_{act}}$ , where  $[a, b]$  is a predetermined integer interval.

– **Function:** On input a vector  $\mathbf{x} = \langle x_1, \dots, x_n \rangle$  where each non- $\perp$  value  $x_i$  is either a  $(DIAL, u_i, u_j)$  request, or a  $(DIALCHECK, u_i)$  request,  $DLN_{abs}$  computes a vector  $\mathbf{y} = \langle y_i \rangle_{u_i \in \mathcal{U}_{act}}$ , as follows:

- Let  $\mathcal{J}_{act} := \{i \mid u_i \in \mathcal{U}_{act}\}$  be the set of indices that refer to active users. For  $i, j \in \mathcal{J}_{act}$ ,  $DLN_{abs}$  samples distinct random integers  $t_{i,j}$  from range  $[a, b]$ .
- For every  $i \in \mathcal{J}_{act}$ :
  - If  $x_i = (DIAL, u_i, u_j)$ , then if there is a  $j \in \mathcal{J}_{act}$  such that  $x_j = (DIALCHECK, u_j)$  and  $i = R_{DLN}^z(j, \mathbf{x})$ , then it sets  $t_i = t_{i,j}$ . Otherwise (i.e., there is no such  $j$ ), it sets  $t_i = t_{i,i}$ . In both cases, it sets  $y_i = t_i$ .
  - If  $x_i = (DIALCHECK, u_j)$ , then if there is a  $j \in \mathcal{J}_{act}$  such that  $j = R_{DLN}^z(i, \mathbf{x}) \neq \perp$ , then it sets  $t_i = t_{i,j}$  and a bit  $c_i = 1$ . Otherwise (i.e., there is no such  $j$ ), it sets  $t_i = t_{i,i}$  and a bit  $c_i = 0$ . In both cases, it sets  $y_i = (t_i, c_i)$ .
- It returns the value vector  $\mathbf{y} := \langle y_i \rangle_{u_i \in \mathcal{U}_{act}}$ .

Figure 2: The Dialing program family  $DLN_{abs} : (D_{DLN_{abs}} \cup \{\perp\})^n \rightarrow Y_{DLN_{abs}}$  with parameter  $z$ , where non- $\perp$  range values are integers sampled from range  $[a, b]$ .

parameterized by  $z$ , that specifies a deterministic program  $R_{DLN}^z(\cdot, \cdot)$  over pairs of inputs to resolve the case where more than one dial requests address the same user/dial checker. The Dialing program family  $DLN_{abs}$  is presented formally in Figure 2.

By the definition of  $DLN_{abs}$ , two active users  $u_i, u_j$  that have submitted matching dialing and dial check requests are going to be provided the same random integer  $t_i = t_j \in \{t_{i,j}, t_{j,i}\}$ , which establishes a rendezvous point. We will refer to these non- $\perp$  values in  $t_1, \dots, t_n$  as *dead drops*. In addition,  $DLN_{abs}$  returns to each dialchecker  $u_i$  a bit  $c_i$  which is 1 iff  $u_i$  has successfully established a rendezvous with some dialer. Such information is reasonable to be provided to a dialchecker, as  $t_i$  might be a random value that is not an actual dead-drop. Hence, the bit  $c_i$  communicates to the dialchecker that she has an incoming call (if nobody calls the dialchecker, then a random dead drop value is returned that nobody else shares with her). On the other hand, a dialer should not be able to infer information about the dial traffic and availability concerning some dialchecker, therefore  $DLN_{abs}$  does not provide this success check to the dialers.

### 3.3.2 The Conversation program family $CNV_{abs}$

Given the establishment of the dead drops, as set by  $DLN_{abs}$ , the Conversation program family  $CNV_{abs}$  realizes the operation of message exchange, where messages lie in some space  $\mathcal{M}$ . The program family  $CNV_{abs}$  is presented in Figure 3.

By the definition of  $\text{CNV}_{\text{abs}}$ , if every dead drop is not shared among three or more users, then two users  $u_i, u_j$  are going to exchange their messages  $m_i, m_j$  only if they provide the same dead drop  $t_i = t_j$ . Recall that if the dead drops are computed as outputs of the Dialing program family  $\text{DLN}_{\text{abs}}$  w.r.t. the same active set  $\mathcal{U}_{\text{act}}$ , then no more than two users share the same dead drop, which implies the correctness of  $\text{CNV}_{\text{abs}}$ . In the other cases, either (i) there is no matching dead drop or (ii) more than 2 matching dead drops exist. In case (ii), the parameter  $z$  specifies a deterministic program  $R_{\text{CNV}}^z$  among inputs which in turn determines the pair of matching dead drops. In any case, when a message exchange fails for some user, then  $\text{CNV}_{\text{abs}}$  returns back this message to the user for resubmission in an upcoming round.

*Program family  $\text{CNV}_{\text{abs}}$  parameterized by  $z$*

– **Domain:**  $(D_{\text{CNV}_{\text{abs}}} \cup \{\perp\})^n$ , where

$$D_{\text{CNV}_{\text{abs}}} := \{(\text{CONV}, t_i, m_i)\}_{u_i \in \mathcal{U}}^{t_i \in [a, b], m_i \in \mathcal{M}}$$

Namely, let  $\mathcal{U}_{\text{act}} := \{u_i \in \mathcal{U} \mid x_i \neq \perp\}$ ; a valid input for user  $u_i$  consists of a  $(\text{CONV}, t_i, m_i)$  request for rendezvous point tagged by  $t_i$  for sending message  $m_i$ .

For a vector of inputs  $\mathbf{x}$ , define  $N_i(\mathbf{x}) = \{j \mid x_j = (\text{CONV}, t_i, m_j)\}$ . Parse  $z$  as a deterministic program  $R_{\text{CNV}}^z$ , such that for any  $\mathbf{x}$  if  $N_i(\mathbf{x}) \neq \emptyset$  then  $R_{\text{CNV}}^z(i, \mathbf{x}) \in N_i(\mathbf{x})$ , else it is equal to  $\perp$ .

– **Range:**  $\{m_i \mid m_i \in \mathcal{U}_{\text{act}}\}_{u_i \in \mathcal{U}_{\text{act}}}$ .

– **Function:** On input a vector  $\langle x_1, \dots, x_n \rangle$  where each non- $\perp$  value  $x_i$  is a  $(\text{CONV}, t_i, m_i)$  request,  $\text{CNV}_{\text{abs}}$  returns a value  $\mathbf{y} = \langle y_i \rangle_{u_i \in \mathcal{U}_{\text{act}}}$ , as follows:

- Let  $\mathcal{J}_{\text{act}} := \{i \mid u_i \in \mathcal{U}_{\text{act}}\}$  be the set of indices that refer to active users. For every  $i \in \mathcal{J}_{\text{act}}$ : if  $j = R_{\text{CNV}}^z(i, \mathbf{x}) \neq \perp$ , then it sets  $y_i = m_j$ . Otherwise, it sets  $y_i = m_i$ .
- It returns the value vector  $\mathbf{y} = \langle y_i \rangle_{u_i \in \mathcal{U}_{\text{act}}}$ .

Figure 3: The Conversation program family  $\text{CNV}_{\text{abs}} : (D_{\text{CNV}_{\text{abs}}} \cup \{\perp\})^n \rightarrow Y_{\text{CNV}_{\text{abs}}}$  with parameter  $z$ , where non- $\perp$  dead drop values are integers sampled from a predetermined interval  $[a, b]$  and messages are taken from space  $\mathcal{M}$ .

### 3.4 Anonymous Messaging Systems

An anonymous messaging system is a pair of protocols that realize any two members of the families  $\text{DLN}_{\text{abs}}$  and  $\text{CNV}_{\text{abs}}$  under the security guarantee provided in Definition 1. Given such realization, anonymous communication can be achieved as a continuous sequence of interleaved invocations of dialing and conversation. In principle, dialing can be more infrequent compared to conversation, e.g., perform only a single dialing every certain number of conversation “rounds.” We note that the value of our relaxation of MPC security is on the fact that we can realize any member of the respective families.

### 3.5 Sharemind as a secure MPC platform

As already discussed, Sharemind will be the building platform for the implementation of our anonymous messaging scheme. As shown in [Bog13], Sharemind is information theoretically secure against a passive (honest-but-curious) adversary that corrupts 1-out-of-3 MPC servers. Subsequent work [PL15] provides interesting directions regarding the active security of Sharemind, even specifically for novel oblivious sorting algorithms [LP16]. However, in our implementation, we consider the case of passive security.

In more detail, let  $\mathbb{S}$  be the class of programs that can be written in Sharemind’s supporting language SecreC. In our analysis, we claim that Sharemind operates as a  $(1, 3)$ -secure MPC platform for any program family member of the class  $\mathbb{S}$  against passive adversaries, as in Definition 1. Using the above claim, we provide two SecreC programs and prove that they realize two members of the families  $DLN_{\text{abs}}$  and  $CNV_{\text{abs}}$ , (cf. Sections 5 and 6) hence obtaining an anonymous messaging system.

### 3.6 Alternative MPC platforms

For the purpose of the proposed anonymous messaging, Sharemind can be viewed as a black box providing MPC functionality. Hence, it is also possible to swap Sharemind for another MPC implementation providing different deployment or security properties. For example, recently, Furukawa et al. proposed a highly-optimised protocol for computation with an honest majority and security for malicious adversaries [FLNW17], that was further improved by Araki et al. [ABF<sup>+</sup>17]. Similarly, it is possible to support more than three computation parties. SPDZ [DPSZ12] is a practical MPC implementation that provides statistical security against an active adversary that corrupts up to  $m - 1$  parties. Its on-line computation and communication complexities are both  $O(m|C| + m^3)$ , where  $|C|$  stands for the computable arithmetic circuit size. In our setting, the lower bound for this circuit size is the number of users,  $n$ . Both actively secure MPC implementations mentioned here work in a preprocessing (i.e. offline/online) model.

## 4 System Architecture

Our work is presented in a manner that makes it easy to implement using any of the aforementioned MPC protocols in Section 2 and with any number of servers. However, for the sake of presentation, we assume three MPC servers, denoted by  $\text{Ser}_1, \text{Ser}_2, \text{Ser}_3$ . As a general idea, the protocol works in rounds, where in each round users break their input into shares and forward the shares to the servers, with each server receiving one share. Then, the servers interactively compute the desired output shares, which are in turn returned to the respective users. In our description, for simplicity we choose additive secret sharing, but other sharing schemes would not affect the functionality of our architecture.

Besides the MPC servers, the complete architecture of our system comprises an *entry* and an *output* server used to handle user requests. The entry and output servers may be located on the same or on different physical machines and are only trusted to relay messages.

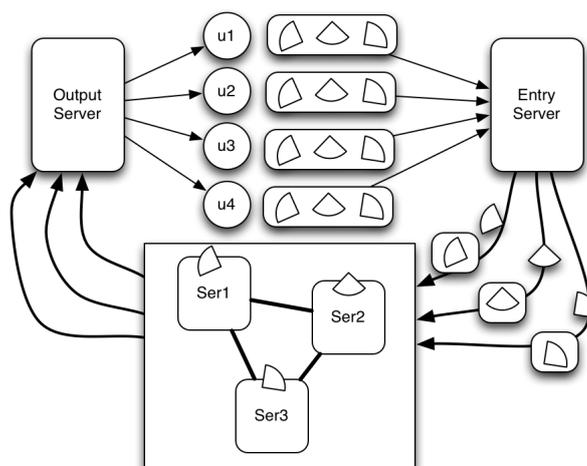


Figure 4: MCMix abstract architecture.

The complete architecture of our system, as shown in Fig. 4 includes the secure MPC servers, as well as *entry* and an *output* server used to handle user requests. The entry and output servers may be located on the same or on different physical machines and are only trusted to relay messages.

#### 4.1 Registration phase

At the beginning, the MPC servers  $\text{Ser}_1, \text{Ser}_2, \text{Ser}_3$  run the Setup phase of the secure multiple KGC ID-KA protocol (cf. Section 2.3) playing the role of three KGCs:  $\text{KGC}_1, \text{KGC}_2, \text{KGC}_3$  generating their partial master secret keys  $\text{msk}_1, \text{msk}_2, \text{msk}_3$ .

Before starting to use the system, each user  $u_i$  registers with a unique username  $\text{UN}_i$  of 64 bits. Then, each MPC server  $\text{Ser}_\ell, \ell \in \{1, 2, 3\}$  generates  $u_i$ 's partial secret key  $\text{sk}_{i,\ell}$  and sends it  $u_i$ . Upon receiving  $\text{sk}_{i,1}, \text{sk}_{i,2}, \text{sk}_{i,3}$ ,  $u_i$  combines the partial keys to obtain her ID-KA secret key  $\text{sk}_i$  as output of the secret key derivation algorithm. In addition, by performing standard key exchange operation,  $u_i$  obtains a symmetric key  $k_{i,\ell}$  for communication with each of  $\text{Ser}_\ell, \ell \in \{1, 2, 3\}$ . From this point on, any authentication and communication between  $u_i$  and the servers is performed using symmetric key cryptography. In the client-side,  $u_i$  can compute  $u_j$ 's ID-KA public key  $\text{pk}_j$  as a function of her username  $\text{UN}_j$  and agree on the ID-KA key  $K_{i,\ell}$ . In the rest of this paper, we set the length of the usernames  $\text{UN}_1, \dots, \text{UN}_n \in \mathcal{UN}$  to be 64 bits.

#### 4.2 Main phase

The main phase of the protocol for each round  $r$ , consists of the following steps:

- 1. Encoding:** Each user  $u_i$  generates a request  $a_i$ , as input to the MPC that is to be executed. All requests are padded to a fixed length specified by the running protocol to hide the content size.
- 2. Secret sharing:** Each user  $u_i$  creates three shares of the request using additive secret sharing, so that  $a_i = a_{i,\text{Ser}_1} + a_{i,\text{Ser}_2} + a_{i,\text{Ser}_3}$  holds. Note that the subscripts denote the MPC server that will process the share. Then each of the three shares intended for one of the MPC servers is encrypted with the respective symmetric key  $k_{i,\ell}$  using authenticated encryption. The result is a triple of the form  $a'_i = (a'_{i,\text{Ser}_1}, a'_{i,\text{Ser}_2}, a'_{i,\text{Ser}_3})$ , where  $a'_{i,\text{Ser}_\ell} := \text{Enk}_{k_{i,\ell}}(a_{i,\text{Ser}_\ell}), \ell = \{1, 2, 3\}$ . Then each user sends the encrypted shares along with her username  $\text{UN}_i$ , as a package to the entry server.
- 3. MPC input preparation:** Before the start of round  $r$ , the entry server groups the packages received already and sends each share along with its associated username to the respective MPC servers. It is important to note that the use of an entry server is only to synchronize the MPC servers and to provide the shares in the same order to each of them. For notation simplicity and without loss of generality, we assume that the entry server arranges  $u_i$  as the user that submitted the  $i$ -th input. Then, each MPC server  $\text{Ser}_\ell$  receives a sequence of the form  $a'_{\text{Ser}_\ell} = \langle a'_{1,\text{Ser}_\ell}, \dots, a'_{n,\text{Ser}_\ell} \rangle$ . We denote as  $n$  the number of users that provided an input in round  $r$ . In addition to  $a'_{\text{Ser}_\ell}$ , the MPC servers also receive a sequence of the users' usernames in corresponding order, that is a sequence of the form  $\text{UN} = \langle \text{UN}_1, \dots, \text{UN}_n \rangle$ , where  $\text{UN}_i$  is the registered username of the user that provided input  $i$ .
- 4. Order check:** Each MPC server computes a hash of the usernames in the order they appear in its input sequence, as  $H(\text{UN}_1 || \dots || \text{UN}_n)$ , and exchanges it with the other MPC servers. In case the three hashes do not match, it is implied that the order of the usernames provided to the three servers was different. Thus, a denial of service attack has taken place by either the entry server or one of the MPC servers (considering they reported a false hash). This step is optional when considering only privacy implications of a malicious entry server.
- 5. Decryption and authentication:** At this point, authentication is performed implicitly by each server via decrypting the received share with the symmetric key corresponding to the username that came with

the share. Thus shares  $a_{\text{Ser}_\ell} = \langle a_{\text{Ser}_\ell,1}, \dots, a_{\text{Ser}_\ell,n} \rangle$ , with  $a_{\text{Ser}_\ell,i} := \text{Dec}_{k_{i,\ell}}(a'_{\text{Ser}_\ell,i})$  are ready for the MPC.

