

Privacy-Preserving Epidemiological Modeling on Mobile Graphs

Daniel Günther
guenther@crypto.cs.tu-
darmstadt.de

Technical University of Darmstadt

Helen Möllering
moellering@crypto.cs.tu-
darmstadt.de

Technical University of Darmstadt

Benjamin Judkewitz
benjamin.judkewitz@charite.de
Charité-Universitätsmedizin

Benny Pinkas
benny@pinkas.net
Bar-Ilan University

Marco Holz
holz@crypto.cs.tu-darmstadt.de
Technical University of Darmstadt

Thomas Schneider
schneider@crypto.cs.tu-
darmstadt.de
Technical University of Darmstadt

Ajith Suresh
suresh@crypto.cs.tu-darmstadt.de
Technical University of Darmstadt

ABSTRACT

Over the last two years, governments all over the world have used a variety of containment measures to control the spread of COVID-19, such as contact tracing, social distance regulations, and curfews. Epidemiological simulations are commonly used to assess the impact of those policies before they are implemented in actuality. Unfortunately, their predictive accuracy is hampered by the scarcity of relevant empirical data, concretely detailed social contact graphs. As this data is inherently privacy-critical, there is an urgent need for a method to perform powerful epidemiological simulations on real-world contact graphs without disclosing sensitive information.

In this work, we present RIPPLE, a privacy-preserving epidemiological modeling framework that enables the execution of a wide range of standard epidemiological models for any infectious disease on a population's most recent real contact graph while keeping all contact information private locally on the participants' devices. In this regard, we also present PIR-SUM, a novel extension to private information retrieval that allows users to securely download the sum of a set of elements from a database rather than individual elements. Our theoretical constructs are supported by a proof-of-concept implementation in which we show that a 2-week simulation over a population of half a million can be finished in 7 minutes with each participant consuming less than 50 KB of data.

KEYWORDS

Decentralized Epidemiological Modeling, Privacy, Private Information Retrieval, COVID-19

1 INTRODUCTION

The current pandemic had and has a significant impact on people's daily lives, posing significant challenges such as increased mental illness, parents having to balance child care, homeschooling, and work, an increase in domestic abuse cases, and many more [88, 117, 126]. Governments all over the world have taken a variety of steps in the last two years to restrict the spread of the virus in order to save human lives and keep the economic system working. Those range from closing institutions, such as schools, to country-wide lockdowns. Despite these courageous efforts, the global number of infections skyrocketed, and COVID-19 claimed far too many

lives. Aside from highly lethal diseases like COVID-19, many other infectious diseases have emerged and have had a significant impact on human life over time. For example, incidences of monkey pox in Europe have lately increased to the point that quarantine measures have been implemented [22, 33, 52].

In the context of COVID-19, contact tracing apps are being used all over the world to notify contacts of potential infections [3, 63, 86, 104, 109, 120, 121, 124]. Unfortunately, there is a fundamental limitation to contact tracing: It only notifies contacts of an infected person *after* the infection has been detected, i.e., typically after a person develops symptoms, is tested, receives the test result, and is able to connect with contacts [82, 122]. Tupper et al. [122] report that in British Columbia in April 2021, this process ideally took 5 days, reducing new cases by only 8% compared to not using contact tracing. They conclude that contact tracing must be supplemented with multiple additional containment measures in order to effectively control disease spread.

Epidemiological modelling allows to predict the spread of an infectious disease in the *future* and has received a lot of attention [35, 55, 76, 112, 113, 118, 131]. It allows to assess the effectiveness of containment measures by mathematically modelling their impact on the spread. As a result, it can be an extremely valuable tool for governments to select effective containment measures [118]. For example, Davis et al. [39] predicted in early 2020 that COVID-19 would infect 85% of the British population without any containment measures in place, causing a massive overload of the health system (13-80× the capacity of intensive care units). Their forecast also indicated that short-term interventions such as school closures, social distancing, and so on will not be effective in reducing the number of cases. As a result, the British government decided to implement a lockdown in March 2020, effectively reducing transmissions and stabilising the health system [118].

With access to detailed information about a population's size, density, the transportation, and health care system, epidemiological modelling could accurately forecast disease transmission in a variety of situations [2]. Especially precise, up-to-date information about movements and physical interactions in space and time are crucial for precisely forecasting transmission as well as the impact of various control measures before being implemented [73]. In practice, these simulations may quickly model future disease's spread,

calculate the projected number of infections when specific actions are taken, and divert disease spread in specific areas. However, data on personal encounters is very scarce, thus the impact of containment measures can only be approximated so far [2, 49, 73]. This lack of data is primarily owing to the fact that encounter data has generally been acquired by surveys, which do not accurately reflect reality [45, 73], e.g., random encounters in public transport or shopping malls. Moreover, social interaction patterns change over time and sometimes even rapidly as we have seen with social distancing measures, rendering collected contact information outdated. None of the existing data permits realistic simulations on the actual person-to-person social contact graph. From a modeler’s perspective, epidemiologists would ideally like to have access to the complete physical interaction graph of a population. Furthermore, the right for privacy reflected in strict data protection regulations of liberal states make accurate tracking of interpersonal contacts unacceptable.

To address the issue of obtaining the most recent contact data while protecting individuals’ privacy, we present RIPPLE, a practical privacy-preserving framework for epidemiological modelling that allows precise simulations of disease spread based on current participant data while taking into account deployed control measures and without leaking any information about individuals contacts. RIPPLE provides a privacy-preserving method for collecting real-time physical encounters and can compute arbitrary compartment-based epidemiological models¹ on the most recent contact graph in a privacy-preserving manner. RIPPLE can be used to investigate the effect of containment measures not only for COVID-19, but for any infectious diseases. We anticipate that our framework’s privacy guarantee will encourage more people to participate, allowing epidemiologists to compute more accurate simulations that will eventually help to develop effective containment measures against diseases in the future.

1.1 Our Contributions

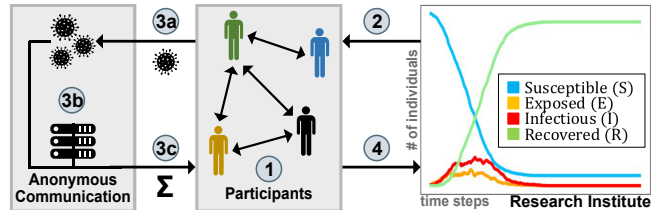
This paper introduces RIPPLE, a framework for expanding the scope of privacy research from contact tracing to epidemiological modelling. While the former only warns about potential infections in the past, epidemiological modelling can predict the spread of infectious diseases in the future. Anticipating the effects of various control measures allows for the development of informed epidemic containment strategies and political interventions prior to their implementation.

RIPPLE uses a fully decentralised system similar to the federated learning paradigm [90] to achieve high acceptance and trust in the system and to motivate many participants to join the system for generating representative contact information. All participant data, such as encounter location, time, distance, and so on, are kept locally on the participants’ devices. Participants in RIPPLE communicate with one another through anonymous communication channels enabled by a group of *semi-honest* central servers.

RIPPLE comprises of two methods for achieving privacy-preserving epidemiological modelling, each of which caters to a different use case. The first is RIPPLE_{TEE}, which assumes that each participant’s mobile device has a Trusted Execution Environment (TEE).

¹ The implementation of concrete simulation functions is outside the scope of this work and referred to medical experts. More details on epidemiological modelling are given in §2.

The second method, RIPPLE_{PIR}, eliminates this assumption by utilising cryptographic primitives such as Private Information Retrieval (PIR). Along the way, we develop a multi-server PIR extension that allows a client to directly retrieve the sum of a set of elements (in our case infection likelihoods) from a database without learning anything about the individual entries.



- ① Mobile apps collect anonymous encounter tokens during interactions.
- ② Research Institute begins the simulation by providing initialization parameters.
- ③a Participants upload infection likelihood securely to servers.
- ③b Servers securely compute cumulative infection likelihood per participant.
- ③c Participants retrieve their cumulative infection likelihood.
- ④ The aggregate results (#S, #E, #I, #R) are sent to the Research Institute.

Figure 1: Overview of RIPPLE Framework.

We assess the practicality of our methods by benchmarking core building blocks using a proof of concept implementation (cf. §6). Our findings indicate that, with adequate hardware, both protocols can scale up to millions of participants. For instance, a simulation of 14 days with 1 million participants can be completed in less than half an hour. We summarize our contributions as follows:

- (1) We present RIPPLE, the *first* privacy-preserving framework to perform epidemiological modelling on contact information stored on mobile devices.
- (2) RIPPLE formalises the notion of *privacy-preserving* epidemiological modelling and defines privacy requirements in the presence of both semi-honest and malicious participants.
- (3) For epidemiological simulations using real-world contact data acquired with participants’ mobile devices, we present two techniques – RIPPLE_{TEE} and RIPPLE_{PIR}, that combine anonymous communication techniques with either TEEs or PIR and anonymous credentials.
- (4) We propose PIR-SUM, an extension to existing PIR schemes, that allows a client to download the sum of τ distinct database entries without learning the values of individual entries or revealing which entries were requested.
- (5) We demonstrate the practicality of our framework by providing a detailed performance evaluation of RIPPLE using open source code as a benchmark.

2 RELATED WORK & BACKGROUND INFORMATION

This section provides necessary background information on contact tracing and epidemiological modelling while an overview of the (cryptographic) primitives and other techniques used in this work is presented in §A.

Contact Tracing. A plethora of contact tracing systems were introduced and deployed in the last two years [3, 34, 109]. They

either use people’s location (GPS or telecommunication provider information) or measure proximity (via Bluetooth LE). Most systems can be categorized into centralized and decentralized designs [124]. In a centralized contact tracing system (e.g., [65, 116]), computations such as the generation of the tokens exchanged during physical encounters are done by a central party. Similarly, this central party may also store some contact information depending on the concrete system design. In contrast, in decentralized approaches (e.g., [23, 104, 121]), computation and encounter information remain (almost completely) locally at the participants’ devices.

Presence Tracking and Hotspot Detection. The recent work of CrowdNotifier [86] notifies visitors of (large) events about an infection risk when another visitor reported SARS-CoV-2 positive after the event even if they have not been in close proximity of less than 2 meters. To protect user privacy, it follows a distributed approach where location and time information are stored encrypted on the user’s device. Bampoulidis et al. [12] introduce a privacy-preserving two-party set intersection protocol that detects infection hotspots by intersecting infected patients, input by a health institute, with customer data from mobile network operators.

Quarantine and Social Distancing. Al-Turjman and David Deebak [4] integrate privacy-protecting health monitoring into a Medical Things device that monitors the health status (heart rate, oxygen saturation, temperature, etc.) of users in quarantine with moderate symptoms. Only in the event of an emergency medical personnel is notified. Pezzutto et al. [103] optimize the distribution of a limited set of tests to identify as many positive cases as possible, which are then quarantined. Their system can be deployed in a decentralized, privacy-preserving environment to identify individuals who are at high risk of infection. Barocchi et al. [13] develop a privacy-preserving architecture for indoor social distancing based on a privacy-preserving access control system. When users visit public facilities (for example, a supermarket or an airport), their mobile devices display a route recommendation for the building that maximizes the distance to other people. Bozedemir et al. [18] suggest privacy-preserving trajectory clustering to identify typical movements of people as well as to detect forbidden gatherings when contact restrictions are in place.

Contact tracing, presence tracing, and hotspot detection are concerned with “flattening the curve” in relation to infections in the past. In contrast, epidemiological modelling is a tool for decision makers to evaluate the efficacy of containment measures like social distancing in the future, allowing them to “get ahead of the wave”.

Epidemiological Modeling. There are several options to mathematically model a disease. The popular compartment models [19, 20, 51, 57, 60, 62, 110, 131] capture the spread with a few continuous variables linked by simple differential equations. A prominent example is the SEIR model [42, 62, 71, 131] with four compartments to which people are assigned, namely, susceptible (S), exposed (E), infectious (I), and recovered (R). For each simulated time interval, the number of people assigned to each class are computed. While such models are useful for capturing macroscopic trends and also used in state-of-the-art epidemiological research, e.g., [31, 115], the basic approaches condense complex individual behaviour into few variables, thus, limiting the simulation’s predictive power [89, 100].

Agent-based epidemiological models [50], on the other hand, initialise a large number of agents with a set of individual properties (e.g., location, age, etc.). These agents then interact according to a set of interaction rules (e.g., location-based, age-based, etc.) to simulate disease spread. The simulations are carried out in many time steps. Combining both directions, i.e., using agents in a compartment model, allows for a more realistic model of individual behaviour for forecasting disease transmission in a population. Many such simulations with varying parameters are run in parallel to simulate the effect of various policy interventions (e.g., reducing interactions between agents of a certain age, capping the maximum number of allowed contacts, or vaccinating a selected group of agents). The aggregated number of agents assigned to the same “infection class” (e.g., susceptible, exposed, infected, and recovered for the SEIR model) is then computed for each simulation step.

A crucial question is how to model the agents’ individual contact behaviour. Older models relied on survey-based contact matrices, which included information such as the average number of contacts in a given age range [73]. This is already a significant improvement over treating all people the same. However, aggregated network statistics cannot recreate the dynamics of a real complex network graph, as evidenced by the prevalence of super-spreaders with far more contacts than the average [77]. Thus, using the real-world contact graph between all individual members of the population would be ideal from an epidemiological standpoint.

Contact Tracing for Privacy-Preserving Epidemiological Modeling. If contact information collected through contact tracing apps was centralised, an up-to-date full contact graph could be constructed for epidemiological simulations. Contact information, on the other hand, is highly sensitive information that should not be shared. Contact information collected via mobile phones can reveal who, when, and whom people meet which is by itself sensitive information and must be protected. Beyond, those information also enable to derive indications about the financial situation [15, 87, 114], personality [94], life-partners [6], and ethnicity [6]. One can think about many more examples: By knowing which medical experts are visited by a person information about the health condition can be anticipated; contact to members of a religious minority as well as visits to places related to a religion might reveal a religious orientation, etc. Thus, it would be ideal to enable precise epidemiological simulations without leaking individual contact information.

One way to achieve privacy-preserving epidemiological modelling from contact tracing apps is to let each participant (i.e., each device using the contact tracing app) secretly share its contact information between a set of non-colluding servers, which can then jointly run simulations using techniques like secure multi-party computation (cf. §A). In fact, a recent paper showed how to efficiently run graph algorithms on secret shared graphs via MPC [9]. Even though such a non-collusion assumption is common in the crypto community, the general public in some countries may have difficulty trusting a system in which all contact information is disclosed if the servers collude. As a result, RIPPLE aims to distribute trust by involving all participants in such a way that they can keep their own contact information local while simulating the spread of a disease by sending messages to each other anonymously. Furthermore, only aggregated simulation results will be shared with a

research institute, so no data directly relating to a single identity will be shared. This approach mimics the baseline idea of Federated Learning [90] and prominent contact tracing designs supported by Google and Apple.² The increased trust level of a distributed design fosters the crucial broad adoption of such a system in the population. To the best of our knowledge, RIPPLE is the first framework that allows the execution of any agent-based compartment model on the real contact graph while maintaining privacy.

3 THE RIPPLE FRAMEWORK

RIPPLE’s primary goal is to enable the evaluation of the impact of multiple combinations of potential containment measures defined by epidemiologists and the government, and to find a balance between the drawbacks and benefits of those measures, rather than to deploy the measures in “real-life” first and then analyse the impact afterwards. Such measures may include, for example, the requirement to wear face masks in public places, restrictions on the number of people allowed to congregate, the closure of specific institutions and stores, or even complete curfews and lockdowns within specific regions.

Participants in RIPPLE collect personal encounter data anonymously and locally store it on their mobile devices such as cell phones, similar to privacy-preserving contract tracing apps. However, for epidemiological modelling, RIPPLE must also derive a contact graph without leaking sensitive personal information in order to compute simulations of disease spread, which may involve multiple sets of containment measures for some time period, such as two weeks. In almost every country, we can find a 6-hour period during the night when the majority of the population sleeps and mobile devices are idle, connected to the Internet via WiFi, and possibly charging, i.e., an ideal time window for running RIPPLE simulations. The results can then be analysed by medical experts to learn more about the disease or by political decision makers to determine the most promising containment measures to implement.

To acquire representative and up-to-date physical encounter data, we require widespread public acceptance of RIPPLE, similar to contact tracing apps. One way to encourage this is to piggyback RIPPLE on the official contract tracing applications used by most countries. Politicians, on the other hand, can motivate residents beyond the intrinsic incentive of supporting public health by coupling the use of RIPPLE with additional benefits such as discounted or free travel passes.

System and Threat Model. RIPPLE comprises of p participants, denoted collectively by \mathcal{P} , a research institute RI who is in charge of the epidemiological simulations, and a set of MPC servers C responsible for anonymous communication among the participants.

We assume that the research institute and MPC servers follow the semi-honest security model [56], which means they correctly follow protocol specifications while attempting to gather additional information. Though not the strongest security model, it provides a good trade-off between privacy and efficiency, which is why it is commonly used in the design of several practical privacy-preserving applications such as privacy-preserving machine learning [25, 91, 93, 101], genome/medical research [111, 119, 125], and

localization services [68, 123]. It also protects against passive attacks by curious administrators and accidental data leakage. Furthermore, it is quite often the first step toward developing protocols with stronger privacy guarantees [11, 84]. We believe this is a reasonable assumption in our setting because the research institute and the servers will be controlled/run by generally trusted entities such as governments or (public) medical research centres, potentially in collaboration with NGOs such as the EFF³ or the CCC⁴.

Since it is in general interest to find effective containment measures, we anticipate that many participants will be highly intrinsically motivated to contribute to a successful epidemiological modelling. However, assuming that each and every one of potentially millions of participants is honest is unrealistic. As a result, we also consider a client-malicious security model [24, 81] for the participants in \mathcal{P} , in which some of the participants may deviate from the protocol to gather additional information about their encounters. We should also mention that malicious behaviour can actively try to hamper or even destroy the simulation itself. However, in the scope of this work, we concentrate on the aforementioned deviations for additional information gain, leaving the problem of developing efficient countermeasures against correctness attacks to future work. Tab. 1 summarises the notations used in this work.

	Parameter	Explanation
Entities	\mathcal{P}	Set of all participants; $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_p\}$.
	RI	Research Institute
	C	Communication Servers; $\{S_0, S_1, S_2\}$ in this work
Simulations	param _{sim}	simulation parameters defined by RI
	N_{sim}	# distinct simulations (executed in parallel)
	N_{step}	# steps per simulation
	class _{inf}	infection classes; class _{inf} = $\{\text{class}_{inf}^1, \dots, \text{class}_{inf}^{N_{inf}}\}$
	I_i^s	\mathcal{P}_i ’s infection class in simulation step $s \in [0, N_{step}]$
	\mathcal{E}_i	Encounter tokens of \mathcal{P}_i
Protocols	E_i^{\max}	#max. encounters by \mathcal{P}_i in pre-defined time interval
	κ	computational security parameter; $\kappa = 128$ in this work
	r_e	Unique token for encounter $e \in [0, E^{\max}]$
	$\delta_i^{r_e}$	\mathcal{P}_i ’s infection likelihood w.r.t token r_e
	Δ_i	\mathcal{P}_i ’s cumulative infection likelihood
	m_i^e	metadata of an encounter e by \mathcal{P}_i
	(pk_i, sk_i)	\mathcal{P}_i ’s public/private key pair
	σ_i^e	\mathcal{P}_i ’s signature on message about encounter e

Table 1: Notations used in RIPPLE.

3.1 Phases of RIPPLE

RIPPLE is divided into four phases as shown in Fig. 1: i) Token Generation, ii) Simulation Initialization, iii) Simulation Execution, and iv) Result Aggregation. While our framework can be applied to any compartment-based epidemiological modeling of any infectious disease (cf. §2), we explain RIPPLE using the prevalent Covid-19 and the SEIR model [42, 71] as a running example. For the sake of simplicity, we assume that an app that emulates RIPPLE is installed on each participant’s mobile device, and that the participants locally enter attributes such as workplace, school, regular eateries, and cafes in the app after installing the app.

² <https://covid19.apple.com/contacttracing>

³ <https://www.eff.org> ⁴ <https://www.ccc.de/en/>

① - **Token Generation:** During a physical encounter, participants exchange data via Bluetooth LE to collect anonymous encounter information (Fig. 2), similar to contact tracing [61, 104, 121]. These tokens are stored locally on the devices of the users and do not reveal any sensitive information (i.e., identifying information) about the individuals involved.

In addition to these tokens, the underlying application will collect additional information on the context of the encounter known as "metadata" for simulation purposes. This can include details such as duration, proximity, time, and location, etc. The metadata can be used to include or

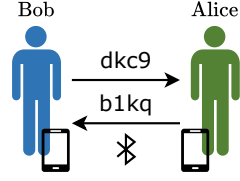


Figure 2: Token Generation.

exclude different encounters in the simulation phase, allowing the effect of containment measures to be modelled (e.g., restaurant closings, by excluding all encounters that happened in restaurants). The token generation phase is not dependent on the simulation phase, so no simulation-dependent infection data is exchanged. The token generation phase is modelled as an ideal functionality \mathcal{F}_{gen} that will be instantiated later in §4.

– As an example, assume that a participant, Alice, takes the bus to pick up her daughter from school. There are several other people on this bus – for simplicity, we call them $\text{Bob}_1, \dots, \text{Bob}_x$. As part of the token generation phase, Alice’s phone exchanges unique anonymous tokens with the devices of the different Bobs. Now, two weeks later, it is night, and the national research institute (RI) wants to run a simulation covering 14 days to see what effect closing all schools would have on the disease’s spread. To accomplish this, the RI notifies all registered participants’ applications to run a simulation using encounter data from the previous two weeks.

② - **Simulation Initialization:** The research institute RI initiates the simulation phase by sending a set of parameters, denoted by $\text{param}_{\text{sim}}$, to the participants in \mathcal{P} . The goal is to "spread" a fictitious infection across N_{sim} different simulation settings. To begin a simulation, each participant \mathcal{P}_i is assigned to an infection class $I_i^{\text{init}} \in \text{class}_{\text{inf}}$ (e.g., {S}usceptible, {E}xposed, {I}nfectious, {R}ecovered for the SEIR model) as specified in $\text{param}_{\text{sim}}$. For each individual simulation, $\text{param}_{\text{sim}}$ defines a set of containment measures, such as school closings and work from home, which the participants will use as filters to carry out the simulation in the next stage. In addition, RI publishes a formula to calculate the infection likelihood, which is denoted by δ . The likelihood is determined by several parameters in the underlying modelling, such as encounter distance and time. For example, this likelihood might range from 0 (no chance of infection) to 100 (certain to get infected).

– In our example, assume Alice is designated as infectious, while Bob_1 is designated as susceptible by RI. The other participants $\text{Bob}_2, \dots, \text{Bob}_x$ are also assigned to an infection class (S, E, I, or R). To simulate containment measures, the RIPPLE-app now employs filters defined in $\text{param}_{\text{sim}}$. Using the information provided by the participants⁵, the application may automatically filter out encounters that would not happen if a containment measure were in place, such as encounters in school while simulating school closings.

⁵ This may also include location data obtained from the mobile app., e.g., Check In and Journal fields in the Corona-Warn contact tracing app.

③ - **Simulation Execution:** Once the RI initialises the simulation, N_{step} simulation steps (③a, ③b, ③c in Fig. 1) are performed for each of the N_{sim} simulation settings (e.g., $N_{\text{step}} = 14$ days). Without loss of generality, consider the first simulation step and let $N_{\text{sim}} = 1$. The simulation proceeds as follows:

1) Participant $\mathcal{P}_i \in \mathcal{P}$ filters out the relevant encounters based on the containment measures defined by RI and let the corresponding encounter tokens be represented by the set \mathcal{E}_i .

2) For each token $r_e \in \mathcal{E}_i$, \mathcal{P}_i computes the infection likelihood $\delta_i^{r_e}$ using the formula from RI, i.e., the probability that \mathcal{P}_i infects the respective participant they met during the encounter with identifier token r_e .

3) Participants use the likelihood values δ obtained in the previous step to execute an ideal functionality called $\mathcal{F}_{\text{esim}}$, which allows them to communicate the δ values anonymously through a set of C communication servers. Furthermore, it allows each participant \mathcal{P}_j to receive a cumulative infection likelihood, denoted by Δ_j , based on all of the encounters they had on the day being simulated, i.e., $\Delta_j = \sum_{r_e \in \mathcal{E}_j} \delta_j^{r_e}$. In this case, $\delta_j^{r_e}$ denotes the infection likelihood

computed by participant \mathcal{P}_f and communicated to \mathcal{P}_j for an encounter between \mathcal{P}_f and \mathcal{P}_j with identifier token r_e . As will be discussed later in §3.2, $\mathcal{F}_{\text{esim}}$ must output the cumulative result rather than individual infection likelihoods because the latter can result in a breach of privacy.

4) Following the guidelines set by the RI, \mathcal{P}_j updates its infection class I_j using the cumulative infection likelihood Δ_j acquired in the previous step.

These steps above are repeated for each of the N_{step} simulation steps in order and across all the N_{sim} simulation settings.

– In our example, let the simulated containment measure be closure of schools. As Alice is simulated to be infectious, Alice’s phone computes the infection likelihood for every single encounter it recorded on the day exactly two weeks ago (Day 1) *except* those that occurred at her daughter’s school. Following that, Alice’s phone combines the computed likelihood of each encounter with the corresponding unique encounter token to form tuples, which are then sent to an entry point for an anonymous communication channel. Using the encounter token as an address, this channel

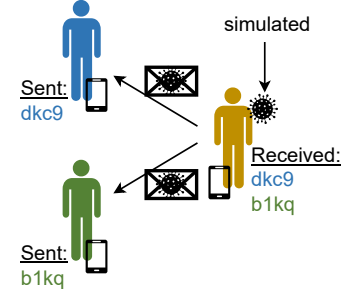


Figure 3: Simulation.

and checks the aggregate result to a threshold defined by the RI to see if he has been infected in the simulation⁶.

⁶ Bob_1 obtains the aggregated likelihood in the actual protocol.

④ - **Result Aggregation:** For a given simulation setting, each participant $\mathcal{P}_i \in \mathcal{P}$ will have its infection class I_i^s updated at the end of every simulation step $s \in [N_{\text{step}}]$. The goal of this phase to allow RI to obtain the aggregated number of participants per class (e.g., #S, #E, #I, #R) for each simulated time step. For this, we rely on a *Secure Aggregation* functionality, denoted by \mathcal{F}_{agg} , which takes a N_{inf} -tuple of the form $\{v_i^1, \dots, v_i^{N_{\text{inf}}}\}^s$ from each participant for every simulation step s and outputs the aggregate of this tuple over all the p participants to RI. In this case, v_i^k is an indicator variable for the k -th infection class, which is set to one if $I_i^s = \text{class}_{\text{inf}}^k$ and zero otherwise. Secure aggregation [47, 78, 83, 83] is a common problem in cryptography these days, particularly in the context of federated learning, and there are numerous solutions proposed for various settings, such as using TEEs, a semi-trusted server aggregating ciphertexts under homomorphic encryption, or multiple non-colluding servers that aggregate secret shares. In this work, we consider \mathcal{F}_{agg} to be a black-box that can be instantiated using any of the existing solutions that are compatible with our framework.

– In our example, all participants will know their updated infection class at the end of Day 1’s simulation round, and they will prepare a 4-tuple of the form $\{v^S, v^E, v^I, v^R\}$ representing their updated infection class in the SEIR model. Participants will then engage in a secure aggregation protocol that determines the number of participants assigned to each infection class, which is then delivered to the RI. Then, the second simulation round begins, which replicates the procedure but this time using encounters from 13 days ago, i.e., Day 2. The RI holds the aggregated number of participants per day per class after simulating all 14 days, i.e., a simulation of how the disease would spread if all schools had been closed in the previous 14 days (cf. graph in Fig. 1).

Complete Protocol: Fig. 4 summarises the phases of the RIPPLE framework in the context of a single simulation setting. Note that multiple simulations can be executed in parallel.

Protocol RIPPLE

① - Token Generation

- $\mathcal{P}_i \in \mathcal{P}$ executes \mathcal{F}_{gen} all the time (on its mobile device), collecting encounter data of the form (r_e, m_e) with $e < E_i^{\text{max}}$.

② - Simulation Initialization

- $\mathcal{P}_i \in \mathcal{P}$ receives $\text{param}_{\text{sim}}$ from RI and locally sets $I_i^1 = I_i^{\text{init}}$.

③ - Simulation Execution

For each simulation step $s \in [N_{\text{step}}]$, $\mathcal{P}_i \in \mathcal{P}$ execute the following:

- Filter out encounters using $\text{param}_{\text{sim}}$ to obtain encounter set \mathcal{E}_i^s .
- For each token $r_e \in \mathcal{E}_i^s$, compute the infection likelihood $\delta_i^{r_e}$ locally using the formula from RI.
- Invoke $\mathcal{F}_{\text{esim}}$ with the input $\{\delta_i^{r_e}\}_{r_e \in \mathcal{E}_i^s}$ and obtain $\Delta_i^s = \sum_{r_e \in \mathcal{E}_i^s} \delta_i^{r_e}$.
- Update the infection class I_i^s using Δ_i^s and the guidelines from RI.

④ - Result Aggregation

For each simulation step $s \in [N_{\text{step}}]$, execute the following:

- $\mathcal{P}_i \in \mathcal{P}$ prepares $\{v_i^1, \dots, v_i^{N_{\text{inf}}}\}^s$ with $v_i^k = 1$ if $I_i^s = \text{class}_{\text{inf}}^k$ and $v_i^k = 0$ otherwise, for $k \in [N_{\text{inf}}]$.

- Invoke \mathcal{F}_{agg} with inputs $\{v_i^1, \dots, v_i^{N_{\text{inf}}}\}^s$ to enable RI obtain the tuple $\{C_{\text{inf}}^1, \dots, C_{\text{inf}}^{N_{\text{inf}}}\}^s$ where $C_{\text{inf}}^k = \sum_{\mathcal{P}_i \in \mathcal{P}} v_i^k$ for $k \in [N_{\text{inf}}]$.

Figure 4: RIPPLE Framework (for one simulation setting).

3.2 Privacy Requirements

A private contact graph necessitates that participants remain unaware of any unconscious interactions. This means they cannot find out if they had unconscious contact with the same person more than once or how often they did. We remark that an insecure variant of RIPPLE, in which each participant \mathcal{P}_i receives the infection likelihood δ_i^e for all of its encounters $e \in E_i$ separately, will be unable to meet this condition.

To demonstrate this, observe that when running multiple simulations (with different simulation parameters $\text{param}_{\text{sim}}$) on the same day, participants will use the same encounter tokens and metadata from the token generation phase in each of the simulations. If a participant \mathcal{P}_i can see the

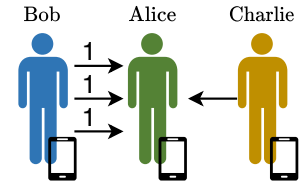


Figure 5: Linking Identities Attack. Alice and Bob had several encounters, but Alice and Charlie only had one.

infection likelihood δ_i of each of their interactions, \mathcal{P}_i can look for correlations between those likelihoods to see if another participant \mathcal{P}_j was encountered more than once. We call this a *Linking Identities Attack* and depict it in Fig. 5, where, for simplicity, the infection likelihood accepts just two values: 1 for high and 0 for low infection likelihood.

Consider the following scenario to help clarify the issue: Alice and Bob work together in the same office. As a result, they have numerous conscious encounters during working hours. However, in their spare time, they may be unaware that they are in the same location (e.g., a club) and may not want the other to know. Even if they do not see each other, their phones are constantly collecting encounters. Assume the RI sent the participants a very simple infection likelihood formula that simply returns 0 (not infected) or 1 (infected). Furthermore, since the data is symmetric, both Alice and Bob have the same metadata (duration, distance, etc.) about their conscious and unconscious encounters.

Let Bob be modelled as infectious in the first simulation. As a result, he will send a 1 for each (conscious and unconscious) encounter he had (including those with Alice). If multiple simulations are run on the same day (i.e., with the same encounters), Alice will notice that some encounters, specifically all conscious and unconscious encounters with Bob, always have the same infection likelihood: If Bob is not infectious, all will return a 0; if Bob is infectious, all will return a 1. Thus, even if Alice had unconscious encounters with Bob, she can detect the correlations between the encounters and, as a result, determine which unconscious encounters were most likely with Bob. The more simulations she runs, the more confident she becomes. Since every participant knows the formula, this attack can be extended to complex infection likelihood functions as well. While it may be more computationally expensive than the simple case, Alice is still able to identify correlations. This attack

also works even if all of the encounters were unconscious. In such situations, Alice may not be able to trace related encounters to a single person (Bob), but she can infer that they were all with the same person (which is more than learning nothing).

To avoid a Linking Identities attack, RIPPLE ensures that in a simulation phase, each participant receives just an aggregation of all infection likelihoods of their encounters. It cannot be avoided that participants can understand that when “getting infected” someone of their contacts must have been in contact with a (simulated) infectious participant. As this is only a simulated infection, we consider this leakage acceptable.

While the Linking Identities Attack is already significant in a semi-honest security model, malicious participants may go even further to circumvent aggregation mechanisms that prevent access to individual infection likelihoods. They could, for example, construct many *sybils*, i.e., multiple identities using several mobile devices, to collect each encounter one by one and then conduct a Linking Identities Attack with the information.

A registration system can be used to increase the costs of performing sybil attacks, i.e., to prevent an adversary from creating many identities. This assures that only legitimate users are allowed to join and participate in the simulation. This might be done in a closed ecosystem, such as a firm, where each member receives exactly one token to participate in the simulation. This might also be done on a larger scale at the national level, with each citizen receiving a token linked to a digital ID card. In such authentication mechanisms, anonymous credentials (cf. §A) can be used to ensure anonymity, and we leave the problem for future work.

4 INSTANTIATING $\mathcal{F}_{\text{esim}}$

In this section, we propose two instantiations of $\mathcal{F}_{\text{esim}}$ that cater to different use cases and offer different efficiency-trust trade-offs. Our first design, RIPPLE_{TEE} (§4.1), assumes the presence of trusted execution environments (TEEs) such as ARM TrustZone on the mobile devices of the participants. In our second design, RIPPLE_{PIR} (§4.2), we eliminate this assumption and provide privacy guarantees using cryptographic techniques such as private information retrieval and secure multi-party computation.

4.1 RIPPLE_{TEE}

The deployment of the entire operation in a single designated TEE would be a simple solution to achieving the ideal functionality $\mathcal{F}_{\text{esim}}$. However, given the massive amount of data that must be handled in a large scale simulation with potentially millions of users, TEE resource limitations are a prohibitive factor. Furthermore, since the TEE would contain the entire population’s contract graph, it would be a single point of failure and an appealing target for an attack on TEE’s known vulnerabilities. RIPPLE_{TEE} (Fig. 7), on the other hand, leverages the presence of TEEs in participants’ mobile devices, but in a decentralised manner, ensuring that each TEE

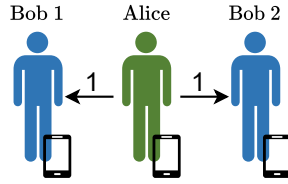


Figure 6: Sybil Attack.

handles only information related to the encounters made by the respective participant.