**6. MPC algorithm:** The MPC servers execute the MPC protocol.

**7. Encryption and return:** Each MPC server encrypts each output share with the respective symmetric key and forwards shares of the form  $b'_{\text{Ser}_\ell} = \langle b'_{1,\text{Ser}_\ell}, \dots, b'_{n,\text{Ser}_\ell} \rangle$ , where each share  $b'_{i,\text{Ser}_\ell}$  is paired with the username  $\text{UN}_i$  of  $u_i$ , to the output server. The output server collects the shares corresponding to the same user and returns a package of the form  $(b'_{i,\text{Ser}_1}, b'_{i,\text{Ser}_2}, b'_{i,\text{Ser}_3})$  to each user  $u_i$ .

**8. Decryption and reconstruction:** Each user decrypts the received shares with the respective symmetric key and adds them, resulting in  $b_i = b_{i,\text{Ser}_1} + b_{i,\text{Ser}_2} + b_{i,\text{Ser}_3}$ , where  $b_{i,\text{Ser}_\ell} = \text{Dec}_{k_{i,\ell}}(b'_{i,\text{Ser}_\ell})$ . The value  $b_i$  is the final output of the MPC protocol for each user  $u_i$  for round  $r$ .

**Remark 1.** The entry and output servers are used for practical reasons. The main function they perform is grouping the received packages of shares and forwarding them to/from the servers. As they have no information about the symmetric keys exchanged between users and servers at the registration phase, they schedule the traffic consisting of encrypted shared data. Hence, if entry and output servers are malicious, they can do no more than an adversary controlling the network.

## 5 The Dialing Protocol

The dialing protocol enables a user  $u_i$  to notify another user  $u_j$  that she wants to start a conversation, much like how the telephone protocol works. The protocol runs in rounds to deter possible timing attacks, where in each round, every online active user will either send a DIAL request or a DIALCHECK request. All requests are mutually indiscriminate. For clarity, we first provide a description of the Dialing protocol steps. Then, we proceed with the efficient program  $\text{DLN}_{\text{sort}}$  implementing it.

### 5.1 Protocol description

The protocol runs in seven steps, where steps 2-6 are executed by the MPC servers. Steps 1 and 7 are executed locally by each user.

**1. Encoding:** The inputs  $x_1, \dots, x_n$  are of the form of  $(\text{DIAL}, u_i, u_j)$  requests,  $(\text{DIALCHECK}, u_i)$  requests, or  $\perp$ , representing the action each user takes for this dialing round. For simplicity, assume that the users are enumerated as  $u_1, \dots, u_n$  consistently with the input sequence  $x_1, \dots, x_n$ , i.e.  $u_i$  is the user that submitted the  $i$ -th input. As a result, the active users that submitted non- $\perp$  values, are enumerated as  $u_1, \dots, u_{\text{act}}$ , where  $\text{act}$  is the size of the active set  $\mathcal{U}_{\text{act}}$ . The inputs of the active users are encoded as triples of the form  $a_i := (a_i[1], a_i[2], a_i[3])$  where the third component is an *input wire ID*  $\text{wid}_i$ . The wire IDs are initially set to zero, but in the following Step 2,  $\text{wid}_i$  will be set unique for  $u_i$ .

In particular, if  $u_i$  wants to dial  $u_j$ , then the  $(\text{DIAL}, u_i, u_j)$  request is encoded as  $(\text{UN}_i, \text{UN}_j, 0)$  where  $\text{UN}_i$  and  $\text{UN}_j$  are the usernames of the dialer and the dialee respectively. If  $u_i$  is a dial checker, then the  $(\text{DIALCHECK}, u_i)$  request is encoded as  $(C, \text{UN}_j, 0)$ , where (i)  $C$  is a special value designated to denote a dial check and is different from any possible username value, and (ii)  $\text{UN}_j$  is the checker's own username.

**2. Assigning wire ID values:** As a first step, the MPC protocol assigns unique wire IDs for each user. This is done by setting the third component  $a_i[3]$  of the encoded triple  $a_i$  to  $i$ . Given the order  $u_1, \dots, u_{\text{act}}$ , for each  $u_i$ , we have that  $\text{wid}_i := i$ . These wire IDs are needed internally for the MPC calculation and express the order in which the inputs were received so that the respective outputs will be delivered in the same order.

**3. Checking input validity:** The protocol then checks if any of the first two members of each triple, denoted by  $a_i[1]$  and  $a_i[2]$ , is equal to the submitter's username. This check ensures that inputs are

encoded in a way that does not compromise the security of the system. The threat here is that a user  $u_i$  might try to impersonate a user  $u_j$  by encoding a DIALCHECK input as  $a_i = (C, \text{UN}_j, \text{wid}_i)$ . That attack would allow user  $u_i$  to receive a dial request that was intended for user  $u_j$ . A similar problem arises when considering a user  $u_i$  encoding a DIAL input as  $a_i = (\text{UN}_l, \text{UN}_j, \text{wid}_i)$ . In this case, user  $u_j$  will think the dial originated from user  $u_l$ . To avert such impersonation attacks, it is enough for the MPC protocol to check that either the first or the second member of an input tuple is equal to the username of the user that submitted that input. This, along with the fact that the input is sent from the user to each MPC server using authenticated encryption (cf. step 2 of the architecture in section 4) guarantees that no impersonation attack can take place.

In more detail, if the input is a DIALCHECK request, then this check ensures that the second member of the tuple is the user's own username. In the case of a DIAL request, the check ensures that a user can only impersonate another user when she dials herself, that is a request of the form  $a_i = (\text{UN}_j, \text{UN}_i, \text{wid}_i)$  is created by user  $u_i$ . In this case, this request does not affect the protocol. If the check fails for the encoded input  $a_i$ , then the input is set to  $a_i = (0, 0, \text{wid}_i)$  and does not affect the protocol.

**4. Sorting by usernames:** The encoded input triples are first sorted according to their second components using the oblivious Quicksort algorithm of [HKI<sup>+</sup>12], implemented according to [BLT14]. Observe that every non-zero second component is either (i) the username  $\text{UN}_j$  of dialee  $u_j$  in a dial request from some user  $u_i$ , or (ii) the username  $\text{UN}_j$  from dial checker  $u_j$ . Thus, when a triple  $(C, \text{UN}_j, \text{wid}_j)$  is adjacent to some triple  $(\text{UN}_i, \text{UN}_j, \text{wid}_i)$  with a non-zero second component, this determines a dial pair between  $u_i, u_j$ . We note that two special conflict cases may appear:

**I.**  $(C, \text{UN}_j, \text{wid}_j)$  is adjacent to two dial triples as

$\dots, (\text{UN}_i, \text{UN}_j, \text{wid}_i), (C, \text{UN}_j, \text{wid}_j), (\text{UN}_{i'}, \text{UN}_j, \text{wid}_{i'}), \dots$

**II.** Two or more adjacent dial triples correspond to  $(C, \text{UN}_j, \text{wid}_j)$ . The sorting would then appear as  $\dots, (\text{UN}_{i'}, \text{UN}_j, \text{wid}_{i'}), (\text{UN}_i, \text{UN}_j, \text{wid}_i), (C, \text{UN}_j, \text{wid}_j), \dots$

**5. Connecting neighbors:** Next, requests are processed individually by looking at both their neighbors' triples to determine if there is a dial for any given dial check request. Of course, requests at the first and last place of the sorted vector need only look at one neighbor. Thus, we can claim that any dial check request will have a suitable dial request as its neighbor or not at all.

In more detail, for every user  $u_i$ , the protocol produces a pair  $b := (b_i[1], b_i[2])$ , where  $b_i[2]$  is  $\text{wid}_i$  and  $b_i[1]$  is either (i) the username  $\text{UN}_j$  of some user  $u_j$  that dialed  $u_i$ , or (ii) 0, if no dial request has been made for  $u_i$ , or  $u_i$  has made a dial request.

**6. Sorting by wire IDs:** As a final sorting step, the protocol needs to sort the processed requests according to their wire IDs in order for the correct requests to be forwarded to each user. The latter sort, performed on  $\langle b_1, \dots, b_{\text{act}} \rangle$  according to the wire IDs can again be implemented by the Quicksort algorithm of [HKI<sup>+</sup>12]. The result of the last sorting is a vector  $\langle \hat{b}_1, \dots, \hat{b}_{\text{act}} \rangle$  where  $\hat{b}_i$  is a pair  $(\hat{b}_i[1], \hat{b}_i[2])$  that corresponds to  $u_i$  and  $\hat{b}_1$  is essentially either (i) a username  $\text{UN}_j$  or (ii) a zero value, in both cases indexed by  $\hat{b}_2 := \text{wid}_i$ .

**7. Computing the dead drops:** After the Quicksort algorithm is completed, the active users  $u_1, \dots, u_{\text{act}}$  are delivered the values  $\hat{b}_1[1], \dots, \hat{b}_1[\text{act}]$  respectively. Then, dialer  $u_i$  that knows  $\text{UN}_j$ , and dial checker  $u_j$  that obtained  $\text{UN}_i$ , can calculate their shared dead drop value for dialing round  $r$  as follows:

$$\begin{aligned} t_i &:= H(K_{i,j}, r), & \text{if } \hat{b}_i[1] = 0 \\ t_j &:= H(K_{j,i}, r), & \text{if } \hat{b}_i[1] = \text{UN}_j \end{aligned}$$

Above,  $H$  is a standard cryptographic hash function, and  $r$  is the round number. The values  $K_{i,j}, K_{j,i}$  are the ID-KA keys that  $u_i$  and  $u_j$  compute by running the key agreement algorithm `GenerateKey` on input  $(\text{sk}_i, \text{UN}_j)$  and  $(\text{sk}_j, \text{UN}_i)$  respectively (cf. Section 2), where  $\text{sk}_i, \text{sk}_j$  are the secret keys of  $u_i$  and  $u_j$ . Recall that ID-KA operations are over a finite multiplicative group of prime order  $q$ .

**Input:** a sequence  $\langle x_1, \dots, x_n \rangle$  where  $x_i$  is either a (DIAL,  $u_i, u_j$ ) request, a (DIALCHECK,  $u_i$ ) request, or  $\perp$ . All  $\perp$  inputs are stacked last.

**Output:** a sequence  $\langle y_i \rangle_{i:x_i \neq \perp}$ , where  $y_i$  either is a  $\kappa$ -bit integer  $t_i$ , if  $x_i = (\text{DIAL}, u_i, u_j)$ , or a pair of a  $\kappa$ -bit integer  $t_i$  and a bit  $c_i$ , if  $x_i = (\text{DIALCHECK}, u_i)$ .

---

```

1. For each  $i \leftarrow 1, \dots, n$ 
if  $x_i = \perp$  then
    Set  $\text{act} := i - 1$  ;
    Break loop ;
else if  $x_i = (\text{DIAL}, u_i, u_j)$  then
    Set  $a_i := (a_i[1], a_i[2], a_i[3]) \leftarrow (\text{UN}_i, \text{UN}_j, 0)$  ;
else if  $x_i = (\text{DIALCHECK}, u_i)$  then
    Set  $a_i := (a_i[1], a_i[2], a_i[3]) \leftarrow (C, \text{UN}_i, 0)$  ;
end if
2. For each  $i \leftarrow 1, \dots, \text{act}$ 
Set  $\text{wid}_i$  as  $a_i[3] \leftarrow i$  ;
3. For each  $i \leftarrow 1, \dots, \text{act}$ 
if  $a_i[1] \neq \text{UN}_i$  AND  $a_i[2] \neq \text{UN}_i$  then
    Set  $a_i[1] = a_i[2] = 0$  ;
end if
4.  $\langle a_i \rangle_{i:x_i \neq \perp}$  according to second coordinate using Quicksort;
5. For each  $i \leftarrow 1, \dots, \text{act}$ 
if  $a_i[1] = C$  AND  $a_i[2] = a_{i-1}[2]$  then
    Set  $b_i := (b_i[1], b_i[2]) \leftarrow (a_{i-1}[1], a_i[3])$  ;
else if  $a_i[1] = C$  AND  $a_i[2] = a_{i+1}[2]$  then
    Set  $b_i := (b_i[1], b_i[2]) \leftarrow (a_{i+1}[1], a_i[3])$  ;
else
    Set  $b_i := (b_i[1], b_i[2]) \leftarrow (0, a_i[3])$  ;
end if
6. Sort tuples  $\langle b_i \rangle_{i:x_i \neq \perp}$  according to second coordinate using Quicksort;
7. For each  $i \leftarrow 1, \dots, \text{act}$ 
if  $a_i[1] = \text{UN}_i$  then
    Set  $t_i \leftarrow H(\text{GenerateKey}(a_i[1], a_i[2]), r)$  ;
    Set  $y_i \leftarrow t_i$  ;
else if  $a_i[1] = C$  AND  $b_i[1] \in \mathcal{UN}$  then
    Set  $t_i \leftarrow H(\text{GenerateKey}(a_i[1], b_i[1]), r)$  ;
    Set  $y_i \leftarrow (t_i, 1)$  ;
else if  $a_i[1] = C$  AND  $b_i[1] = 0$  then
    Pick  $\rho_i \xleftarrow{\$} \{0, 1\}^{64}$  ;
    Set  $t_i \leftarrow H(\text{GenerateKey}(sk_i, \rho_i), r)$  ;
    Set  $y_i \leftarrow (t_i, 0)$  ;
end if
return  $\mathbf{y} := \langle y_i \rangle_{i:x_i \neq \perp}$  .

```

Figure 5: The Dialing program  $DLN_{\text{sort}}$  realizing the Dialing program  $DLN_{\text{abs}}$  for dialing round  $r$ , and users  $u_1, \dots, u_n$  with usernames  $\text{UN}_1, \dots, \text{UN}_n \in \{0, 1\}^{64}$ . The value  $C$  denotes a dial check request.

We stress that the dead drop value is at least 64 bits long to make accidental collisions unlikely, although our system can tolerate them. By the correctness of the ID-KA protocol, it holds that  $K_{i,j} = K_{j,i}$ , hence we have that  $t_i = t_j$ .

On the other hand, if user  $u_i$  dial checked but  $\hat{b}_i[1] = 0$  (no one dialed  $u_i$ ), then for uniformity reasons, she computes a random dead drop as above by inserting a random value  $\rho_i$  in place of  $\text{UN}_j$ , i.e. she sets  $t_i := H(\text{GenerateKey}(\text{sk}_i, \rho_i), r)$ .

Note that if  $u_i$  has dialchecked, then either (i) she established a rendezvous point with  $u_j$ , if  $\hat{b}_1 = \text{UN}_j$ , or (ii) no one dialed her, if  $\hat{b}_1 = 0$ . Thus, she can set a “success” bit  $c_i$  to 1 or 0 respectively, indicating her successful engagement in the dialing round  $r$ . Besides, if  $u_i$  is a dialer that dialed  $u_j$ , then she always computes the value  $t_i := H(\text{GenerateKey}(\text{sk}_i, \text{UN}_j), r)$ , regardless of the success of her dialing request. Hence, she can not infer a success bit.

## 5.2 The Dialing program $\text{DLN}_{\text{sort}}$

The program  $\text{DLN}_{\text{sort}}$  implementing the Dialing protocol is presented in Fig. 5.