Before going into the details of RIPPLE_{TEE}, we will go over the $\mathcal{F}_{\text{anon}}$ functionality (cf. §B.3), which allows two participants, \mathcal{P}_i and \mathcal{P}_j , to send messages to each other anonymously via a set of communication servers denoted by \mathcal{C} . The set \mathcal{C} consists of one server acting as an entry node ($\mathcal{N}_{\text{entry}}$), receiving messages from senders, and one server acting as an exit node ($\mathcal{N}_{\text{exit}}$), forwarding messages to receivers. In $\mathcal{F}_{\text{anon}}$, sender \mathcal{P}_i does not learn to whom the message is sent, and receiver \mathcal{P}_j does not learn who sent it. Similarly, the servers in \mathcal{C} will be unable to relate receiver and sender of a message. Anonymous communication (cf. §A) is an active research area, e.g., [1, 5, 48, 59] and $\mathcal{F}_{\text{anon}}$ in RIPPLE_{TEE} can be instantiated using any of these efficient techniques.

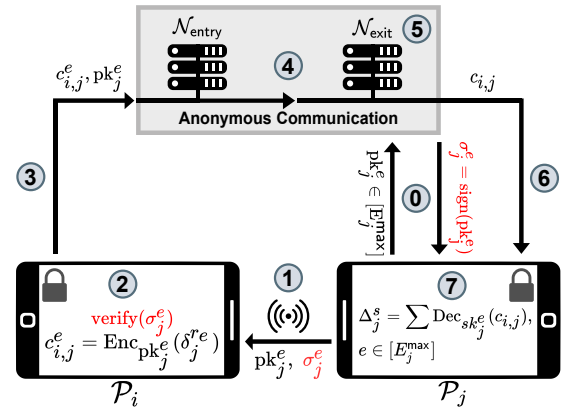


Figure 7: RIPPLE_{TEE} Overview. Messages in red denote additional steps needed for malicious participants.

Token Generation (cf. steps ① to ② in Fig. 7): During the pre-computation phase, the TEE of each participant $\mathcal{P}_i \in \mathcal{P}$ generates a list of fresh unique public/private key pairs (pk_i^e, sk_i^e) for all possible encounters $e \in [E_i^{\text{max}}]$. The keys, for example, can be generated and stored a day ahead of time. The newly generated public keys are then sent by \mathcal{P}_i ’s TEE to the exit node $\mathcal{N}_{\text{exit}}$ (cf. step ① in Fig. 7) to enable anonymous communication (cf. §B.3) via $\mathcal{F}_{\text{anon}}$ later in the protocol’s simulation part.

During a physical encounter e , \mathcal{P}_i and \mathcal{P}_j exchange two unused public keys pk_i^e and pk_j^e (cf. step ① in Fig. 7). Simultaneously, both participants compute and record metadata m_e , such as the time, location, and duration of the encounter, and store this information alongside the received public key.

Additional measures are required for malicious participants to ensure that the participants are exchanging public keys generated by the TEEs. After obtaining the new public keys from \mathcal{P}_i , the exit node $\mathcal{N}_{\text{exit}}$ goes one step further: It signs them and returns the signatures to \mathcal{P}_i after checking that it is connecting directly with a non-corrupted TEE (cf. step ① in Fig. 7 and §A). During a physical encounter, \mathcal{P}_i will provide the corresponding signature, denoted by σ_i^e along with pk_i^e so that the receiver \mathcal{P}_j can verify that the key was correctly generated by \mathcal{P}_i ’s TEE (cf. step ② in Fig. 7).

Simulation Execution (cf. steps ② to ⑦ in Fig. 7): All local computations, including infection likelihood calculation and infection class updates, will be performed within the participants’ TEEs. In

detail, for each encounter e involving participants \mathcal{P}_i and \mathcal{P}_j , the following steps are executed:

- \mathcal{P}_i 's TEE computes δ_i^e and encrypts it using the public key of \mathcal{P}_j obtained during the token generation phase. Let the ciphertext be $c_{i,j}^e = \text{Enc}_{\text{pk}_j^e}(\delta_i^e)$ (cf. step ② in Fig. 7).
- \mathcal{P}_i 's TEE establishes a connection with the anonymous communication servers C via remote attestation and uploads the tuple $(c_{i,j}^e, \text{pk}_j^e)$ to C (cf. step ③ in Fig. 7).
- The tuple $(c_{i,j}^e, \text{pk}_j^e)$ traverse through multiple servers in C and reaches the exit node $\mathcal{N}_{\text{exit}}$ (cf. step ④ in Fig. 7).
- $\mathcal{N}_{\text{exit}}$ discards the tuple if the public key pk_j^e has already been used in this simulation step⁷ (cf. step ⑤ in Fig. 7).
- If not, $\mathcal{N}_{\text{exit}}$ uses pk_j^e to identify the recipient \mathcal{P}_j and the ciphertext $c_{i,j}^e$ is sent to \mathcal{P}_j (cf. step ⑥ in Fig. 7).

After receiving the ciphertexts for all of the encounters, \mathcal{P}_j 's TEE decrypts them and aggregates the likelihoods to produce the desired output (cf. step ⑦ in Fig. 7).

4.1.1 Security of $\text{RIPPLE}_{\text{TEE}}$. First, we consider the case of semi-honest participants. During the token generation phase, since the current architecture in most mobile devices does not allow direct communication with a TEE while working with Bluetooth LE interfaces, participant \mathcal{P}_i can access both the sent and received public keys before they are processed in the TEE. However, unique keys are generated per encounter and do not reveal anything about an encounter's identities due to the security of the underlying \mathcal{F}_{gen} functionality, which captures the goal of several contract tracing apps in use.

The $\mathcal{F}_{\text{anon}}$ functionality, which implements an anonymous communication channel utilising the servers in C , aids in achieving *contact graph privacy* by preventing participants from learning to/from whom they are sending/receiving messages. While the entry node learns who sends messages to it, it does not learn who receives them. Similarly, the exit node $\mathcal{N}_{\text{exit}}$ has no knowledge of the sender but learns the recipient using the public key. Regarding *confidentiality*, participants in $\text{RIPPLE}_{\text{TEE}}$ have no knowledge of the messages being communicated because they cannot access the content of the TEEs and the TEEs communicate directly to the anonymous channel. Furthermore, servers in C will not have access to the messages as they are encrypted.

For the case of malicious participants, they could send specifically crafted keys during the token generation phase instead of the ones created by their TEE. However, this will make the signature verification fail and the encounter will get discarded. Furthermore, a malicious participant may reuse public keys for multiple encounters. This manipulation, however, will be useless because the exit node $\mathcal{N}_{\text{exit}}$ checks that each key is only used once before forwarding messages to participants. During the simulation phase, all data and computation are handled directly inside the TEEs of the participants, so no manipulation is possible other than cutting the network connection, i.e., dropping out of the simulation, ensuring *correctness*. Dropouts occur naturally when working with mobile devices and have no effect on privacy guarantees.

⁷ This step is not required for semi-honest participants.

4.2 $\text{RIPPLE}_{\text{PIR}}$

In the following, we show how to get rid of $\text{RIPPLE}_{\text{TEE}}$'s assumption of each participant having access to a TEE on their mobile devices. If we simply remove the TEE part of $\text{RIPPLE}_{\text{TEE}}$ and run the same protocol, decryption and aggregation of a participant's received infection likelihoods would be under their control. Thus, the individual infection likelihoods of all encounters would be known to them, leaking information about the contact graph (cf. §3.2). To get around this privacy issue, we need to find another way to aggregate the infection likelihoods so that only the sum, not individual values, can be derived by the participants.

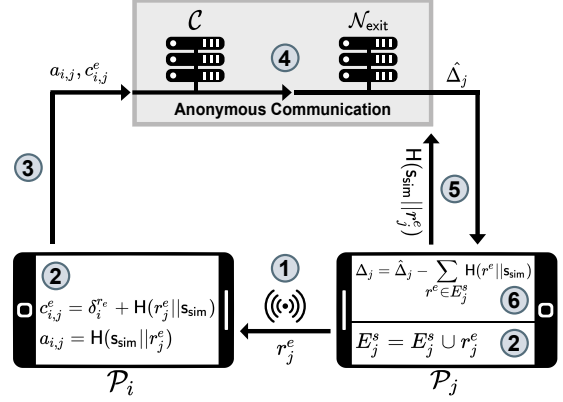


Figure 8: $\text{RIPPLE}_{\text{PIR}}$ Overview.

Private Information Retrieval (PIR, cf. §A) is one promising solution for allowing participants to anonymously retrieve infection likelihoods sent to them. PIR enables the private download of an item from a public database D held by M servers without leaking any information to the servers, such as which item is queried or the content of the queried item. However, the normal textbook PIR is unsuitable for our needs because we need to retrieve the sum of, say, τ items from the database rather than the individual ones. As a result, we introduce the ideal functionality $\mathcal{F}_{\text{pirsum}}$ (Fig. 9), which is similar to a conventional PIR functionality, but returns the sum of τ queried locations of the database as the result. For the remainder of this section, we consider $\mathcal{F}_{\text{pirsum}}$ to be an ideal black-box, and will discuss concrete instantiations in §5.

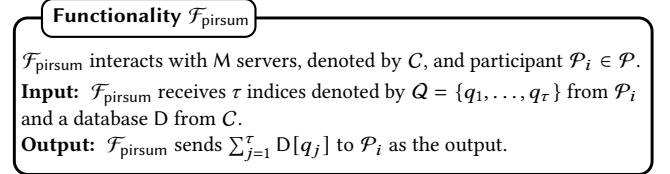


Figure 9: Ideal functionality for PIR-SUM (semi-honest).

We now detail the changes needed in the token generation phase to make it compatible with the rest of the $\text{RIPPLE}_{\text{PIR}}$ protocol.

Token Generation (cf. step ① in Fig. 8): During a physical encounter e among participants \mathcal{P}_i and \mathcal{P}_j , they generate and exchange unique random tokens denoted by r_i^e and r_j^e . Both participants, like in $\text{RIPPLE}_{\text{TEE}}$, record the metadata m^e as well. Thus, at the end of a simulation step $s \in [N_{\text{step}}]$ (e.g., a day), \mathcal{P}_i holds a list of sent encounter tokens, denoted by $E_i^s = \{r_i^e\}_{e \in \mathcal{E}_i}$, and a list of

received tokens, denoted by $R_i^s = \{r_j^e\}_{e \in \mathcal{E}_i}$. Looking ahead, these random tokens will be used as addresses for communicating the corresponding infection likelihood among the participants.

Simulation Execution (cf. steps ② to ⑥ in Fig. 8): Local computations such as encounter filtering and infection likelihood calculation proceed similarly to RIPPLETEE but without TEE protection. The steps for an encounter e among \mathcal{P}_i and \mathcal{P}_j are as follows:

- \mathcal{P}_i blinds each $\delta_i^{r^e}$ computed with the corresponding random token r_j^e received from \mathcal{P}_j and obtains the ciphertext $c_{i,j}^e = \delta_i^{r^e} + H(r_j^e || s_{\text{sim}})$. In addition, it computes the destination address for the ciphertext as $a_{i,j} = H(s_{\text{sim}} || r_j^e)$. Here, $H(\cdot)$ is a cryptographic hash function and $s_{\text{sim}} \in [N_{\text{sim}}]$ denotes the current simulation setting. (cf. step ② in Fig. 8)

- Note that s_{sim} is used in $H(\cdot)$ to ensure that distinct (ciphertext, address) tuples are generated for the same encounters across multiple simulation settings, preventing the exit node N_{exit} from potentially linking messages from different simulations.

- \mathcal{P}_i sends the tuple $(a_{i,j}, c_{i,j}^e)$ anonymously to N_{exit} with the help of the servers in C . N_{exit} discards all the tuples with the same address field $(a_{i,j})$. (cf. step ③ to ④ in Fig. 8)

As a server in C , N_{exit} locally creates the database D for the current simulation step using all of the $(c_{i,j}^e, a_{i,j})$ tuples received (part of step ④ in Fig. 8). A naive solution of inserting $c_{i,j}^e$ using a simple hashing of the address $a_{i,j}$ will not provide an efficient solution in our case since we require only one message to be stored in each database entry in order to have an injective mapping between addresses and messages. This is required in order for the message receiver to precisely download the messages that were sent to them. Using simple hashing, this would translate to a large database size to ensure a negligible probability of collisions. Instead, in RIPPLETEE, we use a novel variant of a garbled cuckoo table, that we call arithmetic garbled cuckoo table (AGCT, cf. §B.4), with $a_{i,j}$ as the insertion key for the database.

Once the database D is created, N_{exit} sends it to the other servers in C based on the instantiation of $\mathcal{F}_{\text{pirsum}}$ (cf. §5). Each $\mathcal{P}_j \in \mathcal{P}$ will then participate in an instance of $\mathcal{F}_{\text{pirsum}}$ with the servers in C acting as PIR servers holding the database D . \mathcal{P}_j uses the addresses of all its sent encounters from E_j^s , namely $H(s_{\text{sim}} || r^e)$, as the input to $\mathcal{F}_{\text{pirsum}}$ and obtains a blinded version of the cumulative infection likelihood, denoted by $\hat{\Delta}_j$, as the output (cf. step ⑤ in Fig. 8). The cumulative infection likelihood, Δ_j , is then unblinded as

$$\Delta_j = \hat{\Delta}_j - \sum_{r^e \in E_j^s} H(r^e || s_{\text{sim}})$$

concluding the current simulation step (cf. step ⑥ in Fig. 8).

4.2.1 Security of RIPPLETEE. Except for the database constructions at exit node N_{exit} and the subsequent invocation of the $\mathcal{F}_{\text{pirsum}}$ functionality for the cumulative infection likelihood computation, the security guarantees for semi-honest participants in RIPPLETEE are similar to those of RIPPLETEE. Unlike RIPPLETEE, N_{exit} in RIPPLETEE will be unable to identify the message’s destination from the address because it will be known only to the receiving participant. Furthermore, each participant obtains the cumulative infection likelihood directly via the $\mathcal{F}_{\text{pirsum}}$ functionality, ensuring

that N_{exit} cannot infer the participant’s encounter details and, thus, *contact graph privacy*.

Malicious participants in RIPPLETEE, as opposed to RIPPLETEE, can tamper with the protocol’s correctness by providing incorrect inputs. However, as stated in the threat model in §3, we assume that malicious participants in our framework will not tamper with the correctness and will only seek additional information. A malicious participant could re-use the same encounter token for multiple encounters during token generation. As a result, the protocol will generate multiple tuples with the same address. However, as stated in the protocol, N_{exit} will discard all such tuples, effectively removing the malicious participant from the system. Another potential information leakage caused by the participant’s aforementioned action is that the entry point of the anonymous communication channel will be able to deduce that multiple participants, say \mathcal{P}_i and \mathcal{P}_j , had an encounter with the same participant. This is not an issue in our protocol because we instantiate the $\mathcal{F}_{\text{anon}}$ functionality using a 3-server oblivious shuffling scheme (cf. §B.3), where all the servers except N_{exit} will not see any messages in the clear, but only see secret shares.

5 PIR-SUM: INSTANTIATING $\mathcal{F}_{\text{pirsum}}$

So far, the discussion has focused on RIPPLETEE as a generic framework composed of multiple ideal functionalities that could be efficiently instantiated using state-of-the-art privacy-enhancing technologies. In this section, we will concentrate on instantiating our novel $\mathcal{F}_{\text{pirsum}}$ functionality (Fig. 9) using three semi-honest MPC servers. In particular, we have three servers S_0, S_1 , and S_2 , and we design the PIR-sum protocol to instantiate the $\mathcal{F}_{\text{pirsum}}$ functionality.

The problem statement in our context is formally defined as follows: Participant $\mathcal{P}_i \in \mathcal{P}$ has a set of τ indices denoted by $Q = \{q_1, \dots, q_\tau\}$ and wants to retrieve $\text{res} = \sum_{q \in Q} D[q]$. In this case, D is a database with N elements of ℓ -bits each that is held in clear by both the servers S_1 and S_2 . The server S_0 aids in the computation performed by the servers S_1 and S_2 . Furthermore, we assume a one-time setup (cf. §B.1) among the servers and \mathcal{P}_i that establishes shared pseudorandom keys among them to facilitate non-interactive generation of random values and, thus, save communication [9, 25, 101].

5.1 Overview of PIR-sum protocol

At a high level, the idea is to use multiple instances of a standard 2-server PIR functionality [17, 32], denoted by $\mathcal{F}_{\text{pir}}^{2S}$, and combine the responses to get the sum of the desired blocks as the output. $D^m = D + m$ denotes a modified version of the database D in which every block is summed with an ℓ -bit mask m , i.e., $D^m[i] = D[i] + m$ for $i \in [N]$. The protocol proceeds as follows:

- S_1 and S_2 non-interactively sample τ random mask values $\{m_1, \dots, m_\tau\}$ such that $\sum_{j=1}^\tau m_j = 0$.

- S_1, S_2 , and \mathcal{P}_i execute τ instances of $\mathcal{F}_{\text{pir}}^{2S}$ in parallel, with servers using D^{m_j} as the database and \mathcal{P}_i using q_j as the query for the j -th instance for $j \in [\tau]$.

- Let res_j denote the result obtained by \mathcal{P}_i from the j -th $\mathcal{F}_{\text{pir}}^{2S}$ instance. \mathcal{P}_i locally computes $\sum_{j=1}^\tau \text{res}_j$ to obtain the desired result.

The details for instantiating $\mathcal{F}_{\text{pir}}^{2S}$ using the standard linear summation PIR approach [32] are provided in §C.1. The approach requires \mathcal{P}_i to communicate $N \cdot \tau$ bits to the servers, which is further reduced in $\text{RIPPLE}_{\text{PIR}}$ as shown in §5.2.

Security of PIR_{sum} . While it is simple to show that the above solution is adequate for semi-honest participants, malicious participants must be dealt with separately. A malicious participant, for example, could use the same query, say q_j , in all τ instances and retrieve only the block corresponding to q_j by dividing the result by τ . We use a simple verification scheme over the $\mathcal{F}_{\text{pir}}^{2S}$ functionality to prevent these manipulations. Its details are presented next.

For malicious participants, we want to ensure that \mathcal{P}_i used a distinct vector \vec{b} during the τ parallel instances. One naive approach is to have S_1 and S_2 compute the bitwise-OR of all the τ bit vectors $\vec{b}_1, \dots, \vec{b}_\tau$, and then run a secure two-party computation protocol to compare the number of ones in the resultant vector to τ . We use the additional server S_0 to further optimize this step. S_1 and S_2 send randomly shuffled versions of their secret shared bit vectors to S_0 , who reconstructs the shuffled vectors and perform the verification locally. This approach leaks no information to S_0 because it has no information about the underlying database D . The verification procedure is as follows:

- S_1 and S_2 non-interactively agree on a random permutation, denoted by π .
- S_u sends $\pi([\vec{b}_j]_u)$ to S_0 for $j \in [\tau]$ and $u \in \{1, 2\}$.
- S_0 locally reconstructs $\pi(\vec{b}_j) = \pi([\vec{b}_j]_1) \oplus \pi([\vec{b}_j]_2)$, for $j \in [\tau]$. If all the τ bit vectors are correctly formed and distinct, it sends Accept to S_1 and S_2 . Else, it sends abort.

Note that the verification using \mathcal{P}_0 will incur a communication of $2\tau N$ bits among the servers. Furthermore, the above verification method can be applied to any instantiation of $\mathcal{F}_{\text{pir}}^{2S}$ that generates a boolean sharing of the query bit vector among the PIR servers and computes the response as described above, e.g., the PIR schemes of [16, 17, 32].

5.2 Reducing participant’s communication

PIR_{sum} is realized in $\text{RIPPLE}_{\text{PIR}}$ using two different approaches, each with its own set of trade-offs, with a goal of minimizing the communication at the participant’s end. While the first approach, denoted by $\text{PIR}_{\text{sum}}^I$, sacrifices computation for better communication, the second approach, denoted by $\text{PIR}_{\text{sum}}^{II}$, reduces both the computational and communication overhead of the participant in $\text{PIR}_{\text{sum}}^I$ with the help of additional server $S_0 \in C$.

5.2.1 $\text{PIR}_{\text{sum}}^I$ (Fig. 18 in §C.2.2). In this approach, we instantiate $\mathcal{F}_{\text{pir}}^{2S}$ using PIR techniques based on Function Secret Sharing (FSS) [16, 17, 36]. To retrieve the q -th block from the database, \mathcal{P}_i uses FSS on a Distributed Point Function (DPF) [54] that evaluates to a 1 only when the input q is 1 and to 0 otherwise. \mathcal{P}_i generates two DPF keys k_1 and k_2 that satisfy the above constraint and sends one key to each of the servers S_1 and S_2 . The servers S_1 and S_2 can then locally expand their key share to obtain their share for the bit vector \vec{b} and the rest of the procedure proceeds similarly to the naive linear summation method discussed in §5.1 (more details on

Linear Summation PIR are given in §C.1). The key size for a database of size N records using the optimised DPF construction in [17] is about $\lambda \log(N/\lambda)$, where $\lambda = 128$ for an AES-based implementation. We provide more details about FSS-based PIR in §C.1.1.

Security. For semi-honest participants, the security of our method directly reduces to that of the 2-server PIR protocol in [17]. However, as mentioned in [17], a malicious participant could generate incorrect DPF keys, putting the scheme’s security as well as the correctness at risk. To prevent this type of misbehaviour, Boyle et al. [17] present a form of DPF called “verifiable DPF”, which can assure the correctness of the DPF keys created by \mathcal{P}_i at the cost of an increased *constant* amount of communication between the servers.

While using verifiable DPFs in $\text{PIR}_{\text{sum}}^I$ ensures that the τ bit vectors generated by \mathcal{P}_i are valid, it does not ensure that the bit vectors $\vec{b}_1, \dots, \vec{b}_\tau$ correspond to τ distinct locations in the database D . However, we leverage the correctness guarantee of verifiable DPFs to reduce the communication cost for verification, as discussed in §5.1. In detail, all τ bit vectors $\vec{b}_1, \dots, \vec{b}_\tau$, i.e., the PIR queries, that are available in a secret-shared form among S_1 and S_2 are now guaranteed to have exactly one 1 in them, with the remaining bit positions being 0. To ensure distinctness, S_1 and S_2 XORs all their respective τ shares locally to obtain the secret-share of a single vector $\vec{b}_c = \bigoplus_{k=1}^{\tau} \vec{b}_k$. The problem now boils down to determining whether or not \vec{b}_c has exactly τ bit positions set to 1. This can be accomplished by servers S_1 and S_2 agreeing on a random permutation π and reconstructing $\pi(\vec{b}_c)$ to S_0 and allowing S_0 to perform the check, as in the naive approach (cf. §5.1).

Computation Complexity (#AES operations). In $\text{PIR}_{\text{sum}}^I$, the participant \mathcal{P}_i must perform $4 \cdot \log(N/\lambda)$ AES operations as part of the key generation algorithm for each of the τ instances of $\mathcal{F}_{\text{pir}}^{2S}$ over a database of size N , where $\lambda = 128$ for an AES-based implementation. Similarly, S_1 and S_2 must perform $\log(N/\lambda)$ AES operations for each of the N DPF evaluations. We refer to Table 1 in [17] for more specifics.

5.2.2 $\text{PIR}_{\text{sum}}^{II}$ (Fig. 10). In this approach, we use the server S_0 to reduce the computation and communication of the participant \mathcal{P}_i in $\text{PIR}_{\text{sum}}^I$. The idea is that S_0 plays the role of \mathcal{P}_i for the PIR-protocol in $\text{PIR}_{\text{sum}}^I$. However, \mathcal{P}_i cannot send its query, say q , to S_0 in clear because it would violate privacy. As a result, \mathcal{P}_i selects random values $q', \theta_q \in [N]$ such that $q = q' + \theta_q$. In this case, q' is a *shifted version* of the index q , and θ is a *shift correction* for q . \mathcal{P}_i sends q' to S_0 and θ_q to both S_1 and S_2 . The remainder of the computation until output retrieval will now take place solely among the servers.

The servers run a DPF instance [17] with S_0 acting as the client and input query q' . At the end of the computation, S_1 and S_2 obtain the bit vector $\vec{b}_{q'}$, which corresponds to q' . However, as discussed in $\text{PIR}_{\text{sum}}^I$, the servers require an XOR sharing corresponding to the actual query q in order to continue the computation. S_1 and S_2 do this by using the shift correction value θ_q received from \mathcal{P}_i . Both S_1 and S_2 will perform a right cyclic shift of their $\vec{b}_{q'}$ shares by θ_q units. A negative value for θ_q indicates a cyclic shift to the left.

It is easy to see that the XOR shares obtained after the cyclic shift correspond to the bit vector \vec{b}_q . To further optimise \mathcal{P}_i ’s communication, \mathcal{P}_i and servers S_1, S_2 non-interactively generate random

shift correction values θ_q first using the shared-key setup (cf. §B.1), and only the corresponding q' values are communicated to S_0 . The rest of the protocol is similar to $\text{PIR}_{\text{sum}}^{\text{I}}$, and the formal protocol is shown in Fig. 10. In terms of malicious participants, $\text{PIR}_{\text{sum}}^{\text{II}}$ has an advantage over $\text{PIR}_{\text{sum}}^{\text{I}}$ as there is no need to use a verifiable DPF to protect against malicious \mathcal{P}_i , because the semi-honest server S_0 generates the DPF key instead of \mathcal{P}_i .

Protocol $\text{PIR}_{\text{sum}}^{\text{II}}$	
<i>Input(s):</i> i) $S_1, S_2 : D; D = N$, ii) $\mathcal{P}_i : Q = \{q_1, \dots, q_\tau\}$, and iii) $S_0 : \perp$.	<i>Output:</i> $\mathcal{P}_i : \text{res} = \sum_{q \in Q} D[q]$
Computation S_1 and S_2 sample τ random mask values $\{m_1, \dots, m_\tau\} \in \mathbb{Z}_{2^\ell}^\tau$ such that $\sum_{j=1}^\tau m_j = 0$. For each $q \in Q$, execute the following: <ol style="list-style-type: none"> 1. S_1, S_2 locally compute $D^{m_q} = D + m_q$, i.e., $D^{m_q}[j] = D[j] + m_q$, for $j \in [N]$. 2. \mathcal{P}_i, S_1, S_2 sample random $\theta_q \in [N]$. 3. \mathcal{P}_i computes and sends $q' = q - \theta_q$ to S_0. 4. Servers execute DPF protocol [17] with S_0 as client with input q'. Server S_u obtains $[\vec{b}_{q'}]_u$ with $b_{q'}^j = 1$ for $j = q'$ and $b_{q'}^j = 0$ for $j \neq q'$, for $u \in \{1, 2\}$. 5. S_u locally applies θ_u on $[\vec{b}_{q'}]_u$ to generate $[\vec{b}_q]_u$, for $u \in \{1, 2\}$. 	
Verification Let $\{\vec{b}_{q_1}, \dots, \vec{b}_{q_\tau}\}$ denote the bit vectors whose XOR-shares are generated during the preceding steps. <ol style="list-style-type: none"> 6. S_k computes $[\vec{b}_c]_k = \bigoplus_{q \in Q} [\vec{b}_q]_k$, for $u \in \{1, 2\}$ 7. S_1 and S_2 non-interactively agree on random permutation π. 8. S_u sends $\pi([\vec{b}_c]_u)$ to S_0, for $u \in \{1, 2\}$. 9. S_0 locally reconstructs $\pi(\vec{b}_c) = \pi([\vec{b}_c]_1) \oplus \pi([\vec{b}_c]_2)$. It sends Accept to S_1 and S_2, if $\pi(\vec{b}_c)$ has exactly τ ones. Else, it sends abort. 	
Output Transfer Send \perp to \mathcal{P}_i if S_0 generated abort during verification. Otherwise, proceed as follows: <ol style="list-style-type: none"> 10. S_u sends $[y_q]_u = \bigoplus_{j=1}^N [b_{q'}^j]_u D^{m_q}[j]$ to \mathcal{P}_i, for $q \in Q, u \in \{1, 2\}$. 11. \mathcal{P}_i locally computes $\text{res} = \sum_{q \in Q} ([y_q]_1 \oplus [y_q]_2)$. 	

Figure 10: $\text{PIR}_{\text{sum}}^{\text{II}}$ Protocol.

5.2.3 *Summary of communication costs.* Tab. 2 summarises the communication cost for our two PIR_{sum} approaches for instantiating $\mathcal{F}_{\text{pirsum}}$ over a database of size N with τ PIR queries per client. More concrete details along with the security guarantees are provided in §C.

Stage	$\text{PIR}_{\text{sum}}^{\text{I}}$	$\text{PIR}_{\text{sum}}^{\text{II}}$
\mathcal{P}_i to servers in C	$2\tau(\lambda + 2) \log(N/\lambda) + 4\tau\lambda$	$\tau \log N$
Server to server	0	$2\tau(\lambda + 2) \log(N/\lambda) + 4\tau\lambda$
Servers in C to \mathcal{P}_i	$\tau \cdot 2\ell$	$\tau \cdot 2\ell$
+ Verification (mal.)	$2N + 2 + \delta$	$2N + 2$

Table 2: Summary of communication costs for PIR_{sum} . Here, λ denotes the AES key size ($\lambda = 128$ in [16]), ℓ denotes the block size in bits ($\ell = 128$ in this work) and δ denotes the constant involved in the verifiable DPF approach [17] (cf. §C.1.1).

6 EVALUATION

In this section, we evaluate and compare the computation and communication efficiency of our two RIPPLE protocols presented in §4. A fully-fledged implementation, similar to existing contact tracing apps, would necessitate collaboration with industry partners to develop a real-world scalable system for national deployment. Instead, we carry out a proof-of-concept implementation and provide micro benchmark results for all major building blocks.⁸ We focus on the simulation phase for benchmarking, which is separate from the token generation phase. The simulations can ideally be done overnight while mobile phones are charging and have access to a high-bandwidth WiFi connection. According to studies [127, 129], sleeping habits in various countries provide a time window of several hours each night that can be used for this purpose.

Setup and Parameters. We run the benchmarks on the server-side with three servers (two for FSS-PIR and one as helper server as discussed in §5.2) with Intel Core i9-7960X CPUs@2.8 GHz and 128 GB RAM connected with 10 Gbit/s LAN and 0.1 s RTT. The client is a Samsung Galaxy S10+ with an Exynos 9820@2.73 GHz and 8GB RAM. As Android does not allow third-party developers to implement applications for Android’s TEE Trusty [7], we use hardware-backed crypto operations already implemented by Android instead. We use the code of [70] to instantiate FSS-PIR. We implement the AGCT in C++ and follow previous work on cuckoo hashing [108] by using tabulation hashing for the hash functions.

We instantiate our protocols in RIPPLE with $\kappa = 128$ bit security. We use RSA-2048 as the encryption scheme in RIPPLE_{TEE} since Android offers a hardware-backed implementation. We omit the overhead of remote attestation for the sake of simplicity. For RIPPLE_{PIR}, we use the FSS-PIR scheme of [17, 70] as the baseline and the addresses are hashed with SHA-256 and trimmed to $40 - 1 + \log_2(p \cdot E^{\text{avg}})$, where p represents the number of participants and E^{avg} represents the average number of encounters per participant per simulation step. We set $E^{\text{avg}} = 100$ while benchmarking based on numbers provided by research on epidemiological modeling [40, 95]. To avoid cycles when inserting n messages into the AGCT (cf. §B.4), we set its size to $10n$. This can be further improved as discussed in §B.4 [105, 107, 108]. A typical simulation step would be one day such that 14 simulation steps can simulate two weeks.

6.1 Communication Costs

In this section, we look at the communication costs that our protocols incur. We consider participant sizes, denoted by p , ranging from thousand (1K) to twenty million (20M) to analyse the scalability of our protocols. Tab. 3 summarises the communication costs incurred by each participant as well as the communication servers (C) for one simulation step in a specific simulation. It should be noted that one simulation step includes all protocol steps, beginning with participants locally computing their infection likelihood (δ) and ending with them obtaining their cumulative infection likelihood (Δ) for that step.

⁸ Note that we are not attempting to create the most efficient instantiation. More optimizations will undoubtedly improve efficiency, and our protocols can be heavily parallelized with a large number of servers. Instead, our goal here is to demonstrate the viability of RIPPLE protocols for large-scale deployment.

Entities	Protocol	Population (p)									
		1K	10K	50K	100K	500K	1M	2M	5M	10M	20M
Participants in \mathcal{P} (in KB)	RIPPLE _{TEE} (§4.1)	16.00	16.00	16.00	16.00	16.00	16.00	16.00	16.00	16.00	16.00
	RIPPLE _{PIR} : PIR _{sum} ^I (§5.2.1)	51.63	62.42	69.97	73.22	80.77	84.02	87.27	91.56	94.81	98.06
	RIPPLE _{PIR} : PIR _{sum} ^{II} (§5.2.2)	3.45	3.49	3.52	3.53	3.56	3.57	3.59	3.60	3.62	3.63
Servers in \mathcal{C} (in GB)	RIPPLE _{TEE} (§4.1)	0.02	0.19	0.96	9.60	19.20	1.92	38.40	96.00	192.00	384.00
	RIPPLE _{PIR} (§5)	0.01	0.10	0.48	0.96	4.80	9.60	19.20	48.00	96.00	192.00

Table 3: Communication costs per simulation step in our RIPPLE instantiations.

Participant Communication. A participant in RIPPLE_{TEE} requires just 16KB of total communication every simulation step, as indicated in Tab. 3, and is independent of the population on which the simulation is done. This is because each participant will only send and receive infection likelihood messages related to its encounters. While the value in the table corresponds to an average of 100 encounters ($E^{\text{avg}} = 100$), we depict the participants’ communication in Fig. 11 with varied E^{avg} ranging from 10 to 500 for a population of 10M. Note that a 2-week simulation with $E^{\text{avg}} = 500$ can be completed by a participant in RIPPLE_{TEE} with roughly 1MB of communication.

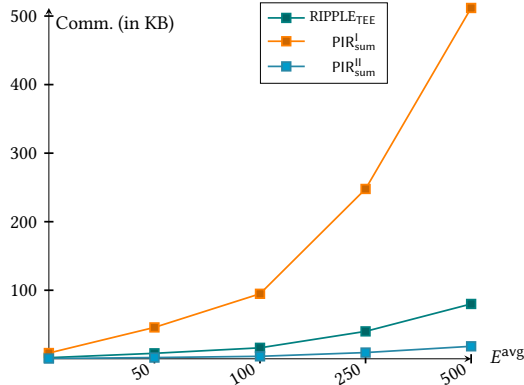


Figure 11: Participant’s communication with varying E^{avg} for a population of 10M.