Following Subsection 3.3.1, we show that  $\text{DLN}_{\text{sort}}$  realizes the member of the Dialing program family  $\text{DLN}_{\text{abs}}$  that corresponds to our sorting process. Namely, in Step 4 of  $\text{DLN}_{\text{sort}}$  (Sorting by usernames), the inputs are arranged according to an ordering of their second coordinate. Thus, we set the index  $z$  that parameterizes the family  $\text{DLN}_{\text{abs}}$  to be the string  $z_{\text{qs2}}$  as follows:  $z_{\text{qs2}}$  is parsed as the deterministic program  $R_{\text{DLN}}^{z_{\text{qs2}}}$  that takes as input an index  $i$  and array of triples  $\mathbf{x}$  in encoded form, and outputs the index  $j$  so that when the array is sorted according to Quicksort ordering on the second coordinate,  $x_i$  is the left neighbor of the encoded  $x_j$ . We prove the correctness of the Dialing program in the following theorem:

**Theorem 1.** *Let  $n$  be the number of users,  $\kappa \geq 64$  be the dead drop string length and  $q$  be the prime order of the underlying ID-KA group. Let  $H$  be the cryptographic hash function modeled as a random oracle. Then, the Dialing program  $\text{DLN}_{\text{sort}}$  described in Fig. 5 implements the member of the Dialing program family  $\text{DLN}_{\text{abs}}$  described in Fig. 2 for parameter  $z_{\text{qs2}}$  with correctness error  $\frac{n^4}{4q} + \frac{n}{2\kappa}$ .*

*Proof.* By the fact that in the arrangement of input sequence  $\langle x_1, \dots, x_n \rangle$  the  $\perp$  values are stacked last, if the first  $\perp$  input corresponds to index  $i$ , then all users  $u_1, \dots, u_{i-1}$  are exactly the ones that are active. Therefore, the set of indices  $\mathcal{J}_{\text{act}}$  referring to active users is set simply as  $\{1, \dots, \text{act}\}$ .

The inputs of all active users  $u_1, \dots, u_{\text{act}}$  are encoded into triples  $a_1, \dots, a_{\text{act}}$ . Observe that input validity (Step 3) always holds in the correctness setting. Besides, by the correctness of Quicksort, if  $u_i$  is a dial checker and at least one user has made a dial request for  $u_i$ , then  $u_i$  will receive the username of the dialer whose encoded input is at an adjacent position. By the symmetric property of ID-KA key generation algorithm, we have that if  $u_i, u_j$  have received each others public keys, then

$$\begin{aligned} t_i &:= H(\text{GenerateKey}(\text{sk}_i, \text{UN}_j), r) = \\ &= H(\text{GenerateKey}(\text{sk}_j, \text{UN}_i), r) = t_j. \end{aligned}$$

Let  $\mathbf{x} = \langle x_i, \dots, x_n \rangle$  be an input vector and let  $M_i(\mathbf{x}) := \{j \mid x_j = (\text{DIAL}, u_j, u_i)\}$ . By the above, we have that for every  $i \in \mathcal{J}_{\text{act}}$ , the program  $\text{DLN}_{\text{sort}}$  sets  $t_i$  such that

- If  $x_i = (\text{DIAL}, u_i, u_j)$ , and running Quicksort on the second coordinate of the encoded triples sets  $x_i$  as the left neighbor of the dial checker  $x_j$ , then it sets  $t_i = t_j$ . The latter is equivalent to  $i = R_{\text{DLN}}^{z_{\text{qs2}}}(j, \mathbf{x})$ . Otherwise, it sets  $t_i$  to be a pseudorandom value.

- If  $x_i = (\text{DIALCHECK}, u_i)$  and for some  $j \in M_i(\mathbf{x})$ , the encoded triple of dialer  $u_j$  is the left neighbor of the encoded  $x_i$  when running Quicksort on the second coordinate, then it sets  $t_i = t_j$  and a bit  $c_i = 1$ . The latter is equivalent to  $j = R_{\text{DLN}}^{z_{\text{qs2}}}(i, \mathbf{x})$ . Otherwise, it sets  $t_i$  to be a pseudorandom value and a bit  $c_i = 0$ .

Hence, in both cases,  $t_i$  is computed consistently with the description of the Dialing program family  $\text{DLN}_{\text{abs}}$  in Fig. 2 for parameter  $z_{\text{qs}2}$ . The implementation is perfect (no error) if all the dead drops are distinct, so it remains to bound the probability that the bad event of dead drop collision does not happen.

Since the hash function  $H$  is modeled as a random oracle, the computed dead drops are not distinct only if

1. Two distinct user pairs  $u_i, u_j$  and  $u_{i'}, u_{j'}$  agree on the same key value. Here, we make use of the fact that in the multiple KGC version of the SOK ID-KA protocol that we apply (cf. Section 2.3, Construction 1), key agreement is generated as shown below:

$$\text{sk}_i := H_1(\text{UN}_i)^x = g^{xy_i}, \text{sk}_j := H_1(\text{UN}_j)^x = g^{xy_j} \xrightarrow{\text{SOK}} e(g, g)^{xy_i y_j},$$

where (i)  $g$  is a generator of the cyclic group  $\mathbb{G}$  of pairing  $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$  with order  $q$ , (ii)  $x$  is the combined master secret key from all KGCs partial master secret keys, and (iii)  $y_i, y_j \in \mathbb{Z}_1$  are the unique exponents s.t.  $H_1(\text{UN}_i) = g^{y_i}$  and  $H_1(\text{UN}_j) = g^{y_j}$  respectively.

By the fact that the group is prime order cyclic, distinct user pairs  $u_i, u_j$  and  $u_{i'}, u_{j'}$  agree on the same key value when for the respective secret keys  $\text{sk}_i = g^{xy_i}, \text{sk}_j = g^{xy_j}, \text{sk}_{i'} = g^{xy_{i'}}, \text{sk}_{j'} = g^{xy_{j'}}$  it holds that

$$xy_i y_j \equiv xy_{i'} y_{j'} \pmod{q} \stackrel{x \neq 0}{\Leftrightarrow} y_i y_j \equiv y_{i'} y_{j'} \pmod{q} \quad (2)$$

Since  $H_1$  is modeled as random oracle, the exponents  $y_i, y_j, y_{i'}, y_{j'}$  are randomly selected, thus the probability that Eq. (2) holds for some fixed  $i, j, i', j'$  is  $\frac{1}{q}$ . By the union bound, the probability that there exist distinct  $i, j, i', j' \in [n]$  s.t. Eq. (2) holds is upper bounded by

$$\sum_{(\{i,j\}, \{i',j'\})} \frac{1}{q} = \frac{n(n-1)}{2} \cdot \frac{(n-2)(n-3)}{2} \cdot \frac{1}{q} = \frac{n!}{4q(n-4)!}.$$

2. A dial checker  $u_i$  that did not pair with any dialer chooses a random value  $\rho_i$  s.t.  $\text{GenerateKey}(\text{pk}_i, \rho_i)$  matches some other key value. Since the group is cyclic, the probability that this happens is no more than  $\frac{n}{q}$ .
3. All agreed key values are distinct but two of them have the same hash for round  $r$ . The probability that this happens is no more than  $\frac{n}{2^\kappa}$ .

We conclude that the overall correctness error is no more than

$$\frac{n!}{4q(n-4)!} + \frac{n}{q} + \frac{n}{2^\kappa} \leq \frac{n^4}{4q} + \frac{n}{2^\kappa}.$$

□

**Remark 2.** The correctness error  $\frac{n^4}{4q} + \frac{n}{2^\kappa}$  is typically a negligible value in our setting. To provide intuition, consider the case with a number of  $n = 100000 < 2^{17}$  users, dead drop size  $\kappa = 64$  bits and group size  $q \geq 2^{128}$ . The error for this case is less than  $\frac{2^{17 \cdot 4}}{2^{128}} + \frac{2^{17}}{2^{64}} \approx 2^{-47}$ .

## 6 The Conversation Protocol

The Conversation protocol facilitates the actual exchange of messages associated with the same  $t$  dead drop value, which represents a rendezvous point computed in the final step of a Dialing protocol execution. It is expected that no more than two messages will have the same  $t$  value due to its large bit-size,

although our system can handle collisions as we will see later. As in the previous section, we first provide a description of the Conversation protocol and then the corresponding program labeled  $\text{CNV}_{\text{sort}}$  that implements it. At this point, we have to highlight our assumption that a valid message  $m_i$  at the input has its least significant bit (LSB) equal to 0. This flag which could also be a discrete fourth member of our tuple, is useful at (i) conflict resolution when more than two dead drops are identical and (ii) the parallelization of our protocol discussed in Section 10.

## 6.1 Protocol description

The protocol is executed via the following steps, where steps 2-6 are executed by the MPC servers. Step 1 is executed locally by each user.

**1. Encoding:** The inputs are of the form of  $(\text{CONV}, t_i, m_i)$  requests, or  $\perp$ . Again, we assume that the users are enumerated as  $u_1, \dots, u_n$  consistently with the order they submitted their input sequence  $x_1, \dots, x_n$ , hence all  $\perp$  values are stacked last. Active users' inputs are encoded as triples of the form  $a_i := (a_i[1], a_i[2], a_i[3])$  where the third component is an input wire ID  $\text{wid}_i$  that will be uniquely assigned in the following step. In particular, if  $u_i$  wants to engage in conversation, then the  $(\text{CONV}, t_i, m_i)$  request is encoded as  $(t_i, m_i, 0)$ . In case  $u_i$  is not engaging in conversation the request will use a random dead drop value and a random message.

**2. Assigning wire ID values:** As a first step, the MPC protocol assigns unique wire IDs for each user. This is done by setting the third component  $a_i[3]$  of the encoded triple  $a_i$  to  $i$ . Thus, for each  $u_i$ , we have that  $\text{wid}_i := i$ .

**3. Sorting by dead drops:** The encoded input triples are first sorted according to their first components using the oblivious Quicksort algorithm of [HKI<sup>+</sup>12]. As a result, the inputs of any two users that share the same dead drop value will become adjacent.

**4. Exchanging adjacent messages:** By construction, two inputs with the same dead drop value indicate a pair of users  $u_i$  and  $u_j$  that wish to communicate. Thus, the protocol generates a vector  $\langle b_1, \dots, b_n \rangle$ , where each  $b_i$  is a pair  $(b_i[1], b_i[2])$ , of which the second component is  $\text{wid}_i$  and the first component is either (i) the message of some adjacent encoded input, or (ii) the original message  $m_i$ , if message exchange did not take place for  $u_i$  because there was no matching dead drop or due to conflict (three or more equal dead drops). As already mentioned, the LSB of two exchanged messages is set to 1. In the special conflict case where three or more values share the same dead drop  $t$ , an arrangement would be as follows:

$$\dots, (t', m_k, k), (t, m_j, j), (t, m_i, i), (t, m_{i'}, i'), \dots$$

In this case, the messages of  $u_i$  and  $u_j$  will be exchanged and  $u_{i'}$  will obtain back his message at the end of the protocol, notifying him to resubmit.

**5. Sorting by wire IDs:** As in the Dialing protocol (Step 5), the Conversation protocol performs a Quicksort on the processed requests according to their mutually distinct wire IDs in order for the correct requests to be forwarded to each user. The result is a vector  $\langle \hat{b}_1, \dots, \hat{b}_n \rangle$  where  $\hat{b}_i$  is a pair  $(\hat{b}_i[1], \hat{b}_i[2])$  that corresponds to  $u_i$  and is either (i) a message  $m_j$  from some user  $u_j$  or (ii) the original message  $m_i$ , in both cases indexed by  $\text{wid}_i$ .

**6. Forwarding messages:** At the end, the protocol discards the wire IDs and creates the output vector  $\mathbf{y} = \langle y_1, \dots, y_n \rangle := \langle \hat{b}_1[1], \dots, \hat{b}_n[1] \rangle$ . Thus, each  $y_i$  is either (i) a message  $m_j$  from some user  $u_j$  or (ii) the self-generated message  $m_i$ . Finally, the users  $u_1, \dots, u_n$  are delivered the values  $y_1, \dots, y_n$ .

**Remark 3.** In reality, the dead drop value  $t_i$  of some user  $u_i$  is not exactly the value she received from a dialing protocol execution. For conversation round  $r$  it is computed as  $t_i := H(t_{(\text{dialing})_i}, r)$ , where  $t_{(\text{dialing})_i}$  is the dead drop for  $u_i$ , generated by the dialing protocol and acts as the seed for the creation of an ephemeral dead drop for each conversation round.

**Remark 4.** Due to the size of dead drops values, the probability that a collision on randomly generated dead drop values will occur can be made very small. Even in the case of a collision, the client of the user that was affected would just resend that message in the next round, as it would know that a collision occurred because it received a message it could not decrypt.

*The Conversation Program  $\text{CNV}_{\text{sort}}$*

**Input:** a sequence  $\langle x_1, \dots, x_n \rangle$  where  $x_i$  is either a  $(\text{CONV}, t_i, m_i)$  request, or  $\perp$ . All  $\perp$  inputs are stacked last.

**Output:** a sequence of messages  $\langle y_i \rangle_{x_i \neq \perp}$ .

---

```

1. For each  $i \leftarrow 1, \dots, n$ 
if  $x_i = \perp$  then
  Set  $\text{act} := i - 1$  ;
  Break loop ;
end if
if  $x_i = (\text{CONV}, t_i, m_i)$  then
  Set  $a := (a_i[1], a_i[2], a_i[3]) \leftarrow (t_i, m_i, 0)$  ;
end if
2. For each  $i \leftarrow 1, \dots, \text{act}$ 
  Set  $\text{wid}_i$  as  $a_i[3] \leftarrow i$  ;
3. Sort tuples  $\langle a_i \rangle_{x_i \neq \perp}$  according to first coordinate  $a_i[1]$  using Quicksort;
4. For each  $i \leftarrow 1, \dots, \text{act} - 1$ 
if  $a_i[1] = a_{i+1}[1]$  AND  $\text{LSB}(a_i[2]) = \text{LSB}(a_{i+1}[2]) = 0$  then
  Set the LSB of  $a_i[2]$  and  $a_{i+1}[2]$  to 1 ;
  Set  $b_i \leftarrow (a_{i+1}[2], a_i[3])$  ;
  Set  $b_{i+1} \leftarrow (a_i[2], a_{i+1}[3])$  ;
end if
5. Sort tuples  $\langle b_i \rangle_{i: x_i \neq \perp}$  according to second coordinate (which is the wire id) using Quicksort;
6. For each  $i \leftarrow 1, \dots, \text{act}$ 
  Set  $y_i \leftarrow b_i[1]$  ;
return  $y := \langle y_i \rangle_{i: x_i \neq \perp}$  .

```

Figure 6: The Conversation program  $\text{CNV}_{\text{sort}}$  realizing the Conversation program  $\text{CNV}_{\text{abs}}$  for conversation round  $r$ , dead drop size  $\kappa \geq 64$  and users  $u_1, \dots, u_n$  with messages taken from space  $\mathcal{M}$ .

### 6.1.1 The Conversation program $\text{CNV}_{\text{sort}}$

The program  $\text{CNV}_{\text{sort}}$  implementing the Conversation protocol is presented in Fig. 6.

Following Section 3.3.2, we show that  $\text{CNV}_{\text{sort}}$  realizes the member of the Conversation program family  $\text{CNV}_{\text{abs}}$  that corresponds to our sorting process. Namely, in Step 3 of  $\text{CNV}_{\text{sort}}$  (Sorting by dead drops), the inputs are arranged according to an ordering of their first coordinate. Thus, we set the index  $z$  that parameterizes the family  $\text{CNV}_{\text{abs}}$  to be the string  $z_{\text{qs1}}$  as follows:  $z_{\text{qs1}}$  is parsed as the deterministic program  $R_{\text{CNV}}^{z_{\text{qs1}}}$  that takes as input an index  $i$  and array of triples  $\mathbf{x}$  in encoded form, and outputs the index  $j$  so that when the array is sorted according to Quicksort ordering on the first coordinate, the encoded triple of  $u_i$  (or resp.  $u_j$ ) has no neighbors on the left of the sorted array and the encoded triple of  $u_j$  (or resp.  $u_i$ ) is the right neighbor of the encoded triple of  $u_i$  (or resp.  $u_j$ ). Formally, we prove the following theorem.

**Theorem 2.** Let  $n$  be the number of users and  $\kappa \geq 64$  be the dead drop string length. The Conversation program  $\text{CNV}_{\text{sort}}$  described in Fig. 6 implements the member of the Conversation program family  $\text{CNV}_{\text{abs}}$  described in Fig. 3 for parameter  $z_{\text{qs1}}$ .

*Proof.* By the fact that  $\perp$  values are stacked last, if the first  $\perp$  input corresponds to index  $i$ , then all users  $u_1, \dots, u_{i-1}$  are exactly the ones that are active. Therefore, the set of indices  $\mathcal{J}_{\text{act}}$  referring to active users is set simply as  $\{1, \dots, \text{act}\}$ .

The rest of the proof follows by the correctness of the Quicksort algorithm. Specifically, if for two user  $u_i, u_j$  it holds that when sorted according to first coordinate the encoded triple of  $u_i$  (resp.  $u_j$ ) is the leftmost (resp. rightmost) neighbor of the encoded triple of  $u_j$  (resp.  $u_i$ ), then  $u_i$  and  $u_j$  constitute the first pair that share the same dead drop value  $t$  and will exchange their messages. The latter is equivalent to  $j = R_{\text{CNV}}^{z_{\text{qs1}}}(i, \mathbf{x})$  and  $i = R_{\text{CNV}}^{z_{\text{qs1}}}(j, \mathbf{x})$ . In any other case,  $u_i$  and  $u_j$  will receive back their messages  $m_i, m_j$ .

As a result, in both cases message exchange is done consistently with the description of the Conversation program family  $\text{CNV}_{\text{abs}}$  in Fig. 3 for parameter  $z_{\text{qs1}}$ . The implementation is perfect.  $\square$

## 7 The MCMix Anonymous Messaging System

Having presented the general architecture of our system in Section 4 and the Dialing and Conversation protocols and programs in Sections 5 and 6 respectively, we now show how these programs are implemented in our architecture. Our system consists of two MPC instances of the general architecture in Section 4, executing one after the other or independently in parallel. One implements the Dialing protocol and the other the Conversation protocol. Below, we specify the operations of general architecture for each of our two protocols. We note with the prime symbol, e.g. **1'**, the specification of the respective step, e.g. **1**, of the general architecture.

### 7.1 Dialing

The execution of the Dialing protocol for round  $r$  follows the steps of section 4 with the following particularities:

**1'. Encoding:** The input of user  $u_i$  is encoded as  $a_i = (\text{UN}_i, \text{UN}_j, 0)$ , in the case of a dial to user  $u_j$ , or as  $a_i = (C, \text{UN}_i, 0)$  in the case of a dial request, as specified by Step 1 of the Dialing program  $\text{DLN}_{\text{sort}}$  in Fig. 5.

**6'. MPC algorithm:** The MPC server secure computation consists of Steps 2-6 of  $\text{DLN}_{\text{sort}}$ .

**8'. Decryption and reconstruction:** The reconstructed value  $b_i$  received by user  $u_i$  is the output  $b_i$  of Step 6 of  $\text{DLN}_{\text{sort}}$ .

**9'. Dead drop calculation:** As an extra step, the dead drop value  $t_i$  is calculated by each user by performing Step 7 of  $\text{DLN}_{\text{sort}}$ .

### 7.2 Conversation

The execution of the conversation protocol for round  $r$  follows the steps of Section 4 with the following particularities:

**1'. Encoding:** Input is encoded as  $a_i = (t_i, m_i, 0)$ , with  $t_i$  being a dead drop calculated by the final step of a previous dialing round in the case of a real conversation request (also taking into account Remark 3), or a random value in the case the user does not want to send a message (but still wants to protect her privacy), according to the Conversation program  $\text{CNV}_{\text{sort}}$  in Fig. 6.

**6'. MPC algorithm:** The MPC server secure computation consists of Steps 2-6 of  $\text{CNV}_{\text{sort}}$ .

**8'. Decryption and reconstruction:** The reconstructed value  $b_i$  received by the user that provided input  $i$  is the output  $y_i$  of Step 6 of  $\text{CNV}_{\text{sort}}$  and is the message intended for this user.

### 7.3 Security of MCMix

We prove our security theorem for the general  $\theta$ -out-of- $m$  case, as in Definition 1, using the parameters  $z_{\text{qs2}}$  and  $z_{\text{qs1}}$  defined in Sections 5 and 6 respectively. We first remark that in the Dialing program  $\text{DLN}_{\text{sort}}$  described in Fig. 2 and in the Conversation program  $\text{CNV}_{\text{sort}}$  described in Fig. 3, the server computation part that we implement in an MPC manner consists of the Steps 2-6 of these programs. Therefore, for the clarity of the statement and proof of our theorem we define the following programs and their related abstracted program families.

- The program  $\text{DLN}_{\text{sort}}^{2-6}$ , which executes the Steps 2-6 of  $\text{DLN}_{\text{sort}}$ .
- The program family  $\text{DLN}_{\text{abs}}^{2-6}$  parameterized by  $z$  that has as input a sequence of encoded triples  $a_1, \dots, a_{\text{act}}$  as in Step 2 of  $\text{DLN}_{\text{sort}}$ , and returns a sequence of pairs  $\langle b_1, \dots, b_{\text{act}} \rangle$ , where  $b_i[2] = \text{wid}_i := i$  and
  - If  $a_i = (\text{UN}_i, \text{UN}_j, i)$ , then it sets  $b_i[1] = a_i[2]$ .
  - If  $a_i = (C, \text{UN}_i, i)$ , then if there is a  $j \in \mathcal{J}_{\text{act}}$  such that  $j = R_{\text{DLN}}^z(i, \mathbf{x}) \neq \perp$ , then it sets  $b_i[1] = a_j[1]$ . Otherwise (i.e., there is no such  $j$ ), it sets  $b_i[1] = 0$ .
- The program  $\text{CNV}_{\text{sort}}^{2-6}$ , which executes the Steps 2-6 of  $\text{CNV}_{\text{sort}}$ .
- The program family  $\text{CNV}_{\text{abs}}^{2-6}$  parameterized by  $z$  that has as input a sequence of encoded triples  $a_1, \dots, a_{\text{act}}$  of the form  $a_i = (t_i, m_i, 0)$  as in Step 2 of  $\text{CNV}_{\text{sort}}$ , and outputs a value vector  $\langle y_1, \dots, y_{\text{act}} \rangle$  as  $\text{CNV}_{\text{abs}}$  on parameter  $z$ . Namely, if there is a  $j \in \mathcal{J}_{\text{act}}$  such that  $j = R_{\text{CNV}}^z(i, \mathbf{x}) \neq \perp$ , then it sets  $y_i = a_j[2]$ . Otherwise, it sets  $y_i = a_i[2]$ .

**Theorem 3.** *Let  $\kappa$  be the dead drop size,  $n$  be the number of users,  $m$  be the number of servers and  $q$  the size of the underlying Diffie-Hellman group, where  $n, m$  are polynomial in  $\lambda$ ,  $\kappa = \Theta(\lambda)$  and  $q = \Omega(2^\lambda)$ . Let  $\mathbb{P}$  be an MPC protocol with  $n$  users and  $m$  servers ( $\theta, m$ )-secure MPC protocol w.r.t. program families  $\text{DLN}_{\text{sort}}^{2-6}$  and  $\text{CNV}_{\text{sort}}^{2-6}$ . Then, MCMix implemented over  $\mathbb{P}$  is an anonymous messaging system by securely realizing the program families  $\text{DLN}_{\text{abs}}$  and  $\text{CNV}_{\text{abs}}$  for parameters  $z_{\text{qs2}}$  and  $z_{\text{qs1}}$  respectively.*

*Proof.* We denote by  $\text{DLN}_{\text{abs}}^{z_{\text{qs1}}}$  and  $\text{CNV}_{\text{abs}}^{z_{\text{qs2}}}$  the members of the families  $\text{DLN}_{\text{abs}}$  and  $\text{CNV}_{\text{abs}}$  for parameters  $z_{\text{qs1}}$  and  $z_{\text{qs2}}$ , respectively. The proof consists of two parts as follows:

**I. Security of the Dialing Protocol  $\text{DLN}_{\text{sort}}$ :** Let  $\mathcal{A}_{\text{DLN}}$  be a passive adversary against the security of  $\text{DLN}_{\text{sort}}$  in a dialing round  $r$  that corrupts the input and output server and up to  $\theta$  servers of the MPC protocol  $\mathbb{P}$ . We will show that there exists an ideal adversary  $\mathcal{S}_{\text{DLN}}$  s.t. for every input vector  $\mathbf{x} = (x_1, \dots, x_n)$ ,

$$\text{EXEC}_{\mathcal{S}_{\text{DLN}}, \mathbf{x}}^{\mathcal{F}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{A}_{\text{DLN}}, \mathbf{x}}^{\mathbb{P}[\text{DLN}_{\text{sort}}]}(\lambda),$$

where  $\mathbb{P}[\text{DLN}_{\text{sort}}]$  denotes the implementation of  $\text{DLN}_{\text{sort}}$  over  $\mathbb{P}$ .

The program  $\text{DLN}_{\text{sort}}^{2-6}$  consists of the following steps:

- It receives an input of triples  $\langle a_i := (a_i[1], a_i[2], a_i[3] := 0) \rangle_{i \in [\text{act}]}$ , that encode dial or dialcheck requests as described in Step 1 of  $\text{DLN}_{\text{sort}}$ .
- It executes the Steps 2-6 of  $\text{DLN}_{\text{sort}}$ .

- It outputs the vector of pairs  $\langle b_1, \dots, b_{\text{act}} \rangle$  sorted according to the second coordinate using Quick-sort, where  $b_i[2] = \text{wid}_i$  and
  - If  $a_i = (\text{UN}_i, \text{UN}_j, 0)$ , then it sets  $b_i[1] = a_i[2]$ .
  - If  $a_i = (C, \text{UN}_i, 0)$ , then if there is a  $j \in \mathcal{J}_{\text{act}}$  such that  $j = R_{\text{DLN}}^{\text{zqs}^2}(i, \mathbf{x}) \neq \perp$ , then it sets  $b_i[1] = a_j[1]$ . Otherwise (i.e., there is no such  $j$ ), it sets  $b_i[1] = 0$ .

Given  $\text{DLN}_{\text{sort}}^{2-6}$ , we write the realization of  $\text{DLN}_{\text{sort}}$  compatibly with the MCMix architecture as follows: let  $\mathbb{P}[\text{DLN}_{\text{sort}}^{2-6}]$  be the implementation of  $\text{DLN}_{\text{sort}}^{2-6}$  in the MPC protocol  $\mathbb{P}$ . Let  $\pi_d$  be a protocol that “completes” the description of  $\text{DLN}_{\text{sort}}$ , by having access to  $\text{DLN}_{\text{sort}}^{2-6}$  and defining the steps run in the users’ side, as well as in the entry and output servers. The communication among the MPC servers and the users via the entry and output servers is done in a share-wise authenticated encryption manner described in Section 4.

An execution of the above protocol, denoted by  $\pi_d^{[\text{DLN}_{\text{sort}}^{2-6}]}$ , for some round  $r$ , is as follows:

- Upon receiving  $(\text{send\_input}, x_i)$ , each user  $u_i$  encodes  $x_i$  as  $a_i := (a_i[1], a_i[2], 0)$  according to Step 1 of  $\text{DLN}_{\text{sort}}$  and sends  $(a_i, \text{UN}_i)$  to the entry server.
- The entry server packs the messages provided by the active users  $a_1, \dots, a_{\text{act}}$  and forwards them as the message  $\langle (\text{UN}_1, a_1), \dots, (\text{UN}_{\text{act}}, a_{\text{act}}) \rangle$  to the MPC servers. Next, the MPC servers in  $\mathbb{P}$  perform the order check for the matching of  $H(\text{UN}_1 || \dots || \text{UN}_{\text{act}})$  described in Section 4.2 (Step 4). If the check is successful, then the servers initiate an MPC execution of  $\mathbb{P}[\text{DLN}_{\text{sort}}^{2-6}]$  on input  $\langle a_1, \dots, a_{\text{act}} \rangle$ .
- Upon finishing computation,  $\mathbb{P}[\text{DLN}_{\text{sort}}^{2-6}]$  sends its output  $\langle (\text{UN}_1, b_1), \dots, (\text{UN}_{\text{act}}, b_{\text{act}}) \rangle$  to the output server which in turn, issues  $b_1, \dots, b_{\text{act}}$  to the users  $u_1, \dots, u_{\text{act}}$ .
- Upon receiving  $b_i$ , each user  $u_i$  can compute her dead drop as an ID-KA session key, like it is described in Step 7 of  $\text{DLN}_{\text{sort}}$ .

By the parameter setting of the statement of Theorem 1,  $\text{DLN}_{\text{sort}}$  realizes the program  $\text{DLN}_{\text{abs}}$  with parameter  $z_{\text{qs}1}$  with error no more than

$$\frac{n^4}{4q} + \frac{n}{2^\kappa} = \frac{\text{poly}(\lambda)}{\Omega(2^\lambda)} + \frac{\text{poly}(\lambda)}{2^{\Theta(\lambda)}} = \text{negl}(\lambda).$$

The above equation implies that  $\text{DLN}_{\text{sort}}^{2-6}$  realizes  $\text{DLN}_{\text{abs}}^{2-6}$  for parameter  $z_{\text{qs}2}$  except some  $\text{negl}(\lambda)$  correctness error. Therefore, since  $\mathbb{P}$  is a  $(\theta, m)$ -secure MPC protocol that implements  $\text{DLN}_{\text{sort}}^{2-6}$ , by Definition 1, we can “replace”  $\mathbb{P}[\text{DLN}_{\text{sort}}^{2-6}]$  with  $\mathcal{F}$  for the program family  $\text{DLN}_{\text{abs}}^{2-6}$ , denoted by  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$ . In more detail, the hybrid  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$ -protocol  $\pi_d^{[\text{DLN}_{\text{abs}}^{2-6}]}$  operates similar to  $\pi_d^{[\text{DLN}_{\text{sort}}^{2-6}]}$ , with the difference that instead of interacting with  $\mathbb{P}[\text{DLN}_{\text{sort}}^{2-6}]$ , it does so with  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$ . The functionality  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$  that considers a single user  $U_{\text{in}}$  for receiving input and a single user  $U_{\text{out}}$  for providing output (in real-world, realized by the entry and output servers) and is described as follows:

- Upon receiving ‘start’ from  $\mathcal{S}$ , it sets the status to ‘input’ and initializes two lists  $L_{\text{input}}$  and  $L_{\text{corr}}$  as empty. It forwards the message
- Upon receiving  $(\text{corrupt}, u_i)$  from  $\mathcal{S}$ , it adds  $u_i$  to  $L_{\text{corr}}$ .
- Upon receiving  $(\text{send\_input}, \langle (u_1, a_1), \dots, (u_{\text{act}}, a_{\text{act}}) \rangle)$  from  $U_{\text{in}}$ , it sets  $\mathcal{U}_{\text{act}} = \{1, \dots, \text{act}\}$ . For every  $u_i \in \mathcal{U}$ , if  $u_i \in L_{\text{corr}}$ , then it sends  $(\text{send\_input}, u_i, a_i)$  to  $\mathcal{S}$ . If  $u_i \notin L_{\text{corr}}$ , then it sends (i)  $(\text{send\_input}, u_i, \text{abstain}_i(a_i))$  to  $\mathcal{S}$ , where  $\text{abstain}_i(\cdot)$  is defined in Eq. (1).