Unlike RIPPLE_{TEE}, participant communication in both PIR_{sum}^I and PIR_{sum}^{II} increases for larger populations as the corresponding database size increases. The communication, however, is only sub-linear in the database size⁹. In particular, the participant’s communication in PIR_{sum}^I ranges from 51.63KB to 98.06KB, with the higher cost over RIPPLE_{TEE} attributed to the size of DPF keys used in the underlying FSS-PIR scheme [17], as discussed in §5. The communication in PIR_{sum}^{II}, on the other hand, is in the range of 3.5KB for all of the participants size under consideration. This reduced communication is due to the optimization in PIR_{sum}^{II}, which offloads the DPF key generation task to the helper server S_0 (cf. §5.2.2). A

⁹ The DB size and communication costs of RIPPLE_{PIR} can be reduced by optimizing the database size by extending the database by only $d + \lambda$ bins, where d is the upper bound of double collisions and λ is an error parameter, instead of $10n$ (cf. §B.4). [105]

participant in PIR_{sum}^I must communicate approximately 7MB of data for a 2-week simulation for a 10M population with $E^{\text{avg}} = 500$, whereas it is only 0.25MB in the case of PIR_{sum}^{II}.

Population (p)	Protocol	E^{avg}				
		10	50	100	250	500
100K	RIPPLE _{TEE} (§4.1)	1.60	8.00	16.00	40.00	80.00
	RIPPLE _{PIR} : PIR _{sum} ^I (§5.2.1)	6.24	34.99	73.22	193.79	403.83
	RIPPLE _{PIR} : PIR _{sum} ^{II} (§5.2.2)	0.35	1.76	3.53	8.87	17.81
1M	RIPPLE _{TEE} (§4.1)	1.60	8.00	16.00	40.00	80.00
	RIPPLE _{PIR} : PIR _{sum} ^I (§5.2.1)	7.32	40.38	84.02	220.78	457.81
	RIPPLE _{PIR} : PIR _{sum} ^{II} (§5.2.2)	0.35	1.78	3.57	8.98	18.01
10M	RIPPLE _{TEE} (§4.1)	1.60	8.00	16.00	40.00	80.00
	RIPPLE _{PIR} : PIR _{sum} ^I (§5.2.1)	8.40	45.78	94.81	247.77	511.79
	RIPPLE _{PIR} : PIR _{sum} ^{II} (§5.2.2)	0.36	1.80	3.62	9.08	18.22

Table 4: Communication per participant (in KB) in a simulation step over different E^{avg} and population sizes.

Tab. 4 provides the communication cost for a participant for multiple population sizes while varying the average number of encounters (E^{avg}) per simulation step from 10 to 500. The communication cost in RIPPLE_{TEE} is independent of population size and grows linearly in E^{avg} . A similar trend can be seen in RIPPLE_{PIR} with the exception that the cost increases sub-linearly with population size due to the use of FSS-based PIR scheme in RIPPLE_{PIR}.

Server Communication. The servers’ communication is primarily attributed to the anonymous communication channel that they have established, which provides unlinkability and, thus, privacy to the messages of the participants. As discussed in §B.3, in order to communicate M messages through the channel, the servers must communicate $2M$ messages in RIPPLE_{TEE}, and $3M$ messages in RIPPLE_{PIR}. When it comes to concrete values, however, the server communication in RIPPLE_{PIR} is half that of RIPPLE_{TEE}, as shown in Tab. 3. This is due to the larger message size in RIPPLE_{TEE} as a result of the use of public-key encryption.

For a population of 10M, the servers in RIPPLE_{TEE} must communicate 192GB of data among themselves, whereas RIPPLE_{PIR} requires 96GB. Setting the proper bit length for the address field in the messages can further reduce communication. For example, a population of 20M with $E^{\text{avg}} = 100$ can be accommodated in a 70-bit address field. Using this optimization will result in an additional

23 percent reduction in communication at the servers, as shown in Tab. 7 (cf. §D). Fig. 12 captures these observations better, and Tab. 7 and Tab. 4 in §D provide a detailed analysis of the concrete communication costs.

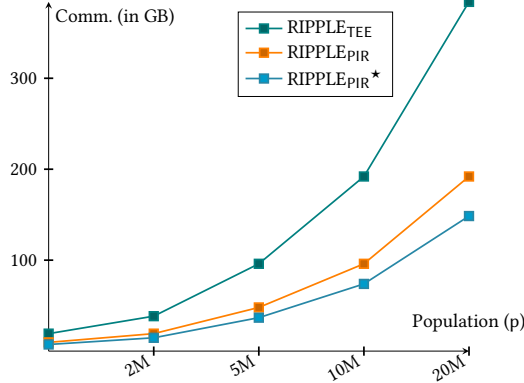


Figure 12: Communication costs for servers per simulation step for varying population. * denotes the results for optimized bit addresses in $\text{RIPPLE}_{\text{PIR}}$ (cf. Tab. 7).

6.2 Computation Costs and Runtime

This section focuses on the runtime, which includes time for computation as well as communication between entities. Tab. 5 summarizes the computation time with respect to a participant \mathcal{P}_i for a two-week simulation over a half-million population. The longer computation time in $\text{RIPPLE}_{\text{TEE}}$, as shown in Tab. 5, is due to the public key encryption and decryption that occurs within the mobile device’s TEE. This cost, however, is independent of population size and scales linearly with the average number of encounters, denoted by E^{avg} . In particular, for a 14-days simulation with a population of half million, \mathcal{P}_i in $\text{RIPPLE}_{\text{TEE}}$ needs approximately 43.7 seconds of computation time to perform the encryption and decryption tasks and may require additional time for the remote attestation procedure, which is not covered in our benchmarks. \mathcal{P}_i ’s computation time in $\text{RIPPLE}_{\text{PIR}}$, on the other hand, is significantly lower and is at most 5 milliseconds for the case of $\text{PIR}_{\text{sum}}^{\text{I}}$, while it increases to around 165 milliseconds for the case of $\text{PIR}_{\text{sum}}^{\text{I}}$. The increased computation time in $\text{PIR}_{\text{sum}}^{\text{I}}$ is due to DPF key generation, which scales sub-linearly with population size.

	Per Simulation Step			Per Simulation ($N_{\text{step}} = 14$)		
	Message Generation (in ms)	PIR Queries (in ms)	Output Computation (in ms)	Message Generation (in sec)	PIR Queries (in sec)	Output Computation (in sec)
$\text{RIPPLE}_{\text{TEE}}$	80.00	-	3040.00	1.12	-	42.56
$\text{PIR}_{\text{sum}}^{\text{I}}$	0.30	11.73	4.8e-2	4.26e-3	0.16	6.72e-4
$\text{PIR}_{\text{sum}}^{\text{II}}$	0.30	3.0e-3	4.8e-2	4.26e-3	4.2e-5	6.72e-4

Table 5: Average participant computation times per simulation step distributed across various tasks. Values are obtained using a mobile for a population of 500K with $E^{\text{avg}} = 100$.

Fig. 13 plots the overall runtime of the two instantiations in RIPPLE over population ranging from 1K to 500K. Similarly, in

Fig. 14, we plot the overall runtime of our two instantiations in RIPPLE for a full simulation of 2 weeks over various populations ranging from 1K to 500K. After a population of 100K, the runtime of $\text{RIPPLE}_{\text{PIR}}$ begins to exceed that of $\text{RIPPLE}_{\text{TEE}}$ due to an increase in database size, which results in longer data transfer times. More details regarding computation time is presented in Tab. 6 in §D.

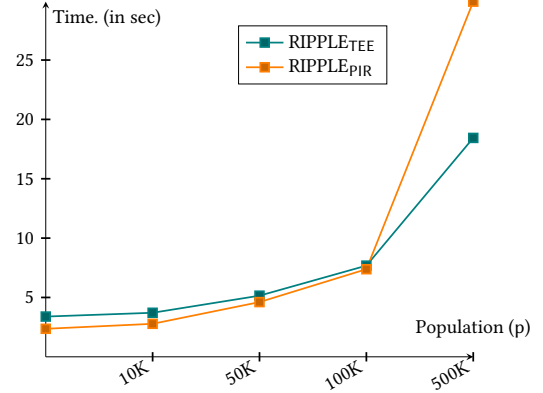


Figure 13: Runtime per simulation step in RIPPLE (1 day).

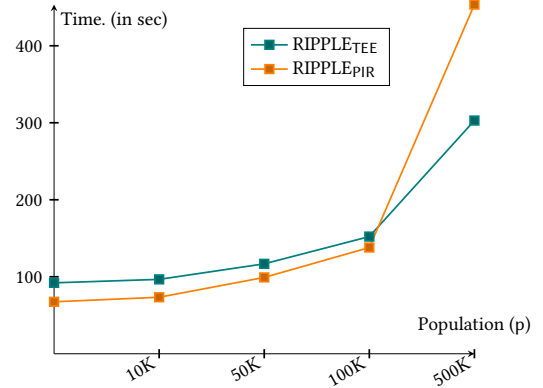


Figure 14: Runtime per simulation in RIPPLE (14 days).

Our benchmarking using the proof-of-concept implementation demonstrated the RIPPLE framework’s viability for real-world adaptation. One of the key benefits of our approaches is that participants have very little work to do. The system’s efficiency can be increased with appropriate hardware and optimized implementations.

Code availability. Available at DOI: [10.5281/zenodo.6595449](https://doi.org/10.5281/zenodo.6595449).

ACKNOWLEDGEMENTS

This project received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 850990 PSOTI). It was co-funded by the Deutsche Forschungsgemeinschaft (DFG) within SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251805230, and by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within ATHENE.

REFERENCES

- [1] Ittai Abraham, Benny Pinkas, and Avishay Yanai. 2020. Blinder: MPC Based Scalable and Robust Anonymous Committed Broadcast. In *ACM CCS*.
- [2] David Adam. 2020. Special report: The simulations driving the world's response to COVID-19. *Nature* (2020).
- [3] Nadeem Ahmed, Regio A Michelin, Wanli Xue, Sushmita Ruj, Robert Malaney, Salil S Kanhere, Aruna Seneviratne, Wen Hu, Helge Janicke, and Sanjay K Jha. 2020. A Survey of COVID-19 Contact Tracing Apps. *IEEE Access* (2020).
- [4] Fadi Al-Turjman and Bakkiam David Deebak. 2020. Privacy-Aware Energy-Efficient Framework Using the Internet of Medical Things for COVID-19. *IEEE Internet Things Mag.* 3, 3 (2020).
- [5] Nikolaos Alexopoulos, Aggelos Kiayias, Riivo Talviste, and Thomas Zacharias. 2017. MCMix: Anonymous Messaging via Secure Multiparty Computation. In *USENIX Security*.
- [6] Yaniv Altshuler, Nadav Aharoni, Micky Fire, Yuval Elovici, and Alex Pentland. 2012. Incremental Learning with Accuracy Prediction of Social and Individual Properties from Mobile-Phone Data. In *International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*.
- [7] Android. unk. Third-party Trusty applications. <https://source.android.com/security/trusty>.
- [8] Sebastian Angel, Hao Chen, Kim Laine, and Srinath Setty. 2018. PIR with Compressed Queries and Amortized Query Processing. In *S&P*.
- [9] Toshinori Araki, Jun Furukawa, Kazuma Ohara, Benny Pinkas, Hanan Rosemarin, and Hikaru Tsuchida. 2021. Secure Graph Analysis at Scale. In *ACM CCS*.
- [10] ARM. 2009. ARM security technology building a secure system using TrustZone technology. <https://developer.arm.com/documentation/genc009492/c>.
- [11] Yonatan Aumann and Yehuda Lindell. 2010. Security Against Covert Adversaries: Efficient Protocols for Realistic Adversaries. *Journal of Cryptology* (2010).
- [12] Alexandros Bampoulidis, Alessandro Bruni, Lukas Helming, Daniel Kales, Christian Rechberger, and Roman Walch. 2020. Privately Connecting Mobility to Infectious Diseases via Applied Cryptography. <https://ia.cr/2020/522>.
- [13] Paolo Barsocchi, Antonello Calabrò, Antonino Crivello, Said Daoudagh, Francesco Furfari, Michele Girolami, and Eda Marchetti. 2021. COVID-19 & privacy: Enhancing of indoor localization architectures towards effective social distancing. *Array* 9 (2021).
- [14] Sebastian P. Bayerl, Tommaso Frassetto, Patrick Jauernig, Korbinian Riedhammer, Ahmad-Reza Sadeghi, Thomas Schneider, Emmanuel Stäpf, and Christian Weinert. 2020. Offline Model Guard: Secure and Private ML on Mobile Devices. *Design, Automation & Test in Europe Conference & Exhibition* (2020).
- [15] Joshua Blumenstock, Gabriel Cadamuro, and Robert On. 2015. Predicting poverty and wealth from mobile phone metadata. *Science* (2015).
- [16] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. 2021. Lightweight Techniques for Private Heavy Hitters. In *S&P*.
- [17] Elette Boyle, Niv Gilboa, and Yuval Ishai. 2016. Function Secret Sharing: Improvements and Extensions. In *CCS*.
- [18] Beyza Bozdemir, Sébastien Canard, Orhan Ermis, Helen Möllering, Melek Önen, and Thomas Schneider. 2021. Privacy-preserving Density-based Clustering. In *ASIACCS*.
- [19] Fred Brauer. 2008. Compartmental models in epidemiology. In *Mathematical Epidemiology*.
- [20] Fred Brauer, Carlos Castillo-Chavez, and Zhilan Feng. 2019. Simple Compartmental Models for Disease Transmission. In *Mathematical Models in Epidemiology*.
- [21] Megha Byali, Harsh Chaudhari, Arpita Patra, and Ajith Suresh. 2020. FLASH: Fast and Robust Framework for Privacy-preserving Machine Learning. *PETS* (2020).
- [22] Clea Caulcutt. 2022. Belgium introduces quarantine for monkeypox cases. *Politico* (2022). <https://www.politico.eu/article/belgium-introduce-quarantine-monkeypox-case/>.
- [23] Justin Chan, Dean Foster, Shyam Gollakota, Eric Horvitz, Joseph Jaeger, Sham Kakade, Tadayoshi Kohno, John Langford, Jonathan Larson, Puneet Sharma, Sudheesh Singanamalla, Jacob Sunshine, and Stefano Tessaro. 2020. PACT: Privacy Sensitive Protocols and Mechanisms for Mobile Contact Tracing. <https://arxiv.org/pdf/2004.03544.pdf>.
- [24] Nishanth Chandran, Divya Gupta, Sai Lakshmi Bhavana Obbattu, and Akash Shah. 2022. SIMC: ML Inference Secure Against Malicious Clients at Semi-Honest Cost. In *USENIX Security*.
- [25] Harsh Chaudhari, Ashish Choudhury, Arpita Patra, and Ajith Suresh. 2019. ASTRA: High Throughput 3PC over Rings with Application to Secure Prediction. In *ACM CCSW@CCS*.
- [26] Harsh Chaudhari, Rahul Rachuri, and Ajith Suresh. 2020. Trident: Efficient 4PC Framework for Privacy Preserving Machine Learning. In *NDSS*.
- [27] David Chaum. 1985. Security without Identification: Transaction Systems to Make Big Brother Obsolete. In *Communications of the ACM*.
- [28] David Chaum. 1988. The Dining Cryptographers Problem: Unconditional Sender and Recipient Untraceability. *Journal of Cryptology* (1988).
- [29] David L Chaum. 1981. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Commun. ACM* (1981).
- [30] Guoxing Chen, Yinqian Zhang, and Ten-Hwang Lai. 2019. OPERA: Open Remote Attestation for Intel's Secure Enclaves. In *CCS*.
- [31] Yi-Cheng Chen, Ping-En Lu, Cheng-Shang Chang, and Tzu-Hsuan Liu. 2020. A Time-Dependent SIR Model for COVID-19 With Undetectable Infected Persons. *Transactions on Network Science and Engineering* (2020).
- [32] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. 1995. Private Information Retrieval. In *FOCS*.
- [33] Adam Durbini Christy Cooney. 2022. High-risk monkeypox contacts advised to isolate. *BBC* (2022). <https://www.bbc.com/news/uk-61546480>.
- [34] Matteo Ciucci and Frédéric Gouardères. 2020. National COVID-19 contact tracing apps. *EPRS: European Parliamentary Research Service* (2020).
- [35] Ian Cooper, Argha Mondal, and Chris G Antonopoulos. 2020. A SIR model assumption for the spread of COVID-19 in different communities. *Chaos, Solitons & Fractals* (2020).
- [36] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. 2015. Riposte: An Anonymous Messaging System Handling Millions of Users. In *S&P*.
- [37] Henry Corrigan-Gibbs and Dmitry Kogan. 2020. Private Information Retrieval with Sublinear Online Time. In *EUROCRYPT*.
- [38] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P. Smart. 2013. Practical Covertly Secure MPC for Dishonest Majority – Or: Breaking the SPDZ Limits. In *ESORICS*.
- [39] Nicholas G Davies, Adam J Kucharski, Rosalind M Eggo, Amy Gimma, W John Edmunds, Thibaut Jombart, Kathleen O'Reilly, Akira Endo, Joel Hellewell, Emily S Nightingale, et al. 2020. Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: a modelling study. *The Lancet Public Health* (2020).
- [40] Sara Y Del Valle, James M Hyman, Herbert W Hethcote, and Stephen G Eubank. 2007. Mixing patterns between age groups in social networks. *Social Networks* (2007).
- [41] Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *NDSS*.
- [42] Odo Diekmann, Hans Heesterbeek, and Tom Britton. 2012. *Mathematical Tools for Understanding Infectious Disease Dynamics*. Princeton University Press.
- [43] Changyu Dong, Liqun Chen, and Zikai Wen. 2013. When private set intersection meets big data: an efficient and scalable protocol. In *CCS*.
- [44] Wenliang Du. 2001. A study of several specific secure two party computation problems. *USA: Purdue University* (2001).
- [45] W John Edmunds, CJ O'callaghan, and DJ Nokes. 1997. Who mixes with whom? A method to determine the contact patterns of adults that may lead to the spread of airborne infections. *Proceedings of the Royal Society of London. Series B: Biological Sciences* (1997).
- [46] J. Ekberg, K. Kostiainen, and N. Asokan. 2014. The Untapped Potential of Trusted Execution Environments on Mobile Devices. In *S&P*.
- [47] Z. Erkin, J. R. Troncoso-pastoriza, R. L. Lagendijk, and F. Perez-Gonzalez. 2013. Privacy-Preserving Data Aggregation in Smart Metering Systems: An Overview. In *Signal Processing Magazine*.
- [48] Saba Eskandarian and Dan Boneh. 2022. Clarion: Anonymous Communication from Multiparty Shuffling Protocols. In *NDSS*.
- [49] Neil Ferguson. 2005. What would happen if a flu pandemic arose in Asia? *Nature* (2005).
- [50] Neil M Ferguson, Derek AT Cummings, Christophe Fraser, James C Cajka, Philip C Cooley, and Donald S Burke. 2006. Strategies for mitigating an influenza pandemic. *Nature* (2006).
- [51] Jesús Fernández-Villaverde and Charles I Jones. 2022. Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Citie. *Journal of Economic Dynamics and Control* (2022).
- [52] European Centre for Disease Prevention and Control. 2022. Epidemiological update: Monkeypox outbreak. (2022). <https://www.ecdc.europa.eu/en/news-events/epidemiological-update-monkeypox-outbreak>.
- [53] Craig Gentry and Shai Halevi. 2019. Compressible FHE with Applications to PIR. In *TCC*.
- [54] Niv Gilboa and Yuval Ishai. 2014. Distributed Point Functions and Their Applications. In *EUROCRYPT*.
- [55] Giulia Giordano, Franco Blanchini, Raffaele Bruno, Patrizio Colaneri, Alessandro Di Filippo, Angela Di Matteo, and Marta Colaneri. 2020. Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. *Nature Medicine* (2020).
- [56] Oded Goldreich. 2009. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press.
- [57] Alison Gray, David Greenhalgh, Liangjian Hu, Xuerong Mao, and Jiafeng Pan. 2011. A Stochastic Differential Equation SIS Epidemic Model. *SIAM J. Appl. Math.* (2011).
- [58] Daniel Günther, Maurice Heymann, Benny Pinkas, and Thomas Schneider. 2022. GPU-accelerated PIR with Client-Independent Preprocessing for Large-Scale Applications. In *USENIX Security*.