- Upon receiving (receive\_input,  $u_i, \tilde{a}_i$ ) from  $\mathcal{S}$ , if (i) the status is ‘input’ and (ii)  $(u_i, \cdot) \notin L_{\text{input}}$ , then if  $u_i \notin L_{\text{corr}}$ , it adds  $(u_i, a_i)$  to  $L_{\text{input}}$ , else it adds  $(u_i, \tilde{a}_i)$  to  $L_{\text{input}}$ .
- Upon receiving (compute,  $z$ ) from  $\mathcal{S}$ , it runs  $\text{DLN}_{\text{abs}}^{2-6}$  with parameter  $z$  and on input  $\langle a_1, \dots, a_{\text{act}} \rangle$ .
- It returns the vector of pairs  $\langle (u_1, b_1), \dots, (u_n, b_{\text{act}}) \rangle$  to  $U_{\text{out}}$ , where  $b_i := (b_i[1], b_i[2])$  is defined as in  $\text{DLN}_{\text{sort}}^{2-6}$ , and also sends  $\langle (u_i, b_i) \rangle_{u_i \in L_{\text{corr}}}$  to  $\mathcal{S}$ .

By the security of  $\mathbb{P}$ , there is an adversary  $\mathcal{S}_{\text{DLN}}^{2-6}$  s.t.

$$\text{EXEC}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{x}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{A}_{\text{DLN}}, \mathbf{x}}^{\pi_d^{\mathbb{P}[\text{DLN}_{\text{sort}}^{2-6}]}}(\lambda) \equiv \text{EXEC}_{\mathcal{A}_{\text{DLN}}, \mathbf{x}}^{\mathbb{P}[\text{DLN}_{\text{sort}}]}(\lambda). \quad (3)$$

We now construct the ideal adversary  $\mathcal{S}_{\text{DLN}}$  for  $\mathcal{A}_{\text{DLN}}$ . Namely,  $\mathcal{S}_{\text{DLN}}$  interacts with  $\mathcal{F}$  for the program family  $\text{DLN}_{\text{abs}}$ , invokes  $\mathcal{S}_{\text{DLN}}^{2-6}$  and simulates an interaction of  $\mathcal{S}_{\text{DLN}}^{2-6}$  with  $\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}$  as follows:

- Upon receiving ‘start’ from  $\mathcal{S}_{\text{DLN}}^{2-6}$ , it forwards ‘start’ to  $\mathcal{F}$  and initializes two lists  $L_{\text{input}}^{2-6}$  and  $L_{\text{corr}}^{2-6}$  as empty, playing the role of  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$ .
- Upon receiving (corrupt,  $u_i$ ) from  $\mathcal{S}_{\text{DLN}}^{2-6}$ , it forwards ‘(corrupt,  $u_i$ )’ to  $\mathcal{F}$  and adds  $u_i$  to  $L_{\text{corr}}^{2-6}$ . In addition, since  $\mathcal{A}_{\text{DLN}}$  corrupts the entry and output servers, then so will  $\mathcal{S}_{\text{DLN}}^{2-6}$  in the simulation of  $\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}$ .
- Upon receiving (send\_input,  $u_i, x_i$ ) from  $\mathcal{F}$  for some  $u_i \in L_{\text{corr}}$ , it adds  $u_i$  to  $\mathcal{U}_{\text{act}}$  and encodes  $x_i$  to the triple  $a_i := (a_i[1], a_i[2], 0)$ .
- Upon receiving (send\_input,  $u_i, Q_i$ ) from  $\mathcal{F}$  or some  $u_i \notin L_{\text{corr}}$ , where  $Q_i$  is the abstain bit for  $u_i$ , if  $Q_i = 1$ , then it adds  $u_i$  to  $\mathcal{U}_{\text{act}}$ , and sets a simulated encoded triple  $a_i := (C, \text{UN}_i, 0)$  as the input for  $u_i$ . Namely, it sets any active honest user to be a dialchecker by default.
- Upon receiving (receive\_input,  $u_i, \hat{a}_i$ ) from  $\mathcal{S}_{\text{DLN}}^{2-6}$ , if  $(u_i, \cdot) \notin L_{\text{input}}$ , then
  - If  $u_i \notin L_{\text{corr}}$ , then it sets  $\tilde{a}_i := a_i$  and adds  $(u_i, \tilde{a}_i)$  to  $L_{\text{input}}^{2-6}$ .
  - If  $u_i \in L_{\text{corr}}$ , then it sets  $\tilde{a}_i := \hat{a}_i$  and adds  $(u_i, \tilde{a}_i)$  to  $L_{\text{input}}^{2-6}$ .

We note that the above operation, captures the adversarial power of  $\mathcal{S}_{\text{DLN}}^{2-6}$ , controlling a part of the user set and passively corrupting the entry server. Namely, the adversary may choose the inputs of the users it controls, but when the round begins (see below), the users’ inputs will be delivered to  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$ .

Next,  $\mathcal{S}_{\text{DLN}}$  decodes  $\hat{a}_i$  to  $\hat{x}_i$  (if decoding is not a dial or dialcheck request for  $u_i$ , then it sets  $\hat{x}_i = \perp$ ) and sends (receive\_input,  $u_i, \hat{x}_i$ ) to  $\mathcal{F}$ .

- Upon receiving (send\_input,  $\langle (u_1, a_1), \dots, (u_{\text{act}}, a_{\text{act}}) \rangle$ ) from  $\mathcal{S}_{\text{DLN}}^{2-6}$ , where the latter acts as  $U_{\text{in}}$  by corrupting the entry server,  $\mathcal{S}_{\text{DLN}}$  plays the role of  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$  with the following modification; it first runs the order check  $H(\text{UN}_1 || \dots || \text{UN}_{\text{act}})$  as the honest MPC servers would do in the real-world setting, and if it fails, then it aborts simulation. Otherwise, it continues normally as dictated by  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$ .
- Upon receiving (compute,  $z$ ) from  $\mathcal{S}_{\text{DLN}}^{2-6}$ , it first sends the message (send\_input,  $\langle (u_1, \tilde{a}_1), \dots, (u_{\text{act}}, \tilde{a}_{\text{act}}) \rangle$ ) on behalf of  $U_{\text{in}}$  to the (simulated)  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$ . Then, it forwards the message (compute,  $z$ ) to  $\mathcal{F}$ .

- Upon receiving the value vector  $\langle (u_1, b_1), \dots, (u_{\text{act}}, b_{\text{act}}) \rangle$  from  $\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]$  on behalf of  $U_{\text{in}}$ , it forwards  $\langle (u_i, b_i) \rangle_{u_i \in L_{\text{corr}}}$  to  $\mathcal{S}_{\text{DLN}}^{2-6}$ .
- It returns the output of  $\mathcal{S}_{\text{DLN}}^{2-6}$ .

By the description of the simulation above, we have that, except some negligible error of the order check that  $\mathcal{S}_{\text{DLN}}$  performs, for input vector  $\mathbf{x}$  and corrupted set  $L_{\text{corr}}$ ,  $\mathcal{S}_{\text{DLN}}$  simulates an execution of  $\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}$  for the input vector  $\tilde{\mathbf{x}} = \langle \tilde{x}_1, \dots, \tilde{x}_n \rangle$  defined as follows:

$$\tilde{x}_i := \begin{cases} (\text{DIALCHECK}, u_i), & \text{if } u_i \notin L_{\text{corr}} \text{ and } x_i \neq \perp \\ x_i, & \text{otherwise} \end{cases}$$

Let  $M$  be the number of users s.t.  $u_i \notin L_{\text{corr}}$  and  $x_i \neq \perp$ . We consider a sequence of “hybrid” inputs  $\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_M$ , where  $\mathbf{h}_0 := \mathbf{x}$  and for  $j = 1, \dots, M$ ,  $\mathbf{h}_j$  derives by replacing the input of the  $j$ -th user  $u_{i_j} \in L_{\text{corr}}$  and  $x_{i_j} \neq \perp$  (arranged w.r.t. some order) in  $\mathbf{h}_{j-1}$  with the pair  $(\text{DIALCHECK}, u_{i_j})$ . Clearly, it holds that  $\mathbf{h}_M = \tilde{\mathbf{x}}$ .

Let  $\text{OUT}_{\mathcal{S}, \tilde{\mathbf{w}}}^{\pi}(\lambda)$  denote the output of adversary  $\mathcal{S}$  when interacting with protocol  $\pi$  on input  $\tilde{\mathbf{w}}$ , parameterized by  $\lambda$ . Assume for the sake of contradiction, that for some  $j \in [M]$ , it holds that

$$\text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{h}_{j-1}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda) \stackrel{c}{\not\approx} \text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{h}_j}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda),$$

and let  $\mathcal{D}$  be a distinguisher for the above random variables. Given  $\mathcal{D}$ , we can construct an algorithm  $\mathcal{B}$  that attacks the authenticated encryption used for message communication in the MCMix architecture. In particular,  $\mathcal{B}$  receives as input a triple of authenticated encryption shares of the input of  $u_{i_j}$  that refer to either (i)  $\mathbf{h}_j$  or (ii)  $\mathbf{h}_{j-1}$ . Observe that by definition, the hybrid inputs  $\mathbf{h}_j$  and  $\mathbf{h}_{j-1}$  differ only for  $u_{i_j}$ . Then,  $\mathcal{B}$  simulates an execution of  $\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}$  in the presence of  $\mathcal{S}_{\text{DLN}}^{2-6}$  and returns the output of  $\mathcal{D}$ , when the latter is given the output of  $\mathcal{S}_{\text{DLN}}^{2-6}$ . Therefore, depending on its input,  $\mathcal{B}$  provides  $\mathcal{D}$  with either an input following  $\text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{h}_j}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda)$  or (ii)  $\text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{h}_{j-1}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda)$ . Thus, by a standard reduction argument,  $\mathcal{B}$  has a non-negligible advantage of distinguishing its input, which contradicts the security of the applied authenticated encryption scheme.

By the above and the transitive property of  $\stackrel{c}{\approx}$  we have that

$$\text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{h}_0}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda) \stackrel{c}{\approx} \text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{h}_M}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda) \Leftrightarrow \text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{x}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda) \stackrel{c}{\approx} \text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \tilde{\mathbf{x}}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda) \quad (4)$$

By the description of  $\mathcal{S}_{\text{DLN}}$  and Eq. (4), we have that

$$\text{OUT}_{\mathcal{S}_{\text{DLN}}, \mathbf{x}}^{\mathcal{F}}(\lambda) \stackrel{c}{\approx} \text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{x}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda). \quad (5)$$

Recall that the transcripts  $\text{EXEC}_{\mathcal{S}_{\text{DLN}}, \mathbf{x}}^{\mathcal{F}}(\lambda)$  and  $\text{EXEC}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{x}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda)$  consist of the inputs and outputs of the users, along with  $\text{OUT}_{\mathcal{S}_{\text{DLN}}, \mathbf{x}}^{\mathcal{F}}(\lambda)$  and  $\text{OUT}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{x}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda)$  respectively. Since the initial input vector is  $\mathbf{x}$  and  $\mathcal{S}_{\text{DLN}}$  forwards to  $\mathcal{F}$  the exact adversarial inputs that  $\mathcal{S}_{\text{DLN}}^{2-6}$  dictates in decoded form, we have that  $\mathcal{F}$  and  $\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}$  are going to be executed on the same inputs. By the security of the ID-KA scheme used for the dead drop computation the outputs of the users in  $\text{EXEC}_{\mathcal{S}_{\text{DLN}}^{2-6}, \mathbf{x}}^{\pi_d^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda)$  include pseudorandom

values in  $\mathbb{Z}_q$ , where as in  $\text{EXEC}_{\mathcal{S}_{\text{DLN},\mathbf{x}}}^{\mathcal{F}}(\lambda)$  the respective values are truly random from the same group. By the latter fact and Eq. 5, we have that

$$\text{EXEC}_{\mathcal{S}_{\text{DLN},\mathbf{x}}}^{\mathcal{F}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{S}_{\text{DLN},\mathbf{x}}^{2-6}}^{\pi_c^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}}(\lambda). \quad (6)$$

Finally, by the transitive property of  $\stackrel{c}{\approx}$  and Eq. (3) and (6), we conclude that

$$\text{EXEC}_{\mathcal{S}_{\text{DLN},\mathbf{x}}}^{\mathcal{F}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{A}_{\text{DLN},\mathbf{x}}}^{\mathbb{P}[\text{DLN}_{\text{sort}}]}(\lambda). \quad (7)$$

2). *Security of the Conversation Protocol  $\text{CNV}_{\text{sort}}$* : The proof for the security of  $\text{CNV}_{\text{sort}}$  is similar. Briefly, let  $\mathbb{P}[\text{CNV}_{\text{sort}}]$  denote the implementation of  $\text{CNV}_{\text{sort}}$  over  $\mathbb{P}$ , that is expressed as  $\pi_c^{\mathbb{P}[\text{CNV}_{\text{sort}}^{2-6}]}$ , where  $\pi_c$  is a protocol completing the non-MPC part (Step 1) of  $\text{CNV}_{\text{sort}}$  run in the users' client and the communication handled by entry and output servers.

By Theorem 2, we have that  $\text{CNV}_{\text{sort}}^{2-6}$  realizes  $\text{CNV}_{\text{abs}}^{2-6}$  for parameter  $z_{\text{qs1}}$ . As a result the security of  $\mathbb{P}$  implies that for any passive adversary  $\mathcal{A}_{\text{CNV}}$  corrupting  $\theta$ -out-of- $m$  servers, there is an adversary  $\mathcal{S}_{\text{CNV}}^{2-6}$  such that for every input vector  $\mathbf{x}$ ,

$$\text{EXEC}_{\mathcal{S}_{\text{CNV},\mathbf{x}}^{2-6}}^{\pi_c^{\mathcal{F}[\text{CNV}_{\text{abs}}^{2-6}]}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{A}_{\text{CNV},\mathbf{x}}}^{\pi_c^{\mathbb{P}[\text{CNV}_{\text{sort}}^{2-6}]}}(\lambda) \equiv \text{EXEC}_{\mathcal{A}_{\text{CNV},\mathbf{x}}}^{\mathbb{P}[\text{CNV}_{\text{sort}}]}(\lambda).$$

Subsequently, we can construct an ideal adversary  $\mathcal{S}_{\text{CNV}}$  that interacts with  $\mathcal{F}$  for the program family  $\text{CNV}_{\text{abs}}$ , invokes  $\mathcal{S}_{\text{CNV}}^{2-6}$  and simulates an interaction of  $\mathcal{S}_{\text{CNV}}^{2-6}$  with  $\pi_c^{\mathcal{F}[\text{DLN}_{\text{abs}}^{2-6}]}$ . As previously, based on the security of the authenticated encryption scheme, we can show that

$$\text{EXEC}_{\mathcal{S}_{\text{CNV},\mathbf{x}}}^{\mathcal{F}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{S}_{\text{CNV},\mathbf{x}}^{2-6}}^{\pi_c^{\mathcal{F}[\text{CNV}_{\text{abs}}^{2-6}]}}(\lambda).$$

Finally, by the transitive property of  $\stackrel{c}{\approx}$ , we conclude that

$$\text{EXEC}_{\mathcal{S}_{\text{CNV},\mathbf{x}}}^{\mathcal{F}}(\lambda) \stackrel{c}{\approx} \text{EXEC}_{\mathcal{A}_{\text{CNV},\mathbf{x}}}^{\mathbb{P}[\text{CNV}_{\text{sort}}]}(\lambda). \quad (8)$$

Eq. (7) and Eq. (8) imply the security of MCMix, which completes the proof.  $\square$

**Remark 5 (On forward security of MCMix).** MCMix in its current form does not offer forward security. Nevertheless, it is possible to provide forward security as follows. First, clients could refresh their exchanged keys with the servers in regular time intervals, e.g., once a day. Alternatively, to avoid interaction, forward secure encryption can be used, e.g., see [BY03]. With respect to the dead drop calculation we can obtain forward security by applying our second ID-KA construction with forward secrecy (cf. Section 2.3, Construction 2). The additional communication cost to the Dialing protocol would be one extra random group element per user as now the active inputs  $x_1, \dots, x_n$  for dialing need to be used for the first round of the exchange; they are of the form of  $(\text{DIAL}, u_i, u_j, r_i)$  and  $(\text{DIALCHECK}, u_i, r_j)$ , where  $r_i, r_j$  are random elements from the ID-KA cyclic group. Sorting would still be executed on the users' usernames and the wire IDs as before thus incurring no additional overhead. We omit further details.

**Remark 6 (End-to-end encryption of MCMix).** In our architecture, we made no assumptions about a pre-existing PKI such that each user could securely obtain each other user's public key. Thus, in order to enable encryption on the client side, users need to exchange public keys. To make this exchange as lightweight as possible, we turned to ID-based cryptography where it is sufficient that the dialer knows the dialer's 64bit username (the way this knowledge is established is outside the scope of our construction). By the nature of our ID-based solution, if at least one of the KGCs (MPC servers) is honest, then

an active adversary controlling the other KGCs can not infer any information about messages encrypted under an ID-based key. However, if all KGCs are corrupted, then even a passive adversary can reconstruct all users’ secret keys and thus decrypt any transferred ID-based ciphertext. Consequently, when instantiated as above, MCMix supports end-to-end encryption security as long as a single MPC server remains honest.

Following an alternative approach, we could build upon a public-key setting where the MPC servers realize a distributed authority responsible for forwarding the correct public keys to the users. In more detail, during registration, each user  $u_i$  generates a private and public key pair  $(sk_i, pk_i)$  and uploads  $pk_i$  to the MPC system. During the Dialing protocol, if  $u_i$  wishes to dial user  $u_j$ , then it sends a  $(DIAL, u_i, u_j)$  request encoded as follows  $(pk_i, UN_j, 0, 0)$ , while if  $u_j$  dialchecks then it sends a  $(DIALCHECK, u_j)$  request now encoded as  $(C, UN_j, pk_j, 0)$ . In the end of the MPC protocol execution  $u_j$  would receive  $pk_i$  and  $u_i$  will receive  $pk_j$  and they will complete the dialing protocol by a standard Diffie Hellman key exchange. Note that this will require a “swapping” public-keys step to be executed as part of the MPC protocol which can be achieved as in the case of the conversation protocol. In this way MCMix can support end-to-end encryption security even if all servers are passively corrupted.

## 8 Implementation and Benchmarking

We implemented a prototype of our system using the Sharemind platform and performed extensive evaluation.

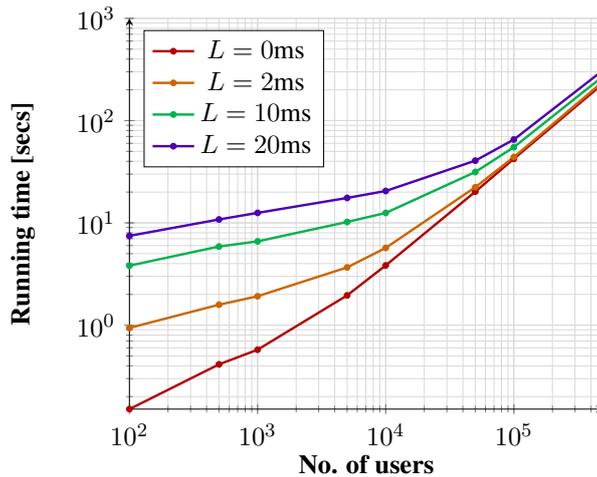


Figure 7: Running time in secs of the Dialing protocol implementation for a number of  $n = 100, 500, 1K, 5K, 10K, 50K, 100K, 500K$  users and latency  $L = 0, 2, 10, 20$  ms. The benchmarks were run with username size 8 Bytes and 1 Gbps network bandwidth.

### 8.1 Experiment setting

Benchmarks were run on a cluster of three machines with point-to-point 1 Gbps network connections using various profiles for network latency aiming to simulate WAN behavior. Each machine has a 12-core 3 GHz Hyper-Threading CPU and 48 GB of RAM. However, even though the hardware supports it, Sharemind MPC protocols are not optimized to use multiple CPU cores or the network layer in a parallel manner. The servers running Sharemind employ only 2 cores, one for executing the computations and another for pseudo-random number generation. To simulate real-world environment, we use the `tc` tool to manipulate the operating system’s network traffic control settings. This tool is used to cap the available network bandwidth and introduce communication latency by adding round-trip delay (ping).

## 8.2 Dialing protocol

We benchmarked our dialing protocol for various numbers of users and various latency values. The results are presented in Fig. 7. As we can see, the dialing protocol has a runtime for each round of around one minute for 100,000 users and around 300 seconds for 500,000 users, considering the worst case of 20 ms of latency. The latter value might still be considered acceptable for some settings, as dialing rounds need not be executed very often. Another interesting observation is that the effect of latency diminishes as the number of users increases, due to the fact that the number of communication rounds of our algorithm scales logarithmically to the number of inputs. This in turn happens because Quicksort needs  $\mathcal{O}(\log(n))$  steps to sort  $n$  inputs when executed in parallel. The vectorized nature of our implementation succeeds in taking advantage of the parallelizable nature of the algorithm. The time a user needs to encode her request and send it, as well as the time required by each MPC server to decrypt the requests it received, have no effect on the per round runtime of our system. This is because these operations are performed in a pipelined fashion. This means that the encoding, encryption and decryption of the requests for round  $r + 1$  takes place while the MPC servers perform the computations for round  $r$ . In the dialing protocol this is acceptable as a user’s intent on whether to dial or perform a dial check might not depend on the output of the previous dialing round.

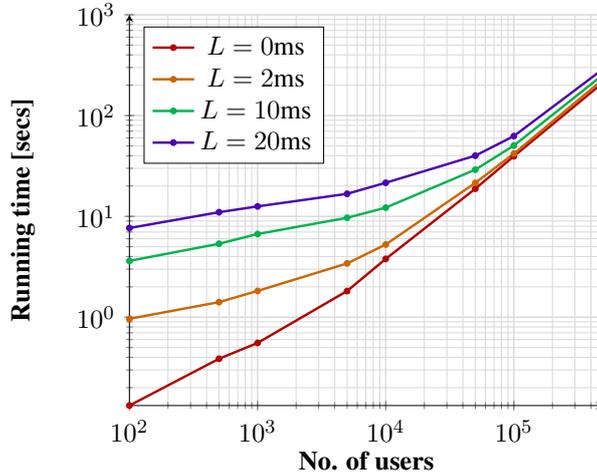


Figure 8: Running time in secs of the Conversation protocol implementation for a number of  $n = 100, 500, 1\text{K}, 5\text{K}, 10\text{K}, 50\text{K}, 100\text{K}, 500\text{K}$  users and latency  $L = 0, 2, 10, 20$  ms. The benchmarks were run with message size 8 Bytes and 1 Gbps network bandwidth.

## 8.3 Conversation protocol

For the conversation protocol we provide extensive benchmarks considering the number of users, the latency of the network, as well as the message size. In Fig. 8, we can see that the running time of the conversation protocol with a very small message size of 8 Bytes (B) is similar to the running time of the dialing protocol. That is, the system can serve 100,000 users in around one minute for maximum latency of 20 ms. Again, we see that latency is a minor performance factor for a large number of users. This fact enables us to claim that our system will have similar running times even with greater latency values.

In Fig. 9, we consider how the message size affects performance. We have benchmarked various message sizes ranging from 8 B to 1 KB messages. No artificial latency has been injected for these experiments. We see that message size affects performance in a significant way as opposed to latency, but the system can still support anonymity sets of tens of thousands of users even with 1KB messages and certainly SMS long messages for hundreds of thousands.

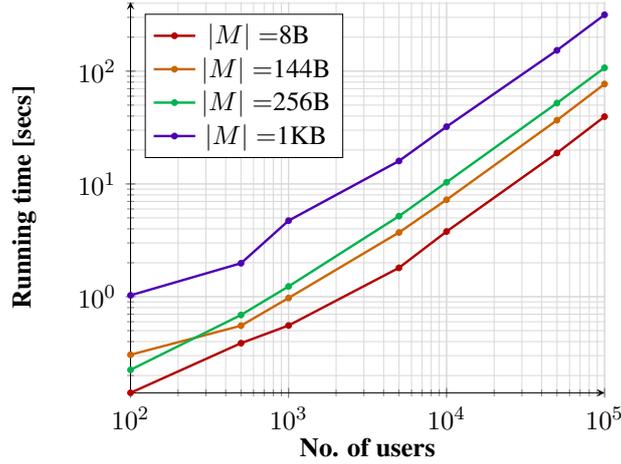


Figure 9: Running time in secs of the Conversation protocol implementation for a number of  $n = 100, 500, 1K, 5K, 10K, 50K, 100K$  users and message size  $|M| = 8, 144, 256, 1K$  Bytes. The benchmarks were run with no latency and 1 Gbps network bandwidth.

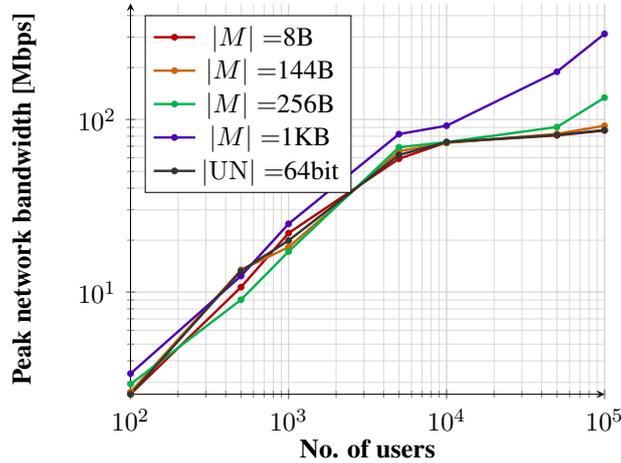


Figure 10: The peak network bandwidth consumption in Mbps during the Dialing protocol for usernames (UNs) of 64bits and the Conversation protocol for message size  $|M| = 8B, 144B, 256B, 1KB$ , given a number of  $n = 100, 500, 1K, 5K, 10K, 50K, 100K$  users. The benchmarks were run with no latency and 1 Gbps network bandwidth.

Finally, in Fig. 10, we provide the peak network bandwidth consumption during the Dialing and Conversation protocols. We note that the total bandwidth is shown, i.e. bytes sent and received and to both other computing nodes. We observe that in both protocols the bandwidth consumption remains at a low level of less than 100Mbps for the Dialing protocol for (usernames of 64bits) as well as the Conversation protocol for messages of up to SMS size. For bigger message sizes and 100,000 users, we get that the total consumption is roughly 150Mbps and 300Mbps for messages of 256B and 1KB respectively, which can be realistic for a large scale setting.

## 9 Client load and adoption incentives

Anonymous communication systems critically rely on having adequately large anonymity sets to be effective. In other words: “Anonymity Loves Company” [DM06], and the usability aspects of anonymous communication systems should be an important design consideration. MCMix strives to offer strong adoption incentives by offering strong security, while minimizing the computation and communication load on the client side.

**Computation load:** For the Dialing protocol, each client performs an ID-KA operation (see section 2) to compute the dead drop value, plus a few symmetric operations to encrypt/decrypt the shares. The Key Exchange operation consists of a single bilinear pairing computation. The computation burden of pairings is an ongoing research topic and recent work [AKL<sup>+</sup>11] using elliptic curves, has offered significant improvements, allowing pairing to be computed in under 2 million cpu cycles in commodity desktop computers. In [AGH15], symmetric pairing time is estimated at 14.9 ms running on a commodity device, or around three times the time needed for a modular exponentiation in the corresponding cyclic group. For the Conversation protocol, the load is very low, with the client symmetrically encrypting the shares to the MPC servers and being additionally able to encrypt/decrypt the message sent with the shared ID-KA key agreed in the dialing phase. Assuming Dialing and Conversation rounds are executed once per minute, the computational cost on the client is expected to be minimal and easily handled by a mobile device.

**Communication load:** For the Dialing protocol, each client sends 3 encrypted shares containing 128 bits of information, plus her 64 bit username, one to each server. Given that the fixed block size for AES is 128 bits, we can assume each packet payload is 32 B in size. As output, each client receives three shares carrying 64 bits of information (the username of a possible caller). Taking into account the block size of AES, this results in packets with a payload of 128 bits, or 16 B. We additionally assume a TCP header of 20 B [Pos81b], along with an IP header again of 20 B [Pos81a], resulting in a single packet of 72 B and 56 B for the input and output shares of the Dialing protocol respectively. In addition a SHA-256 HMAC (32 B) is assumed as the authentication mechanism. Therefore, in total, a client needs  $3 \cdot (72 + 32) + 3 \cdot (56 + 32) = 576$  B of bandwidth for each dialing round. Assuming dialing rounds every 1 minute, the bandwidth cost for each client would be in the order of 24 MB per month. For the Conversation protocol, with a variable message size  $|M|$  B, each client sends tuples of the form  $(t_i, m_i)$  and receives some  $m_j$ . Therefore, assuming a joint TCP/IP header overhead of 40 B and 32 B HMAC, the total bandwidth cost per round of conversation is  $3 \cdot (8 + |M| + 40 + 32) + 3 \cdot (|M| + 40 + 32) = 6 \cdot |M| + 456$  B. The message size is padded so as to be a multiple of 128 bits (AES block size). Taking as an example SMS length messages of 140 B (padded to 144 B due to AES), each conversation round would cost 1320 B, resulting in a monthly bandwidth need of around 54 MB when a client is constantly online. An overview of the total monthly bandwidth costs of the clients is available in Table 1. Dialing and conversation rounds are assumed to be executed every one minute (simultaneously), and MB conversions are base 2. The theoretical analysis of the computational and communication overhead

$ M $ (B)	bandwidth per month (MB)
8	47
144	78
256	106
1K	296

Table 1: Communication costs of clients (Dialing and Conversation combined) w.r.t. message size.

of our system shows, that it is lightweight on the client side and the bandwidth needs of a device to be constantly connected are in the range of tens of MB per month, which we consider easily manageable.

The low communication overhead of MCMix is due to the fact that each client only sends/receives one share to/from each MPC server per round. As a result of the relatively small costs, users can be expected to remain online even if they are not actively using the service, resulting in bigger anonymity sets. While we expect MCMix to be practical for mobile users, further experiments may be needed to compute actual battery consumption and bandwidth usage in a real-world setting. Our theoretical results presented above do not take into account factors like packet loss and network-level packet padding but serve as a guideline.

## 10 Parallelizing the Conversation protocol

As discussed in previous sections, our protocols are provably secure assuming a secure MPC framework and are also scalable enough to support hundreds of thousands of users. While these anonymity sets can accommodate a lot of use cases, we recognize the need for anonymity systems to offer as large an anonymity set as possible. Therefore, we propose a technique that leads to an even more scalable system, by describing a parallel realization of the Conversation protocol, as this is the latency-critical component of our system. The Dialing protocol can be executed independently of the Conversation one, and in much longer time intervals, e.g. every five minutes. Therefore, the implementation of Dialing on a single MPC instance can cover very large anonymity sets, e.g. 500,000 users as seen in Fig. 7.

### 10.1 General Idea

Our main challenge is to come up with a protocol that can run in different MPC instances (islands) in parallel, with minimal communication between those islands, and that achieves very strong privacy. Additionally, the anonymity set should be the whole client population. The problem of anonymous communication, where two users who want to communicate may submit their messages to different islands and still expect to communicate with perfect correctness, while leaking no information at all, is hard to parallelize. In our approach, we choose to maintain the strongest possible privacy standards. As a result, in our parallelized version of MCMix, we relax our quality of service (qos) guarantees. That is, in each round, an adjustably small number of requests that would have been served when using the algorithm of Fig. 6, will fail to do so; a message exchange that would have been performed by our original algorithm will not take place, and the affected clients will have to resend their messages in the next round. The probability of this phenomenon can be made arbitrarily small in the expense of performance, as described later in this section. As evident by the algorithmic representation of our two protocols, the integral part of their function is matching equal values in pairs and performing an action on these pairs (a swap). In this section, we introduce a parallelizable approach to performing this action that benefits from the fact that the values in question (dead drops) are uniformly distributed, being the output of a hash function.