- [59] Thomas Haines and Johannes Müller. 2020. SoK: Techniques for Verifiable Mix Nets. In *CSF*.
- [60] Tiberiu Harko, Francisco SN Lobo, and MK3197716 Mak. 2014. Exact analytical solutions of the Susceptible-Infected-Recovered (SIR) epidemic model and of the SIR model with equal death and birth rates. *Appl. Math. Comput.* (2014).
- [61] Gary F Hatke, Monica Montanari, Swaroop Appadwedula, Michael Wentz, John Meklenburg, Louise Ivers, Jennifer Watson, and Paul Fiore. 2020. Using Bluetooth Low Energy (BLE) Signal Strength Estimation to Facilitate Contact Tracing for COVID-19. <https://arxiv.org/ftp/arxiv/papers/2006/2006.15711.pdf>.
- [62] Shaobo He, Yuexi Peng, and Kehui Sun. 2020. SEIR modeling of the COVID-19 and its dynamics. *Nonlinear Dynamics* (2020).
- [63] Katie Hogan, Briana Macedo, Venkata Macha, Arko Barman, Xiaoqian Jiang, et al. 2021. Contact Tracing Apps: Lessons Learned on Privacy, Autonomy, and the Need for Detailed and Thoughtful Implementation. *JMIR Medical Informatics* (2021).
- [64] Yan Huang, David Evans, and Jonathan Katz. 2012. Private Set Intersection: Are Garbled Circuits Better than Custom Protocols?. In *NDSS*.
- [65] Inria and Fraunhofer AISEC. 2020. *ROBust and privacy-presERving proximity Tracing protocol*. <https://github.com/ROBERT-proximity-tracing/documents>.
- [66] Intel. 2014. Intel® Software Guard Extensions Programming Reference. <https://software.intel.com/sites/default/files/managed/48/88/329298-002.pdf>.
- [67] Intel. unk. Attestation Service for Intel Software Guard Extensions. <https://api.trustedservices.intel.com/documents/sgx-attestation-api-spec.pdf>.
- [68] Kimmo Järvinen, Helena Leppäkoski, Elena-Simona Lohan, Philipp Richter, Thomas Schneider, Oleksandr Tkachenko, and Zheng Yang. 2019. PILOT: Practical Privacy-Preserving Indoor Localization Using Outsourcing. In *EUROS&P*.
- [69] P. Jaurnig, A. Sadeghi, and E. Staff. 2020. Trusted Execution Environments: Properties, Applications, and Challenges. In *S&P*.
- [70] Daniel Kales, Olamide Omolola, and Sebastian Ramacher. 2019. Revisiting User Privacy for Certificate Transparency. In *EuroS&P*. Code: <https://github.com/dkales/dpf-cpp>.
- [71] W. O. Kermack and A. G. McKendrick. 1991. Contributions to the mathematical theory of epidemics—I. In *Bulletin of Mathematical Biology*.
- [72] Adam Kirsch, Michael Mitzenmacher, and Udi Wieder. 2010. More Robust Hashing: Cuckoo Hashing with a Stash. *Journal on Computing* (2010).
- [73] Petra Klepac, Adam J Kucharski, Andrew JK Conlan, Stephen Kissler, Maria L Tang, Hannah Fry, and Julia R Gog. 2020. Contacts in context: large-scale setting-specific social mixing matrices from the BBC Pandemic project. *MedRxiv* (2020).
- [74] Nishat Koti, Mahak Pancholi, Arpita Patra, and Ajith Suresh. 2021. SWIFT: Superfast and Robust Privacy-Preserving Machine Learning. In *USENIX Security*.
- [75] Nishat Koti, Arpita Patra, Rahul Rachuri, and Ajith Suresh. 2022. Tetrad: Actively Secure 4PC for Secure Training and Inference. In *NDSS*.
- [76] Adam J Kucharski, Petra Klepac, Andrew JK Conlan, Stephen M Kissler, Maria L Tang, Hannah Fry, Julia R Gog, W John Edmunds, Jon C Emery, Graham Medley, et al. 2020. Effectiveness of isolation, testing, contact tracing, and physical distancing on reducing transmission of SARS-CoV-2 in different settings: a mathematical modelling study. *The Lancet Infectious Diseases* (2020).
- [77] Kai Kupferschmidt. 2020. Case clustering emerges as key pandemic puzzle. <https://www.science.org/doi/full/10.1126/science.368.6493.808>.
- [78] Klaus Kursawe, George Danezis, and Markulf Kohlweiss. 2011. Privacy-friendly aggregation for the smart-grid. In *PETS*.
- [79] Eyal Kushilevitz and Rafail Ostrovsky. 1997. Replication is NOT Needed: SINGLE Database, Computationally-Private Information Retrieval. In *FOCS*.
- [80] Sven Laur, Jan Willemsen, and Bingsheng Zhang. 2011. Round-Efficient Oblivious Database Manipulation. In *International Conference on Information Security*.
- [81] Ryan Lehmkuhl, Pratyush Mishra, Akshayaram Srinivasan, and Raluca Ada Popa. 2021. MUSE: Secure Inference Resilient to Malicious Clients. In *USENIX Security*.
- [82] Dyani Lewis. 2020. Where Covid contract-tracing went wrong. *Nature* (2020).
- [83] F. Li, B. Luo, and P. Liu. 2010. Secure Information Aggregation for Smart Grids Using Homomorphic Encryption. In *International Conference on Smart Grid Communications*.
- [84] Yehuda Lindell and Benny Pinkas. 2007. An Efficient Protocol for Secure Two-Party Computation in the Presence of Malicious Adversaries. In *EUROCRYPT*.
- [85] Yehuda Lindell, Benny Pinkas, Nigel P Smart, and Avishay Yanai. 2015. Efficient Constant Round Multi-Party Computation Combining BMR and SPDZ. In *CRYPTO*.
- [86] Wouter Lueks, Seda F. Gürses, Michael Veale, Edouard Bugnion, Marcel Salathé, Kenneth G. Paterson, and Carmela Troncoso. 2021. CrowdNotifier: Decentralized Privacy-Preserving Presence Tracing. *PoPETs*. (2021).
- [87] Shaojun Luo, Flaviano Morone, Carlos Sarraute, Matias Travizano, and Hernán A Makse. 2017. Inferring personal economic status from social network location. *Nature Communications* (2017).
- [88] Dominika Maison, Diana Jaworska, Dominika Adamczyk, and Daria Affeltowicz. 2021. The challenges arising from the COVID-19 pandemic and the way people deal with them. A qualitative longitudinal study. *PLoS One* (2021).
- [89] Robert M. May and Alun L. Lloyd. 2001. Infection dynamics on scale-free networks. *Phys. Rev. E* (2001).
- [90] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *International Conference on Artificial Intelligence and Statistics*.
- [91] Pratyush Mishra, Ryan Lehmkuhl, Akshayaram Srinivasan, Wenting Zheng, and Raluca Ada Popa. 2020. Delphi: A Cryptographic Inference Service for Neural Networks. In *USENIX Security*.
- [92] Payman Mohassel and Saeed Sadeghian. 2013. How to Hide Circuits in MPC an Efficient Framework for Private Function Evaluation. In *EUROCRYPT*.
- [93] Payman Mohassel and Yupeng Zhang. 2017. SecureML: A System for Scalable Privacy-Preserving Machine Learning. In *S&P*.
- [94] Yves-Alexandre de Montjoye, Jordi Quoidbach, Florent Robic, and Alex Sandy Pentland. 2013. Predicting Personality Using Novel Mobile Phone-Based Metrics. In *International conference on social computing, behavioral-cultural modeling, and prediction*.
- [95] Joël Mossong, Niel Hens, Mark Jit, Philippe Beutels, Kari Auranen, Rafael Mikolajczyk, Marco Massari, Stefania Salmaso, Gianpaolo Scalia Tomba, Jacco Wallinga, et al. 2008. Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases. *PLoS Medicine* (2008).
- [96] B. Ngabonziza, D. Martin, A. Bailey, H. Cho, and S. Martin. 2016. TrustZone Explained: Architectural Features and Use Cases. In *International Conference on Collaboration and Internet Computing*.
- [97] Olga Ohrimenko, Felix Schuster, Cédric Fournet, Aastha Mehta, Kapil Vaswani, and Manuel Costa. 2016. Oblivious Multi-Party Machine Learning on Trusted Processors. In *USENIX Security*.
- [98] Rasmus Pagh and Flemming Friche Rodler. 2004. Cuckoo Hashing. *Journal of Algorithms* (2004).
- [99] Christian Paquin and Greg Zaveruch. 2013. *U-Prove Cryptographic Specification V1.1 (Revision 3)*. <http://www.microsoft.com/uprove>.
- [100] Romualdo Pastor-Satorras and Alessandro Vespignani. 2002. Immunization of complex networks. *Phys. Rev. E* (2002).
- [101] Arpita Patra, Thomas Schneider, Ajith Suresh, and Hossein Yalame. 2021. ABY2.0: Improved Mixed-Protocol Secure Two-Party Computation. In *USENIX Security*.
- [102] Arpita Patra and Ajith Suresh. 2020. BLAZE: Blazing Fast Privacy-Preserving Machine Learning. In *NDSS*.
- [103] Matthias Pezzutto, Nicolás Bono Rosselló, Luca Schenato, and Emanuele Garone. 2021. Smart Testing and Selective Quarantine for the Control of Epidemics. *Annu. Rev. Control.* 51 (2021), 540–550.
- [104] Benny Pinkas and Eyal Ronen. 2021. Hashomer—Privacy-Preserving Bluetooth Based Contact Tracing Scheme for Hamagen. *Real World Crypto and NDSS Corona-Def Workshop* (2021).
- [105] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. 2020. PSI from PaXoS Fast, Malicious Private Set Intersection. In *EUROCRYPT*.
- [106] Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. 2015. Phasing: Private Set Intersection Using Permutation-based Hashing. In *USENIX Security*.
- [107] Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. 2018. Efficient Circuit-Based PSI via Cuckoo Hashing. In *EUROCRYPT*.
- [108] Benny Pinkas, Thomas Schneider, and Michael Zohner. 2018. Scalable Private Set Intersection based on OT Extension. *TOPS* (2018).
- [109] Leonie Reichert, Samuel Brack, and Björn Scheuermann. 2021. Poster: Privacy-Preserving Contact Tracing of COVID-19 Patients. *S&P* (2021).
- [110] Reinhard Schlickeiser and Martin Kröger. 2021. Analytical Modeling of the Temporal Evolution of Epidemics Outbreaks Accounting for Vaccinations. (2021).
- [111] Thomas Schneider and Oleksandr Tkachenko. 2019. EPISODE: Efficient Privacy-Preserving Similar Sequence Queries on Outsourced Genomic Databases. In *ASIACCS*.
- [112] Gitanjali R Shinde, Asmita B Kalamkar, Parikshit N Mahalle, Nilanjan Dey, Jyotismita Chaki, and Aboul Ella Hassanien. 2020. Forecasting Models for Coronavirus Disease (COVID-19): A Survey of the State-of-the-Art. *SN Computer Science* (2020).
- [113] Petrônio CL Silva, Paulo VC Batista, Hélder S Lima, Marcos A Alves, Frederico G Guimarães, and Rodrigo CP Silva. 2020. COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions. *Chaos, Solitons & Fractals* (2020).
- [114] Vivek K Singh, Laura Freeman, Bruno Lepri, and Alex Sandy Pentland. 2013. Predicting Spending Behavior Using Socio-mobile Features. In *International Conference on Social Computing*.
- [115] Michael Small and Chi K Tse. 2005. Small World and Scale free Model of Transmission of SARS. In *International Journal of Bifurcation and Chaos*.
- [116] Hallam Stevens and Monamie Bhadra Haines. 2020. TraceTogether: Pandemic Response, Democracy, and Technology. <https://www.tracetogogether.gov.sg>.
- [117] Amanda Taub. 2020. A New Covid-19 Crisis: Domestic Abuse Rises Worldwide. *The New York Times* (2020).
- [118] Robin N. Thompson. 2020. Epidemiological models are important tools for guiding COVID-19 interventions. *BMC Medicine* 18, 1 (2020), 152.
- [119] Oleksandr Tkachenko, Christian Weinert, Thomas Schneider, and Kay Hamacher. 2018. Large-Scale Privacy-Preserving Statistical Computations for Distributed Genome-Wide Association Studies. In *ASIACCS*.

- [120] Ni Trieu, Kareem Shehata, Prateek Saxena, Reza Shokri, and Dawn Song. 2020. Epione: Lightweight Contact Tracing with Strong Privacy. arXiv preprint arXiv:2004.13293. <https://arxiv.org/abs/2004.13293>.
- [121] Carmela Troncoso, Mathias Payer, Jean-Pierre Hubaux, Marcel Salathé, James R. Larus, Wouter Lueks, Theresa Stadler, Apostolos Pyrgelis, Daniele Antonioli, Ludovic Barman, Sylvain Chatel, Kenneth G. Paterson, Srdjan Capkun, David A. Basin, Jan Beutel, Dennis Jackson, Marc Roeschlin, Patrick Leu, Bart Preneel, Nigel P. Smart, Aysajan Abidin, Seda Gurses, Michael Veale, Cas Cremers, Michael Backes, Nils Ole Tippenhauer, Reuben Binns, Ciro Cattuto, Alain Barrat, Dario Fiore, Manuel Barbosa, Rui Oliveira, and José Pereira. 2020. Decentralized Privacy-Preserving Proximity Tracing. *IEEE Data Eng. Bull.* (2020).
- [122] Paul Tupper, Sarah P. Otto, and Caroline Colijn. 2021. Fundamental Limitations of Contact Tracing for COVID-19. *FACETS* (2021).
- [123] Christopher van der Beets, Raine Nieminen, and Thomas Schneider. 2022. FAPRIL: Towards Faster Privacy-Preserving Fingerprint-Based Localization. In *SECURITY*.
- [124] Serge Vaudenay. 2020. Centralized or Decentralized? The Contact Tracing Dilemma. *Cryptography ePrint Archive, Report 2020/531*. <https://ia.cr/2020/531>.
- [125] Meilof Veeningen, Supriyo Chatterjee, Anna Zsófia Horváth, Gerald Spindler, Eric Boersma, Peter van der SPEK, Onno Van Der Galiën, Job Gutteling, Wessel Kraaij, and Thijs Veugen. 2018. Enabling Analytics on Sensitive Medical Data with Secure Multi-Party Computation. In *Medical Informatics Europe*.
- [126] Nina Vindegaard and Michael Eriksen Benros. 2020. COVID-19 pandemic and mental health consequences: Systematic review of the current evidence. *Brain, Behavior, and Immunity* (2020).
- [127] Olivia J Walch, Amy Cochran, and Daniel B Forger. 2016. A global quantification of “normal” sleep schedules using smartphone data. *Science advances* (2016).
- [128] Guan Wang, Tongbo Luo, Michael T Goodrich, Wenliang Du, and Zutao Zhu. 2010. Bureaucratic protocols for secure two-party sorting, selection, and permuting. In *ASIACCS*.
- [129] Victoria Woollaston. 2015. *Sleeping habits of the world revealed: The US wakes up grumpy, China has the best quality shut-eye and South Africa gets up the earliest*. <https://www.dailymail.co.uk/sciencetech/article-3042230/Sleeping-habits-world-revealed-wakes-grumpy-China-best-quality-shut-eye-South-Africa-wakes-earliest.html>.
- [130] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets. In *FOCS*.
- [131] Tijana Šušteršič, Andjela Blagojević, Danijela Cvetković, Aleksandar Cvetković, Ivan Lorencin, Sandi Baressi Šegota, Dragan Milovanović, Dejan Baskić, Zlatan Car, and Nenad Filipović. 2021. Epidemiological Predictive Modeling of COVID-19 Infection: Development, Testing, and Implementation on the Population of the Benelux Union. *Frontiers in Public Health* 9 (2021).