The parallel protocol will not have the correctness property of the non-parallel one. The system will (with high probability) never produce a correct output, meaning that there will exist a very small portion of the users that will need to resend their messages in the next round.

In our approach, requests are split obliviously between MPC islands based on the fact that equal dead drop values are likely to be located at roughly the same indexes of different arrays after sorting, considering these values are uniformly distributed. In our illustrations, we use two MPC islands, each one consisting of an MPC system (e.g., a three-server Sharemind implementation), but the method can be applied to any number of islands. The notion of an island is a logical and not a geographical characterization. Islands are MPC systems with each of their servers ideally spread around the world. An MPC island and its clients interact in the same way we have described in the previous sections, i.e. it is transparent to the clients whether or not a parallelized protocol is running.

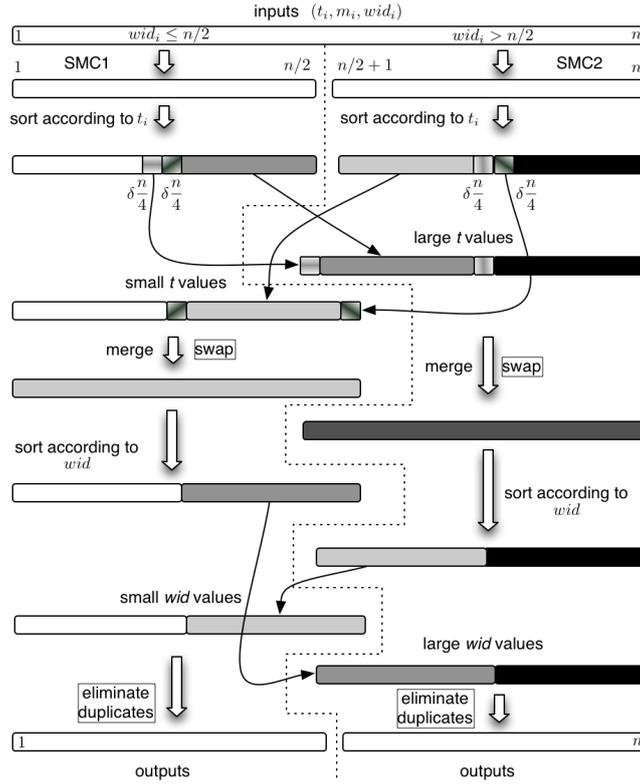


Figure 11: Parallel operation of two MPC islands performing the Conversation Protocol.

## 10.2 Parallelizing the Conversation protocol

In Fig. 11, we can see how we can combine two MPC islands. In our illustration we use two MPC islands, and assign half of the incoming requests ( $\frac{n}{2}$ ) of the form  $(t_i, m_i, wid_i)$  to each of them. Without generality loss, the first island gets requests from clients with  $wid \in \{1, \dots, \frac{n}{2}\}$  and the second with  $wid \in \{\frac{n}{2} + 1, \dots, n\}$ .

The parallelized protocol is executed via the following steps:

- 1. Independent dead drop sort:** Each island independently sorts its requests obliviously, according to their  $t$  coordinate. Having assumed a uniform distribution of those values, small (respectively high) values of  $t$  will end up at the lower (upper) half of the sorted sequences.
- 2. First inter-island communication:** The first island keeps the lower half (plus  $\delta \frac{n}{4}$ ) of its sorted requests and receives the lower half (plus  $\delta \frac{n}{4}$ ) of the sorted requests of the second island. The second island does the opposite, keeping the upper half of its requests and receiving the upper half of the first island's requests. Due to the fact that these  $t$  values are uniformly distributed, equal pairs will end up in the same island with a high probability. Practically, this transmission of data is made on a peer level between each of the servers of the two islands. In detail, the first server of the first island communicates through a secure channel with the first server of the second island etc. The additional ( $d = \delta \frac{n}{2}$ ) requests, apart from the halves assigned to each island, serve the purpose of calibrating the quality of service parameter of the protocol. The bigger the value of  $d$ , the less likely it is that two requests with the same  $t$  values will end up in different islands and communication will not take place.
- 3. Merge sub-sequences:** Each island obliviously merges the two sorted sequences according to their  $t$  values. The result is a fully sorted sequence for each island.
- 4. Exchange messages:** Each island performs the exchange of messages as described in Step 4 of

$CNV_{\text{sort}}$  in Fig. 6. After this step, the message exchange has been performed and the messages must reach their intended recipients, based on their wire ID. The  $t$  values of each tuple are ignored from this step onwards.

**5. Independent wire ID sort:** Each island obviously sorts its requests independently according to their  $wid$  coordinate. After this step, it is guaranteed that the first  $\frac{n}{4} + \delta\frac{n}{4}$  requests of the first island originated from the first island and the rest  $\frac{n}{4} + \delta\frac{n}{4}$  from the second. The same is true for the second island.

**6. Second inter-island communication:** Each message request is sent back to the island it originated from, again using direct secure channels between the respective servers of each island.

**7. Merge sub-sequences:** Each island obviously merges the requests designated for it, according to their  $wid$ , and ends up with  $\frac{n}{2} + d$  requests, some of which have duplicate wire IDs. The duplicate requests must then be combined in a meaningful way before proceeding with the algorithm.

**8. Eliminate duplicates:** (Obsolete if  $d = 0$ ) For a pair of messages  $m_1$  and  $m_2$  that are part of tuples with the same wire ID, one of the following must be true:

- $m_1 = m_2$ , either both of them are carrying the result of a successful message exchange or none of them.
- $m_1 \neq m_2$ , that is one of the two messages carries the result of a successful message exchange and the other an original message that was not exchanged.

So to eliminate duplicates, we apply the algorithm of Fig. 12. This algorithm eliminates duplicates by combining two requests of the form  $(m_1, wid)$ ,  $(m_2, wid)$  with the same wire ID, so that if one of them carries a message that was the result of an exchange operation (LSB=1, see step 4 of  $CNV_{\text{sort}}$ ), then this is the message that makes it to the receiver. At the end of this combination we have the valid message form a tuple of the form  $(m, wid)$  and the invalid one a tuple  $(m', 0)$ , which is discarded at the end.

### Eliminate Duplicates Algorithm

**Input:** a sequence of  $n + d$  tuples  $\langle a_1, \dots, a_{n+d} \rangle = \langle (m_1, wid_1), (m_2, wid_2), \dots, (m_n, wid_n) \rangle$ , which is the sorted output (according to the wire IDs) of the merge step of the algorithm and contains  $d$  duplicates.

**Output:** a sequence  $b$  of size  $n$ , that is the output of the protocol.

---

For each  $i \leftarrow 1, \dots, n + d - 1$

**if**  $a_i[2] = a_{i+1}[2]$  **then**

**if**  $a_i[1] \bmod 2 = 0$  **then**

$a_i[1] = a_{i+1}[1]$

**end if**

$a_{i+1}[1] = 0$

**end if**

Sort sequence  $a$  according to  $a_i[2]$  using Quicksort. Sorted sequence is  $a'$ .

**return** sequence  $b = \langle a'_{d+1}, \dots, a'_n \rangle$

Figure 12: Eliminate Duplicates Algorithm used in step 8 of the parallel conversation protocol

**9. Independent wire ID sort:** (Obsolete if  $d = 0$ ) Finally, each island obviously sorts its requests according to their wire IDs. Then the valid messages are forwarded to their recipients.

**Remark 7.** The combination of two islands, as explained above, can be generalized to the combination of  $S$  islands with each one handling  $\frac{n}{S}$  of the requests. The range of the possible values for  $t$  is therefore also implicitly partitioned in  $S$ . Quality of service may decline with increased island numbers but it can be controlled and is predicted to remain at a very high level. The work done by each island is roughly double the work an MPC system processing the same amount of requests would do. Such an overhead may sound big, but it is actually reasonable, as it is independent of the number of islands that are used. As a consequence of the above, using 2 islands, as illustrated in our figure, will yield no extra performance, and is expected to be even worse than the non-parallel version of our protocol. However, by having 10 islands we would theoretically have a performance gain that would allow us to support anonymity sets 5 times larger than the ones of a single MPC system.

### 10.3 Quality of Service Analysis

As mentioned before, for parallel MCMix, there is a (non negligible) probability that a user's message will not get exchanged, even if there exists a request with a matching dead drop. An analysis of the probability of such an event taking place in the two island case, follows:

Let  $n$  be an even number, representing the number of users, and  $\mathcal{C}$  an arbitrary set of disjoint subsets from  $\binom{[n]}{2}$ , representing the pairs of users currently in conversation. Consider the following probabilistic procedure of assigning rendezvous points to the users in each round: for each  $i \in [n]$ , pick a random value  $v_i \leftarrow \{0, 1\}^\lambda$  and if it happens that  $i \in P = \{i, j\} \in \mathcal{C}$  such that  $t_j$  is defined already, set  $t_i = t_j$ ; Else, set  $t_i = v_i$ . Consider now the vector  $\langle t_1, \dots, t_n \rangle$  and sort it, resulting to the vector  $\langle t'_1, \dots, t'_n \rangle$ . Define the event  $\text{BAD}_\delta$  to be the event that  $\text{MSB}(t'_i) = 1$  for some  $i \leq (1 - \delta)n/2$  where  $\delta \in (0, 1)$  is a parameter, or that  $\text{MSB}(t'_i) = 0$  for some  $i \geq (1 + \delta)n/2$ .

If  $\mathcal{C} = \emptyset$ , it holds that the most significant bit is uniformly distributed over  $\{0, 1\}$  in each draw  $t_i$  and it follows that the mean of number of times it will be selected to be 1 is  $n/2$ . We recall the two-sided Chernoff bound  $\Pr[|X - \mu| \leq \delta\mu] \leq 2 \exp(-\delta^2\mu/3)$  where  $X$  is the Binomial distribution with mean  $\mu$  and  $\delta \in (0, 1)$ . Consider  $X_i$  to be equal to  $\text{MSB}(t_i) = 1$  and we let  $X = \sum_{i=1}^n X_i$ . Observe that indeed  $X$  is following the Binomial distribution with mean  $\mu = n/2$  and that the event  $\text{BAD}_\delta$  can only happen if the number of  $t_i$ 's with  $\text{MSB}(t_i) = 1$  deviate by a factor  $(1 - \delta)$  below or  $(1 + \delta)$  above the mean. It follows immediately that  $\Pr[\text{BAD}_\delta] \leq 2 \exp(-\delta^2n/6)$ .

In the general case, for arbitrary  $\mathcal{C}$ , observe that the most significant bit for each draw  $t_i$  is uniformly selected unless it happens that  $\{i, j\} \in \mathcal{C}$ . It follows we have  $n' = n - |\mathcal{C}|$  draws, where  $0 \leq |\mathcal{C}| \leq n/2$  since the elements of  $\mathcal{C}$  are subsets of size 2 that are disjoint. The  $i$ -th draw selects element  $t_{f(i)}$  where  $f$  is a mapping from  $[n']$  to a subset of equal size in  $[n]$  that drops the largest element from each conversation pair. We denote by  $P$  the set of all elements of  $[n']$  so that  $f(i)$  participates in a conversation in  $\mathcal{C}$ . We now define the random variable  $Y_i$  as the most significant bit of the  $i$ -th draw among the  $n'$  ones we perform and we let  $Y = \sum_{i=1}^{n'} c_i \cdot Y_i$  where  $c_i = 2$  if  $i \in P$  and  $c_i = 1$  otherwise. Observe that if the event  $\text{BAD}_\delta$  happens it should be that  $Y \notin [(1 - \delta)n/2, (1 + \delta)n/2]$ . Note that by linearity of expectation it holds that  $E[Y] = n/2$  hence the mean, compared to the previous case, has not changed. It is easy to extend the Chernoff tail bound to a tail bound and obtain a tail bound for  $Y$  that will be exponentially decreasing with  $n'$  (alternatively one may use the Hoeffding bound). Specifically we can prove the following.

**Proposition 1.** *For any  $n \in \mathbb{N}$ , any  $\delta \in (0, 1)$  and any set  $\mathcal{C}$  of disjoint subsets of cardinality 2 from  $[n]$ , it holds that  $\Pr[\text{BAD}_\delta] \leq 2 \exp(-\delta^2n/12)$ .*

So, for instance, for a total of ten thousand users and two servers and with  $\delta = 0.08$ , the probability that someone will not be served is less than one percent.

However, this is an upper bound and in practice the quality of service is much better. To better assess it we ran experiments on the Sage platform [Sag16] for  $n = 100,000$  users and  $S = 10, 20$

servers, with each experiment run 200 times. These parameters were chosen as they express a possible usage scenario. An assumption made when running the experiments was that half of the clients were communicating with someone and half were idle. This parameter, however did not seem to influence the results in a substantial way. The quality of service for a 10 island system with each island processing 10000 requests is very high (96%) even when no extra requests are taken. With an overhead of 10% (11000 requests per server), the possibility of even a single failure is very close to zero. The 20 island scenario naturally lags behind in the quality of service when no extra requests are taken, beginning at 88%. Nearly perfect service is achieved with an overhead of 40%, that is 7000 requests per server compared to the original 5000.

#### 10.4 Performance of the parallelized protocol

Considering the fact that we did not have access to a great number of physical machines, in order to run the parallelized protocol with a variety of island numbers, we ran the parallel algorithm on a single island for different user numbers and then extrapolated to give predictions for a real multi-island implementation. Except of the running time of the MPC that we measured, we also added the communication time calculated by assuming commodity 100Mbps connections between the islands. Because inter-island communication consists of only two rounds of transmitting comparatively large chunks of data, the latency of the connection becomes irrelevant. In both inter-island communication rounds, each party sends and receives in total  $n/m \cdot (m-1)/m$  elements to/from other parties, where  $n$  is the number of messages and  $m$  is the number of islands. Also, we have not added any overhead for symmetric encryption between the islands, as even a commodity laptop can keep up with encrypting and decrypting data at a rate of 100Mbps. Taking into account the above, we consider the results presented in Fig. 13 closely depict reality for inter-island communication bandwidth 1Gbps and 100Mbps respectively.

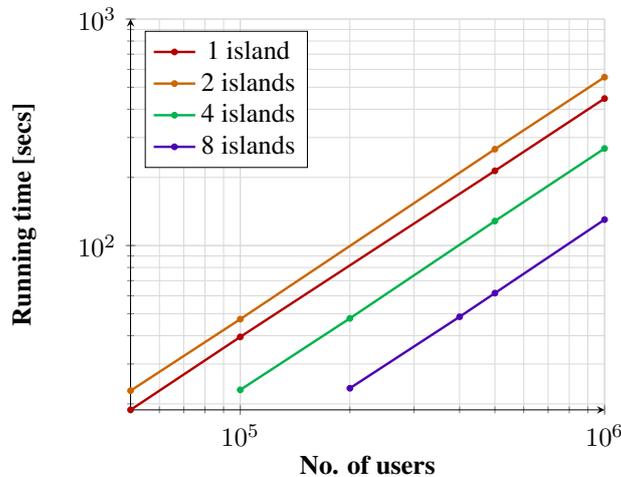


Figure 13: Running time in secs of the Conversation protocol implemented in 1,2,4,8 island setting. The benchmarks were run with no latency, 1Gbps network bandwidth (intra-island) and 64 bit message size. Bandwidth between the islands was modeled at 100Mbps.

From the results of Fig. 13, we can see that as predicted (see Remark 7), employing our system over 2 islands does not provide any additional performance. However, when using 4 or more islands, our parallelization technique yields significant rewards. In the case of 8 islands, the system can support an anonymity set of 500,000 users with a latency of 60 seconds. We expect this trend in performance to continue for even more than 8 islands, thus enabling even larger anonymity sets.

**Remark 8.** A message size of only 64 bits, as benchmarked in this section is not representative of a

realistic deployment. However, the results presented in this section aim at highlighting the scalability of the system

**Remark 9.** An important note is that for the parallel system to be secure, assuming the use of Sharemind as the MPC platform, a non-colluding server majority is needed in every island. The proposed deployment scenario is one in which a single entity, e.g. a university, operates a number of servers that are located in different locations (campuses), which are only connected with commodity internet bandwidths, and one of them is part of each island. In this scenario, the trust model requires that no 2 entities collude to break the privacy of the system.

## 11 Related Work and Comparison

This section attempts to place our work in relation to the state of the art in the expanding field of anonymity-preserving communication systems.

First, regarding Onion-routing based approaches, like POND [Lan15] which uses the Tor network [DMS04], we emphasize that they do not fit the model of a global adversary who can easily defeat them, see e.g., [JWJ<sup>+</sup>13]. Systems that attempt to defeat global adversaries operate in rounds and expect each online user to send encrypted messages in each round. Furthermore, our interpretation of anonymous messaging is one of unobservable bilateral communication. Therefore, unilateral shuffling mechanisms based on mixnets or recent MPC constructions [MSZ15] do not satisfy our application scenario.

Our work is most closely related to the Vuvuzela system [VDHLZZ15] that uses mixnets in addition to dummy messages, to add noise and achieve a differentially private (cf. [Dwo06]) solution to anonymous messaging. By definition, differential privacy protects users as individuals and also allows for some (albeit small) leakage to an observer and thus it is weaker than the simulation-based privacy that we achieve. For example, when all users talk to each other compared to when no user is talking to anyone is completely distinguishable in Vuvuzela, but indistinguishable for MCMix that does not leak any metadata at all. Furthermore, Vuvuzela puts a burden on the client side that requires to finish the dial protocol by downloading a substantial amount of user data (or losing substantially in terms privacy); note that using Bloom filters as described in [LZ16] can help in making this a one time cost. Another drawback of this system is that it cannot scale down in a tight way, due to the burden imposed by the added noise that needs to be always added to maintain acceptable privacy guarantees. On the up side, the system has good architecture and is extremely scalable to millions of users under the assumption of a single honest server, whereas (non-parallelized) MCMix can scale to 100,000 users with similar latency and assuming an honest server majority. However, our parallelized MPC approach can reach that level of performance and in any case, we anticipate that further advances in secure MPC protocols can improve performance substantially even in the non-parallelized version.

Riffle [KLDF15], uses hybrid mixnets and private information retrieval (PIR, [CKGS98]) techniques to implement anonymous messaging. It offers good privacy guarantees, but unlike MCMix and Vuvuzela, it can not handle network churn. During the setup phase of the protocol, client keys are verifiably shuffled by a mixnet. During each communication phase, the same permutations as the ones established in the setup phase are applied to the clients' authenticated messages by the mix servers. As a result of this setup, a single client momentarily leaving or entering the system would require to re-run the expensive setup phase of the protocol.

cMix [CJK<sup>+</sup>16] introduces a mixnet design that can shuffle messages faster than previous work by avoiding public key operations in the real-time phase. cMix provides sender anonymity, yet it may leak the number of messages received by each user, exhibiting a similar security performance as Vuvuzela's dialing protocol.

Dissent [CGF10, WCGFJ12] is based on DC-nets and achieves anonymity sets up to a few thousand users, in an anonymous broadcasting scenario. Riposte [CGBM15] uses PIR techniques to implement a

distributed database that users can anonymously write and read from, assuming no two servers collude (in the efficient scheme). Specifically, the authors implement the write stage on the database as a “reverse” PIR, where a client spreads suitable information for writing in the database. Subsequently, when used for messaging, users can read using PIR from the position in the database that the sender wrote the message (which can be a random position calculated from key information available to the users). Riposte can scale to millions of users but it requires many hours to perform a complete operation; a significant bottleneck is the write-operation that requires  $O(\sqrt{L})$  client communication for an  $L$ -long database which is proportional to the number of users. In contrast, in our system, client bandwidth is minimal, i.e. a single message per server is sent by each user. Additionally, the application scenario is more related to that of Dissent, rather than ours, i.e. anonymous broadcasting, instead of private point to point message exchange, as the authors specify that their approach is suitable “for latency-tolerant workloads with many more readers than writers”. Finally, our technical approach is very different compared to Riposte, as Riposte uses MPC techniques only to detect and exclude malformed client requests, while MCMix offers a native MPC solution for the complete messaging functionality.

BAR [KCB17] uses a “broadcast to all” approach to achieve perfect privacy. A central untrusted server receives all messages in each round and then broadcasts them to all participants. This approach induces a very large communication overhead and therefore anonymity sets are limited to hundreds of users. Pung [AS16] is a system that like BAR operates on fully untrusted setting, while it uses state-of-the-art PIR techniques and smart database organization to scale to a much larger number of users. However, Pung can only implement the equivalent of our conversation functionality and not the dialing functionality, and exhibits substantial client load.

## 12 Conclusions

We presented MCMix the first anonymous real-time point-to-point messaging system that can scale to hundreds of thousands of users, while maintaining security at cryptographic standards. To achieve this, we utilized MPC to securely implement oblivious sorting for (i) establishing the rendezvous points at the dialing round and (ii) pairing the messages that are to be exchanged. In our construction, we applied Sharemind as a secure platform and the oblivious Quicksort algorithm for sorting. We argued about the security of our system under a robust cryptographic framework and implemented a prototype of our system in Sharemind for extensive evaluation and benchmarks. We discussed and compared our system with state-of-the-art related work. Finally, we provided a parallelized implementation of our conversation protocol as theoretical evidence for further improving performance. We strongly believe that MCMix can be seen as a proof of concept for the potential of MPC-based privacy-enhancing technologies that can be applicable in real-world scenarios at a large scale.

## Acknowledgements

Alexopoulos, Kiayias and Zacharias were supported by the Horizon 2020 PANORAMIX project (Grant Agreement No. 653497). Alexopoulos was also supported by the DFG as part of project S1 within the CRC 1119 CROSSING. Talviste was supported by the Estonian Research Council (Grant No. IUT27-1). The authors would like to thank Tim Grube and Chris Campbell for their comments on a previous version of this paper.

## References

- [ABF<sup>+</sup>17] Toshinori Araki, Assi Barak, Jun Furukawa, Tamar Lichter, Yehuda Lindell, Ariel Nof, Kazuma Ohara, Adi Watzman, and Or Weinstein. Optimized Honest-Majority MPC

- for Malicious Adversaries – Breaking the 1 Billion-Gate Per Second Barrier. In *IEEE Symposium on Security and Privacy*, pages 843–862, 2017.
- [AFL<sup>+</sup>16] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *ACM CCS*, pages 805–817, 2016.
- [AGH15] Joseph A Akinyele, Christina Garman, and Susan Hohenberger. Automating fast and secure translations from type-I to type-III pairing schemes. In *ACM CCS*, pages 1370–1381, 2015.
- [AKL<sup>+</sup>11] Diego F Aranha, Koray Karabina, Patrick Longa, Catherine H Gebotys, and Julio López. Faster explicit formulas for computing pairings over ordinary curves. In *EUROCRYPT*, pages 48–68, 2011.
- [AKS83] Miklós Ajtai, János Komlós, and Endre Szemerédi. An  $O(n \log n)$  sorting network. In *ACM STOC*, pages 1–9, 1983.
- [AS16] Sebastian Angel and Srinath Setty. Unobservable communication over fully untrusted infrastructure. In *OSDI*, pages 551–569, 2016.
- [Bat68] Kenneth E Batchner. Sorting networks and their applications. In *Proceedings of the April 30–May 2, 1968, spring joint computer conference*, pages 307–314. ACM, 1968.
- [BDNP08] Assaf Ben-David, Noam Nisan, and Benny Pinkas. Fairplaymp: a system for secure multi-party computation. In *ACM CCS*, pages 257–266, 2008.
- [Bea97] Donald Beaver. Commodity-based cryptography. In *ACM STOC*, pages 446–455, 1997.
- [BGIK16] Amos Beimel, Ariel Gabizon, Yuval Ishai, and Eyal Kushilevitz. Distribution design. In *ITCS*, pages 81–92, 2016.
- [BLR14] Dan Bogdanov, Peeter Laud, and Jaak Randmets. Domain-polymorphic programming of privacy-preserving applications. In *PLAS*, pages 53–65, 2014.
- [BLT14] Dan Bogdanov, Sven Laur, and Riivo Talviste. A Practical Analysis of Oblivious Sorting Algorithms for Secure Multi-party Computation. In *Proceedings of the 19th Nordic Conference on Secure IT Systems, NordSec 2014*, volume 8788 of *LNCS*, pages 59–74. Springer, 2014.
- [Bog13] Dan Bogdanov. *Sharemind: programmable secure computations with practical applications*. PhD thesis, University of Tartu, 2013.
- [BY03] Mihir Bellare and Bennet S. Yee. Forward-security in private-key cryptography. In *CT-RSA*, pages 1–18, 2003.
- [CCS07] Liqun Chen, Zhaohui Cheng, and Nigel P. Smart. Identity-based key agreement protocols from pairings. *Int. J. Inf. Sec.*, 6(4):213–241, 2007.
- [CGBM15] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *IEEE Symposium on Security and Privacy*, pages 321–338, 2015.
- [CGF10] Henry Corrigan-Gibbs and Bryan Ford. Dissent: accountable anonymous group messaging. In *ACM CCS*, pages 340–350, 2010.

- [Cha81] David L Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.
- [Cha88] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of cryptology*, 1(1):65–75, 1988.
- [CJK<sup>+</sup>16] David Chaum, Farid Javani, Aniket Kate, Anna Krasnova, Joeri de Ruiters, and Alan T. Sherman. cMix: Anonymization by high-performance scalable mixing. *IACR Cryptology ePrint Archive*, 2016.
- [CK03] Liqun Chen and Caroline Kudla. Identity based authenticated key agreement protocols from pairings. In *CSFW-16*, pages 219–233, 2003.
- [CKGS98] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *Journal of the ACM (JACM)*, 45(6):965–981, 1998.
- [CSWH01] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing Privacy Enhancing Technologies*, pages 46–66. Springer, 2001.
- [DDM03] George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a type III anonymous remailer protocol. In *IEEE Symposium on Security and Privacy*, pages 2–15, 2003.
- [DGKN09] Ivan Damgård, Martin Geisler, Mikkel Krøigaard, and Jesper Buus Nielsen. Asynchronous multiparty computation: Theory and implementation. In *PKC*, pages 160–179, 2009.
- [DM06] Roger Dingledine and Nick Mathewson. Anonymity loves company: Usability and the network effect. In *WEIS*, 2006.
- [DMS04] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. Technical report, DTIC Document, 2004.
- [DPSZ12] Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, pages 643–662, 2012.
- [Dwo06] Cynthia Dwork. Differential privacy. In *Automata, languages and programming*, pages 1–12. Springer, 2006.
- [FG10] Dario Fiore and Rosario Gennaro. Identity-based key exchange protocols without pairings. *Trans. Computational Science*, 10:42–77, 2010.
- [FLNW17] Jun Furukawa, Yehuda Lindell, Ariel Nof, and Or Weinstein. High-throughput secure three-party computation for malicious adversaries and an honest majority. In *EUROCRYPT*, pages 225–255, 2017.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game. In *ACM STOC*, pages 218–229, 1987.
- [Gün89] Christoph G. Günther. An identity-based key-exchange protocol. In *EUROCRYPT*, pages 29–37, 1989.

- [HKI<sup>+</sup>12] Koki Hamada, Ryo Kikuchi, Dai Ikarashi, Koji Chida, and Katsumi Takahashi. Practically efficient multi-party sorting protocols from comparison sort algorithms. In *Information Security and Cryptology–ICISC 2012*, pages 202–216. Springer, 2012.
- [JWJ<sup>+</sup>13] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. Users get routed: Traffic correlation on tor by realistic adversaries. In *ACM CCS*, pages 337–348, 2013.
- [KCB17] Panayiotis Kotzanikolaou, George Chatzisofofroniou, and Mike Burmester. Broadcast anonymous routing (BAR): scalable real-time anonymous communication. *Int. J. Inf. Sec.*, 16(3):313–326, 2017.
- [KLDF15] Albert Kwon, David Lazar, Srinivas Devadas, and Bryan Ford. Riffle: An efficient communication system with strong anonymity. *PoPETS*, 2016(2):115–134, 2015.
- [Lan15] A Langley. Pond (v0.1.1). <https://github.com/agl/pond>, 2015.
- [LP16] Peeter Laud and Martin Pettai. Secure multiparty sorting protocols with covert privacy. In *NordSec*, pages 216–231, 2016.
- [LWN<sup>+</sup>15] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. Oblivm: A programming framework for secure computation. In *IEEE Symposium on Security and Privacy*, pages 359–376, 2015.
- [LZ16] David Lazar and Nikolai Zeldovich. Alpenhorn: Bootstrapping secure communication without leaking metadata. In *OSDI*, pages 571–586, 2016.
- [MSZ15] Mahnush Movahedi, Jared Saia, and Mahdi Zamani. Shuffle to baffle: Towards scalable protocols for secure multi-party shuffling. In *ICDCS*, pages 800–801, 2015.
- [PL15] Martin Pettai and Peeter Laud. Automatic proofs of privacy of secure multi-party computation protocols against active adversaries. In *CSF*, pages 75–89, 2015.
- [Pos81a] Jon Postel. Internet protocol. RFC, 1981.
- [Pos81b] Jon Postel. Transmission control protocol. RFC, 1981.
- [PS09] Kenneth G. Paterson and Sriramkrishnan Srinivasan. On the relations between non-interactive key distribution, identity-based encryption and trapdoor discrete log groups. *Des. Codes Cryptography*, 52(2):219–241, 2009.
- [Sag16] Sage Developers. *SageMath, the Sage Mathematics Software System (Version 7.2)*, 2016. <http://www.sagemath.org>.
- [SGR97] Paul F Syverson, David M Goldschlag, and Michael G Reed. Anonymous connections and onion routing. In *IEEE Symposium on Security and Privacy*, pages 44–54, 1997.
- [Sha84] Adi Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53, 1984.
- [She59] Donald L. Shell. A high-speed sorting procedure. *Communications of the ACM*, 2(7):30–32, 1959.
- [SKO00] Ryuichi Sakai, Masao Kasahara, and K Oghishi. Cryptosystems based on pairing. SCIS, Okinawa, Japan, 2000.

- [Sma01] Nigel P. Smart. An identity based authenticated key agreement protocol based on the weil pairing. *IACR Cryptology ePrint Archive*, 2001.
- [VDHLZZ15] Jelle Van Den Hooff, David Lazar, Matei Zaharia, and Nickolai Zeldovich. Vuvuzela: Scalable private messaging resistant to traffic analysis. In *SOSP*, pages 137–152, 2015.
- [Wan13] Yongge Wang. Efficient identity-based and authenticated key agreement protocol. *Trans. Computational Science*, 17:172–197, 2013.
- [WCGFJ12] David Isaac Wolinsky, Henry Corrigan-Gibbs, Bryan Ford, and Aaron Johnson. Dissent in numbers: Making strong anonymity scale. In *OSDI*, pages 179–182, 2012.
- [YL05] Quan Yuan and Songping Li. A new efficient id-based authenticated key agreement protocol. *IACR Cryptology ePrint Archive*, 2005.
- [ZSB13] Yihua Zhang, Aaron Steele, and Marina Blanton. Picco: a general-purpose compiler for private distributed computation. In *ACM CCS*, pages 813–826, 2013.