A RELATED PRIMITIVES

In the following, we provide an overview about the (cryptographic) primitives and other techniques used in this work.

Anonymous Communication: To simulate the transmission of the modelled disease, RIPPLE requires anonymous messaging between participants. Mix-nets [29] and protocols based on the dining cryptographer (DC) problem [28] were the first approaches to anonymous messaging. A fundamental technique underlying mix-nets is the execution of an oblivious shuffling algorithm that provides unlinkability between the messages before and after the shuffle. In a mix-net, so-called mix servers jointly perform the oblivious shuffling so that no single mix server is able to reconstruct the permutation performed on the input data. Past research established a wide variety of oblivious shuffle protocols based on garbled circuits [44, 64, 128], homomorphic encryption [64], distributed point functions [1], switching networks [92], permutation matrices [80, §4.1], sorting algorithms [80, §4.2], and re-sharing [80, §4.3+4.4]. Recently, the works of [9] and [48] proposed efficient oblivious shuffling schemes using a small number of mix net servers.

Trusted Execution Environment (TEE): RIPPLE_{TEE} (§4.1) assumes the availability of TEEs on the mobile devices of participants. TEEs are hardware-assisted environments providing secure storage and execution for sensitive data and applications isolated from the normal execution environment. Data stored in a TEE is secure even

if the operating system is compromised, i.e., it offers confidentiality, integrity, and access control [46, 69]. Widely adopted TEEs are Intel SGX [66] and ARM TrustZone [10] (often used on mobile platforms [96]). Using TEEs for private computation has been extensively investigated, e.g., [14, 97]. A process called remote attestation allows external parties to verify that its private data sent via a secure channel is received and processed inside the TEE using the intended code [30, 67].

Private Information Retrieval (PIR): The first computational single-server PIR (cPIR) scheme was introduced by Kushilevitz and Ostrovsky [79]. Recent cPIR schemes [8, 53] use homomorphic encryption (HE). However, single-server PIR suffers from significant computation overhead since compute intensive HE operations have to be computed on each of the database block for each PIR request. In contrast, multi-server PIR relies on a non-collusion assumption between multiple PIR servers and uses only XOR operations [16, 17, 32, 36, 37] making it significantly more efficient than cPIR.

Cuckoo Hashing: In RIPPLE_{PIR} (§4.2), messages of participants have to be stored in a database D . To do so, a hash function H can be used to map an element x into bins of the database: $D[H(x)] = x$. However, as we show in §4.2, RIPPLE_{PIR} requires that at most one element is stored in every database location which renders simple hashing impracticable [106]. Cuckoo hashing uses h hash functions H_1, \dots, H_h to map elements into bins. It ensures that each bin contains exactly one element. If a collision occurs, i.e., if a new element is to be added into an already occupied bin, the old element is removed to make space for the new one. The evicted element, then, is placed into a new bin using another of the h hash functions. If the insertion fails for a certain amount of trials, an element is inserted into a special bin called stash which is the only one that is allowed to hold more than one element. Pinkas et al. [106] show that for $h = 2$ and $n = 2^{20}$ elements inserted to $2.4n$ bins, a stash size of 3 is sufficient to have a negligible error probability.

Garbled Cuckoo Table (GCT): As RIPPLE_{PIR} uses key-value pairs for the insertion into the database, a combination of garbled Bloom filters [43] with cuckoo hashing [72, 98], called Garbled Cuckoo Table [105], is needed. Instead of storing x in one bin as in an ordinary cuckoo table, in a GCT h XOR shares of x are stored at the h locations determined by inputting k into all h hash functions. E.g., with $h = 2$, if one of these two locations is already in use, the XOR share for the other (free) location is set to be the XOR of x and the data stored in the used location. In §B.4, we introduce a variant of GCT called *arithmetic garbled cuckoo table* (AGCT) that uses arithmetic sharing over the ring \mathbb{Z}_{2^t} instead of XOR sharing. For a database with $2.4n$ entries where n is the number of elements inserted, Pinkas et al. [105] show that the number of cycles is maximally $\log n$ with high probability.

Secure Multi-Party Computation (MPC): MPC [130] allows a set of mutually distrusting parties to jointly compute an arbitrary function on their private inputs without leaking anything but the output. In the last years, MPC techniques for various security models have been introduced, extensively studied, and improved, e.g., in [38, 41, 85]. These advancements significantly enhance the efficiency of MPC making it more and more practical for real-world applications. Due to the practical efficiency it can provide, various works [9, 21, 26, 74, 75, 102] have recently concentrated on MPC for a small number of parties, especially in the three and four party

honest majority setting tolerating one corruption. In RIPPLE, we employ MPC techniques across three servers to enable an anonymous communication channel (cf. §B.3) and to develop efficient PIR_{sum} (cf. §5) protocols.

Anonymous Credentials: To protect against sybil attacks (cf. §3.2), i.e., to hinder an adversary from creating multiple identities that can collect encounter information to detect correlations among unconscious encounters, we suggest to use anonymous credentials such that only registered participants can join RIPPLE. In this manner, the registration process can, for example, be linked to a passport. Such a registration system increases the cost to create (fake) identities. Chaum [27] introduced anonymous credentials where a client holds the credentials of several unlinkable pseudonyms. The client can then prove that it possesses the credentials of pseudonyms without the service provider being able to link different pseudonyms to the same identity. Additionally, anonymous credentials allow to certify specific properties like the age. Several instantiations for anonymous credentials have been proposed, e.g., Microsoft U-Prove [99].

B BUILDING BLOCKS IN RIPPLE

This section contains details about the building blocks used in the RIPPLE framework, such as shared-key setup, anonymous communication channels, AGCT tables and Distributed Point Functions.

B.1 Shared-Key Setup

Let $F : \{0, 1\}^{\kappa} \times \{0, 1\}^{\kappa} \rightarrow X$ be a secure pseudo-random function (PRF), with co-domain X being $\mathbb{Z}_{2^\ell}^r$ and $C' = C \cup \{\mathcal{P}_i\}$ for a participant $\mathcal{P}_i \in \mathcal{P}$. The following PRF keys are established among the parties in C' in RIPPLE:

- k_{ij} among every $P_i, P_j \in C'$ and $i \neq j$.
- k_{ijk} among every $P_i, P_j, P_k \in C'$ and $i \neq j \neq k$.
- $k_{C'}$ among all the parties in C' .

To sample a random value $r_{ij} \in \mathbb{Z}_{2^\ell}^r$ non-interactively, each of P_i and P_j can invoke $F_{k_{ij}}(id_{ij})$. In this case, id_{ij} is a counter that P_i and P_j maintain and update after each PRF invocation. The appropriate sampling keys are implied by the context and are, thus, omitted.

B.2 Collision Resistant Hash Function

A family of hash functions $\{H : \mathcal{K} \times \mathcal{L} \rightarrow \mathcal{Y}\}$ is said to be *collision resistant* if, for all probabilistic polynomial-time adversaries \mathcal{A} , given the description of H_k , where $k \in \mathcal{K}$, there exists a negligible function $negl(\cdot)$ such that $\Pr[(x, x') \leftarrow \mathcal{A}(k) : (x \neq x') \wedge H_k(x) = H_k(x')] = negl(\kappa)$, where $x, x' \in_R \{0, 1\}^m$ and $m = \text{poly}(\kappa)$.

B.3 Anonymous Communication Channel

This section describes how to instantiate the $\mathcal{F}_{\text{anon}}$ functionality used by RIPPLE for anonymous communication, as discussed in §4. We start with the protocol for the case of $\text{RIPPLE}_{\text{PIR}}$ and then show how to optimize it for the use in the $\text{RIPPLE}_{\text{TEE}}$ protocol. Recall from §4.2 that in $\text{RIPPLE}_{\text{PIR}}$, participants in \mathcal{P} upload a set of messages from which a database D must be constructed at the end by S_1 and S_2 . The anonymous communication is required to ensure that neither S_1 nor S_2 can link the source of the message even after receiving all messages in clear, which may not be in the

same order. To tackle this problem, we use an approach based on oblivious shuffling inspired by [9, 48], which is formalised next.

Problem Statement. Consider the vector $\vec{m} = \{m_1, \dots, m_\tau\}$ of τ messages with $m_j \in \mathbb{Z}_{2^\ell}^r$ for $j \in [\tau]$. We want servers S_1 and S_2 to obtain $\pi(\vec{m})$, where $\pi(\cdot)$ denotes a random permutation that neither S_1 nor S_2 knows. Furthermore, an attacker with access to a portion of the network and, hence, the ability to monitor network data should not be able to gain any information about the permutation $\pi(\cdot)$.

In $\text{RIPPLE}_{\text{PIR}}$, the vector \vec{m} corresponds to the infection likelihood messages of the form $(a_{i,j}, c_{i,j}^e)$ that each participant $\mathcal{P}_i \in \mathcal{P}$ conveys over the network (cf. §4.2). W.l.o.g., we let \mathcal{P}_i have the complete \vec{m} with them. The protocol makes use of the third server S_0 in our setting and proceeds as follows:

1. \mathcal{P}_i generates an additive sharing of \vec{m} among S_0 and S_1 :
 - a) \mathcal{P}_i, S_0 sample random $\langle \vec{m} \rangle_1 \in \mathbb{Z}_{2^\ell}^r$.
 - b) \mathcal{P}_i computes and sends $\langle \vec{m} \rangle_2 = \vec{m} - \langle \vec{m} \rangle_1$ to S_1 .
2. S_0 and S_1 agree on a random permutation π_{01} and locally apply π_{01} to their shares. Let $\pi_{01}(\vec{m}) = \pi_{01}(\langle \vec{m} \rangle_1) + \pi_{01}(\langle \vec{m} \rangle_2)$.
3. S_0, S_1 perform a *re-sharing* of $\pi_{01}(\vec{m})$, denoted by \vec{m}_{01} , by jointly sampling a random $r_{01} \in \mathbb{Z}_{2^\ell}^r$ and setting $\langle \vec{m}_{01} \rangle_1 = \pi_{01}(\langle \vec{m} \rangle_1) + r_{01}$ and $\langle \vec{m}_{01} \rangle_2 = \pi_{01}(\langle \vec{m} \rangle_2) - r_{01}$.
4. S_1 sends $\langle \vec{m}_{01} \rangle_2$ to S_2 . Now, $(\langle \vec{m}_{01} \rangle_1, \langle \vec{m}_{01} \rangle_2)$ forms an additive sharing of \vec{m}_{01} among S_0 and S_2 .
5. S_0 and S_2 agree on a random permutation π_{02} and apply π_{02} to their shares. Let $\pi_{02}(\vec{m}_{01}) = \pi_{02}(\langle \vec{m}_{01} \rangle_1) + \pi_{02}(\langle \vec{m}_{01} \rangle_2)$.
6. S_0 sends $\pi_{02}(\langle \vec{m}_{01} \rangle_1)$ to S_2 , who reconstructs $\pi_{02}(\vec{m}_{01})$.
7. S_2 generates an *additive-sharing* of $\pi_{02}(\vec{m}_{01})$, denoted by \vec{m}_{02} , among S_1 and S_2 , by jointly sampling $\langle \vec{m}_{02} \rangle_1 \in \mathbb{Z}_{2^\ell}^r$ with S_1 and locally setting $\langle \vec{m}_{02} \rangle_2 = \pi_{02}(\vec{m}_{01}) - \langle \vec{m}_{02} \rangle_1$.
8. S_2 sends $\langle \vec{m}_{02} \rangle_2$ to S_1 , who locally compute the output as $\vec{m}_{02} = \langle \vec{m}_{02} \rangle_1 + \langle \vec{m}_{02} \rangle_2$.

We claim that the protocol described above will produce a random permutation of the vector \vec{m} that neither S_1 nor S_2 is aware of. To see this, note that

$$\vec{m}_{02} = \pi_{02}(\vec{m}_{01}) = \pi_{02}(\pi_{01}(\vec{m}))$$

and both S_1 and S_2 know only one of the two permutations π_{01} and π_{02} , but not both. Furthermore, the re-sharing performed in step 3 and the generation of additive-shares in step 6 above ensures that an attacker observing the traffic cannot relate between messages sent and received.

Anonymous Communication in $\text{RIPPLE}_{\text{TEE}}$. As discussed in §4.1, the server S_2 is only required to have the complete set of messages in clear, but in an unknown random order. As a result, in the case of $\text{RIPPLE}_{\text{TEE}}$, only the first permutation (π_{01} in Step 2) is sufficient and steps 5-8 are no longer required. Furthermore, in addition to the communication by S_1 in step 4, S_0 sends its share of \vec{m}_{01} to S_2 , who can then reconstruct $\vec{m}_{01} = \pi_{01}(\vec{m})$.

B.4 Arithmetic Garbled Cuckoo Table (AGCT)

We design a variant of garbled cuckoo tables ([105], cf. §A) that we term arithmetic garbled cuckoo table (AGCT) to reduce the required size of a PIR database while ensuring a negligible collision likelihood. It uses arithmetic sharing instead of XOR-sharing to share database entries and the details are presented next.

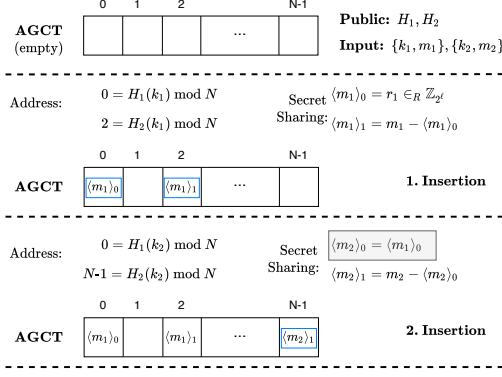


Figure 15: Insertion into the Arithmetic Garbled Cuckoo Table (AGCT). H_1 and H_2 are two hash functions. $\{k_1, m_1\}$ and $\{k_2, m_2\}$ are key-value pairs where the key is used to determine the address of the data in the database.

Let's assume two key-message pairs $\{k_1, m_1\}$ and $\{k_2, m_2\}$ ¹⁰ that shall be added to database D with N bins and two hash function H_1 and H_2 to determine the insertion addresses. The insertion process works as follows:

1. Insertion of $\{k_1, m_1\}$:
 - a) Compute $a_1 = H_1(k_1) \bmod N$ and $a_2 = H_2(k_1) \bmod N$.
 - b) Check if bins a_1 and a_2 are already occupied. Let's assume this is not the case.
 - c) Compute the arithmetic sharing of the message m_1 : $\langle m_1 \rangle_0 = r_1 \in_R \mathbb{Z}_{2^\ell}$ and $\langle m_1 \rangle_1 = m_1 - \langle m_1 \rangle_0$.
 - d) Insert $D[a_1] = \langle m_1 \rangle_0$ and $D[a_2] = \langle m_1 \rangle_1$.
2. Insertion of $\{k_2, m_2\}$:
 - a) Compute $b_1 = H_1(k_2) \bmod N$ and $b_2 = H_2(k_2) \bmod N$.
 - b) Check if bins b_1 and b_2 are already occupied. Let's assume $b_1 = a_1$, i.e., the first bin is already occupied, but bin b_2 is free.
 - c) Compute the arithmetic sharing m_2 with $\langle m_2 \rangle_0 = \langle m_1 \rangle_0$ as $b_1 = a_1$. Then, the other share is $\langle m_2 \rangle_1 = m_2 - \langle m_2 \rangle_0$.
 - d) Insert $D[b_1] = \langle m_2 \rangle_0$ and $D[b_2] = \langle m_2 \rangle_1$.

Double Collision: Now the question is how to handle the insertion of a database entry if both addresses determined by the two hash functions are already occupied. An easy solution is to pick different hash functions s.t. no double collision occurs for all n elements that shall be stored in the database. Alternatively, Pinkas et al. [105] demonstrate for a garbled cuckoo table how to extend the database by $d + \lambda$ bins, where d is the upper bound of double collisions and λ is an error parameter, such that double collisions occur with a negligible likelihood. For details, please refer to [105, §5].

¹⁰ k corresponds to a key and m to a message in our application.

B.5 Distributed Point Functions (DPF)

Consider a point function $P_{\alpha, \beta} : \mathbb{Z}_{2^\ell} \rightarrow \mathbb{Z}_{2^\ell}$ such that for all $\alpha \in \mathbb{Z}_{2^\ell}$ and $\beta \in \mathbb{Z}_{2^\ell}$, $P_{\alpha, \beta}(\alpha) = \beta$ and $P_{\alpha, \beta}(\alpha') = 0$ for all $\alpha' \neq \alpha$. That is, when evaluated at any input other than α , the point function $P_{\alpha, \beta}$ returns 0 and when evaluated at α it returns β .

An (s, t) -distributed point function (DPF) [36, 54] distributes a point function $P_{\alpha, \beta}$ among s servers in such a way that no coalition of at most t servers learns anything about α or β given their t shares of the function. We use $(2, 1)$ -DPFs in RIPPLE to optimize the communication of PIR-based protocols, as discussed in §5.2. Formally, a $(2, 1)$ -DPF comprises of the following two functionalities:

- $\text{Gen}(\alpha, \beta) \rightarrow (k_1, k_2)$. Output two DPF keys k_1 and k_2 , given $\alpha \in \mathbb{Z}_{2^\ell}$ and $\beta \in \mathbb{Z}_{2^\ell}$.
- $\text{Eval}(k, \alpha') \rightarrow \beta'$. Return $\beta' \in \mathbb{Z}_{2^\ell}$, given key k generated using Gen , and an index $\alpha' \in \mathbb{Z}_{2^\ell}$.

A $(2, 1)$ -DPF is said to be *correct* if for all $\alpha, x \in \mathbb{Z}_{2^\ell}$, $\beta \in \mathbb{Z}_{2^\ell}$, and $(k_1, k_2) \leftarrow \text{Gen}(\alpha, \beta)$, it holds that

$$\text{Eval}(k_1, x) + \text{Eval}(k_2, x) = (x = \alpha) \cdot \beta : 0.$$

A $(2, 1)$ -DPF is said to be *private* if neither of the keys k_1 and k_2 leaks any information about α and β . That is, there exists a polynomial time algorithm that can generate a computationally indistinguishable view of an adversary \mathcal{A} holding DPF key k_u for $u \in \{1, 2\}$, when given the key k_u .

As mentioned in [17, 36], a malicious participant could manipulate the Gen algorithm to generate incorrect DPF keys that do not correspond to any point function. While [36] used an external non-colluding auditor to circumvent this issue in the two server setting, [17] formalised this issue and proposed an enhanced version of DPF called Verifiable DPFs. In addition to the standard DPF, a verifiable DPF has an additional function called Ver that can be used to ensure the correctness of the DPF keys. In contrast to Eval , Ver in a $(2, 1)$ -verifiable DPF is an interactive protocol between the two servers, with the algorithm returning a single bit indicating whether the input DPF keys k_1 and k_2 are valid.

A verifiable DPF is said to be *correct* if for all $\alpha \in \mathbb{Z}_{2^\ell}$, $\beta \in \mathbb{Z}_{2^\ell}$, keys $(k_1, k_2) \leftarrow \text{Gen}(\alpha, \beta)$, the verify protocol Ver outputs 1 with probability 1. Ver should ensure that no additional information about α or β is disclosed to the party in possession of one of the DPF keys. Furthermore, the probability that Ver outputs 1 to at least one of the two servers for a given invalid key pair (k'_1, k'_2) is negligible in the security parameter κ .

Communication Complexity. Using the protocol of Boyle et al. [17], a $(2, 1)$ -DPF protocol for a point function with domain size N requires key size to be $(\lambda + 2) \cdot \log(N/\lambda) + 2 \cdot \lambda$ bits, where $\lambda = 128$ for an AES based implementation. The additional cost in the case of verifiable DPF is for executing the Ver function, which has constant number of elements in [17]. Furthermore, as stated in [17], the presence of additional non-colluding servers can improve the efficiency of Ver , and we use S_0 in the case of $\text{PIR}_{\text{sum}}^1$, as discussed in §5.2.1. We refer to [17] for more details regarding the scheme.

C PIR-SUM PROTOCOL DETAILS

This section provides additional details of the PIR_{sum} (cf. §5) protocols proposed in this work. We begin by recalling the security guarantees of a 2-server PIR with respect to our setting [32, 58]. Informally, a single server $S_u \in \{S_1, S_2\}$ should not learn any information about the client’s query in a two-server PIR protocol, where the database D is held by two non-colluding servers S_1 and S_2 . The security requirement is formally captured in Definition C.1.

Definition C.1. (Security of 2-server PIR) A PIR scheme with two non-colluding servers is called secure if each of the server does not learn any information about the query indices.

Let $\text{view}(S_u, Q)$ denote the view of server $S_u \in \{S_1, S_2\}$ with respect to a list of queries, denoted by Q . We require that for any database D , and for any two τ -length list of queries $Q = (q_1, \dots, q_\tau)$ and $Q' = (q'_1, \dots, q'_\tau)$, no algorithm whose run time is polynomial in τ and in computational parameter κ can distinguish the view of the servers S_1 and S_2 , between the case of participant \mathcal{P}_i using the queries in Q ($\{\text{view}(S_u, Q)\}_{u \in \{1,2\}}$), and the case of it using Q' ($\{\text{view}(S_u, Q')\}_{u \in \{1,2\}}$).

C.1 Linear Summation PIR for $\mathcal{F}_{\text{pir}}^{2S}$

This section describes the 2-server linear summation PIR protocol in [32], as well as how to optimize communication using DPF techniques discussed in §B.5. To retrieve the q -th block from database D of size N , the linear summation PIR proceeds as follows:

- Participant \mathcal{P}_i prepares an N -bit boolean string $\vec{b}_q = \{b_q^1, \dots, b_q^N\}$ with $b_q^j = 1$ for $j = q$ and $b_q^j = 0$ and $j \neq q$, for $j \in [N]$.
- \mathcal{P}_i generates a boolean sharing of \vec{b}_q among S_1 and S_2 , i.e., \mathcal{P}_i, S_1 non-interactively sample the random $[\vec{b}_q]_1 \in \{0, 1\}^N$ and \mathcal{P}_i sends $[\vec{b}_q]_2 = \vec{b}_q \oplus [\vec{b}_q]_1$ to S_2 .
- S_u , for $u \in \{1, 2\}$, sends $[y]_u = \bigoplus_{j=1}^N [b_q^j]_u D[j]$ to \mathcal{P}_i .
- \mathcal{P}_i locally computes $D[q] = [y]_1 \oplus [y]_2$.

The linear summation PIR described above requires communication of $N + 2\ell$ bits, where ℓ denotes the size of each data block in D .

C.1.1 Optimizing Communication using DPFs. Several works in the literature [17, 36, 54, 58] have used DPFs (cf. §B.5) as a primitive to improve the communication of a multi-server PIR scheme. The idea is to use a DPF function to allow the servers S_1 and S_2 to obtain the XOR shares of an N -bit boolean string \vec{b} that has a zero in all positions except the one representing the query q . Because DPF keys are much smaller in size than the actual database size, this method aids in the elimination of N -bit communication from \mathcal{P}_i to the servers, as discussed in the linear summation PIR in §C.1.

To query the q -th block from a database D of size N ,

- Participant \mathcal{P}_i executes the key generation algorithm with input q to obtain two DPF keys, i.e., $(k_1, k_2) \leftarrow \text{Gen}(q, 1)$.
- \mathcal{P}_i sends k_u to S_u , for $u \in \{1, 2\}$.
- S_u , for $u \in \{1, 2\}$, performs a DPF evaluation at each of the positions $j \in [N]$ using key k_u and obtains the XOR share corresponding to bit vector b_q .

- S_u expands the DPF keys as $[b_q^j]_u \leftarrow \text{Eval}(k_u, j)$ for $j \in [N]$.
- S_u , for $u \in \{1, 2\}$, sends $[y]_u = \bigoplus_{j=1}^N [b_q^j]_u D[j]$ to \mathcal{P}_i .
- \mathcal{P}_i locally computes $D[q] = [y]_1 \oplus [y]_2$.

For the case of semi-honest participants, we use the DPF protocol of [17] and the key size is $O(\lambda \cdot \log(N/\lambda))$ bits where $\lambda = 128$ is related to AES implementation in [17]. To prevent a malicious participant from sending incorrect or malformed keys to the servers, we use the verifiable DPF construction proposed in [17] for the case of malicious participants. As discussed in §B.5, this will only result in a constant communication overhead over the semi-honest case. Furthermore, as noted in [17], we use the additional server S_0 for a better instantiation of the verifiable DPF, removing the need for interaction with the participant \mathcal{P}_i for verification. We refer the reader to [17] for more details.

C.2 PIR_{sum} Protocol

Fig. 16 presents the ideal functionality for PIR_{sum} in the context of malicious participants. In this case, $\mathcal{F}_{\text{pirsum}}$ first checks whether all the queries made by the participant \mathcal{P}_i are distinct. If yes, the correct result is sent to \mathcal{P}_i ; otherwise, \perp is sent to \mathcal{P}_i .

Functionality $\mathcal{F}_{\text{pirsum}}$

$\mathcal{F}_{\text{pirsum}}$ interacts with servers in C , and participant $\mathcal{P}_i \in \mathcal{P}$.

Input: $\mathcal{F}_{\text{pirsum}}$ receives τ indices denoted by $Q = \{q_1, \dots, q_\tau\}$ from \mathcal{P}_i and a database D from C .

Computation: $\mathcal{F}_{\text{pirsum}}$ sets $y = \sum_{j=1}^\tau D[q_j]$ if all the queries in Q are distinct. Else, it sets $y = \perp$.

Output: $\mathcal{F}_{\text{pirsum}}$ sends y to \mathcal{P}_i .

Figure 16: PIR-SUM functionality (malicious participants).

We use a standard 2-server PIR functionality, denoted by $\mathcal{F}_{\text{pir}}^{2S}$, to instantiate $\mathcal{F}_{\text{pirsum}}$. The guarantees of $\mathcal{F}_{\text{pir}}^{2S}$, however, are insufficient to meet the security requirements of $\mathcal{F}_{\text{pirsum}}$, so we modify $\mathcal{F}_{\text{pir}}^{2S}$ as a sequence of hybrids, denoted by **HYB**. The modification is carried out in such a way that for a malicious participant \mathcal{P}_i , each hybrid is computationally indistinguishable from the one before it. As the first hybrid, $\mathcal{F}_{\text{pir}}^{2S}$ is denoted by **HYB₀**.

HYB₀: Let $\mathcal{F}_{\text{pir}}^{2S}$ denote a 2-server PIR ideal functionality for our case, with servers S_1 and S_2 acting as database holders and \mathcal{P}_i acting as the client. For a database D held by S_1 and S_2 and a query q held by \mathcal{P}_i , $\mathcal{F}_{\text{pir}}^{2S}$ returns $D[q]$ to \mathcal{P}_i , but S_1 and S_2 receive nothing.

HYB₁: We modify $\mathcal{F}_{\text{pir}}^{2S}$ so that it returns $D[q] + r$ to \mathcal{P}_i , and S_1, S_2 receive r , where r is a random value from the domain of database block size. In other words, the modification can be thought of as the standard $\mathcal{F}_{\text{pir}}^{2S}$ being executed over a database $D' = D + r$ rather than the actual database D . This modification leaks no additional information regarding the query to the servers because they will receive random masks that are independent of the query q . Furthermore, from the perspective of \mathcal{P}_i with no prior knowledge of the database D , **HYB₁** will be indistinguishable from **HYB₀** because the values it sees in both cases are from the same distribution. As a result, **HYB₀** \approx **HYB₁**.

HYB₂: Looking ahead, in PIR_{sum} , the servers S_1, S_2 and the participant \mathcal{P}_i run τ instances of $\mathcal{F}_{\text{pir}}^{2S}$ in parallel, one for each query $q \in Q$. As shown in $\mathcal{F}_{\text{pirsum}}$ (Fig. 16), the servers must ensure that all of the queries in Q are distinct. For this, we modify $\mathcal{F}_{\text{pir}}^{2S}$ in **HYB₁** to additionally output a secret share of the query q to each of S_1 and S_2 . Because the servers S_1 and S_2 are assumed to be non-colluding in our setting, this modification will leak no information about the query q to either server. Since the output to \mathcal{P}_i remains unchanged, **HYB₁** \approx **HYB₂** from \mathcal{P}_i 's perspective.

C.2.1 Instantiating $\mathcal{F}_{\text{pirsum}}$. The formal protocol for PIR_{sum} in the case of malicious participants is provided in Fig. 17, assuming the presence of an ideal functionality $\mathcal{F}_{\text{pir}}^{2S}$ as discussed in **HYB₂** above. In PIR_{sum} , the servers S_1, S_2 and the participant \mathcal{P}_i run τ instances of $\mathcal{F}_{\text{pir}}^{2S}$ in parallel, one for each query $q \in Q$. Following the execution, \mathcal{P}_i receives $D[q] + r_q$ whereas S_u receives $r_q, [q]_u$, for $u \in \{1, 2\}$ and $q \in Q$. \mathcal{P}_i then adds up the received messages to get a masked version of the desired output, i.e. $\sum_{q \in Q} D[q] + \text{mask}_Q$ with $\text{mask}_Q = \sum_{q \in Q} r_q$. S_1, S_2 compute mask_Q in the same way.

The protocol could be completed by S_1 and S_2 sending mask_Q to \mathcal{P}_i , then \mathcal{P}_i unmasking its value to obtain the desired output. However, before communicating the mask, the servers must ensure that all queries in Q are distinct, as shown in $\mathcal{F}_{\text{pirsum}}$ (Fig. 16). For this, S_1, S_2 use their share of the queries $q \in Q$ and participate in a secure computation protocol with S_0 . We capture this with an ideal functionality $\mathcal{F}_{\text{verfy}}$, which takes the secret shares of τ values from S_1 and S_2 and returns Accept to the servers if all of the underlying secrets are distinct. Otherwise, it returns abort.

Protocol PIR_{sum}

Input(s): i) $S_1, S_2 : D; |D| = N$, ii) $\mathcal{P}_i : Q = \{q_1, \dots, q_\tau\}$, and iii) $S_0 : \perp$.
Output: $\mathcal{P}_i : \text{res} = \sum_{q \in Q} D[q]$ for distinct queries, else $\text{res} = \perp$.

Computation

1. For each $q \in Q$,
 - a. S_1, S_2 and \mathcal{P}_i invoke $\mathcal{F}_{\text{pir}}^{2S}$ (cf. **HYB₂** in §C.2) with the inputs D, q .
 - b. Let $r_q, [q]_u$ denote the output of S_u , for $u \in \{1, 2\}$ and $D[q] + r_q$ denote the output of \mathcal{P}_i .
2. \mathcal{P}_i computes $\text{res}' = \sum_{q \in Q} (D[q] + r_q)$, while S_1, S_2 computes $\text{mask}_Q = \sum_{q \in Q} r_q$.
3. S_1, S_2 and S_0 invokes $\mathcal{F}_{\text{verfy}}$ on the secret shares of queries, denoted by $\{[q]_u\}_{q \in Q, u \in \{1, 2\}}$, to check the distinctness of the queries in Q .
4. If $\mathcal{F}_{\text{verfy}}$ returns Accept, S_1, S_2 sends mask_Q to \mathcal{P}_i , who computes $\text{res} = \text{res}' - \text{mask}_Q$. Otherwise, abort.

Figure 17: PIR_{sum} Protocol.

LEMMA C.2. Protocol PIR_{sum} (Fig. 17) securely realises the $\mathcal{F}_{\text{pirsum}}$ ideal functionality (Fig. 16) for the case of malicious participants in the $\{\mathcal{F}_{\text{pir}}^{2S}, \mathcal{F}_{\text{verfy}}\}$ -hybrid model.

PROOF. The proof follows with a hybrid argument based on the three hybrids **HYB₀**, **HYB₁**, and **HYB₂** discussed in §C.2. Furthermore, any secure three-party protocol can be used to instantiate $\mathcal{F}_{\text{verfy}}$ in RIPPLE. \square

C.2.2 $\text{PIR}_{\text{sum}}^I$ Protocol Details. Fig. 18 provides the formal details of $\text{PIR}_{\text{sum}}^I$ protocol discussed in §5.2.1.

Protocol $\text{PIR}_{\text{sum}}^I$

Input(s): i) $S_1, S_2 : D; |D| = N$, ii) $\mathcal{P}_i : Q = \{q_1, \dots, q_\tau\}$, and iii) $S_0 : \perp$.
Output: $\mathcal{P}_i : \text{res} = \sum_{q \in Q} D[q]$

Computation S_1 and S_2 sample τ random mask values $\{m_1, \dots, m_\tau\} \in \mathbb{Z}_{2^t}^\tau$ such that $\sum_{j=1}^\tau m_j = 0$. For each $q \in Q$, execute the following:

1. S_1, S_2 locally compute $D^{m_q} = D + m_q$.
2. Execute DPF protocol [17] (verifiable DPF for malicious participants) with \mathcal{P}_i as client with input q . Server S_u obtains $[\vec{b}_q]_u$ with $b_q^j = 1$ for $j = q$ and $b_q^j = 0$ for $j \neq q$, for $u \in \{1, 2\}$.

Verification Let $\{\vec{b}_{q_1}, \dots, \vec{b}_{q_\tau}\}$ denote the bit vectors whose XOR-shares are generated during the preceding steps.

3. Servers verify correctness of $q_j, j \in [\tau]$, by executing the Ver algorithm of the verifiable DPF protocol [17] (cf. §B.5). It outputs Accept to S_1 and S_2 if q_j has exactly 1 one and $(N - 1)$ zeroes. Else, it outputs abort.
4. S_u computes $[\vec{b}_c]_u = \oplus_{q \in Q} [\vec{b}_q]_u$, for $u \in \{1, 2\}$.
5. S_1 and S_2 non-interactively agree on random permutation π .
6. S_u sends $\pi([\vec{b}_c]_u)$ to S_0 , for $u \in \{1, 2\}$.
7. S_0 locally reconstructs $\pi(\vec{b}_c) = \pi([\vec{b}_q]_1) \oplus \pi([\vec{b}_q]_2)$. It sends Accept to S_1 and S_2 , if $\pi(\vec{b}_c)$ has exactly τ ones. Else, it sends abort.

Output Transfer Send \perp to \mathcal{P}_i if verifiable DPF or S_0 generated abort during verification. Otherwise, proceed as follows:

8. S_u sends $[y_q]_u = \bigoplus_{j=1}^N [b_q^j]_u D^{m_q}[j]$ to \mathcal{P}_i , for $q \in Q, u \in \{1, 2\}$.
9. \mathcal{P}_i locally computes $\text{res} = \sum_{q \in Q} ([y_q]_1 \oplus [y_q]_2)$.

Figure 18: $\text{PIR}_{\text{sum}}^I$ Protocol.

C.2.3 $\text{PIR}_{\text{sum}}^{II}$ Protocol Details. This section contains additional details about the $\text{PIR}_{\text{sum}}^{II}$ protocol, which was discussed in §5.2.2. A large amount of communication is used in the $\text{PIR}_{\text{sum}}^{II}$ protocol to verify malicious participants. More specifically, in Step 8 of Fig. 10, $2N$ bits are communicated towards S_0 to ensure the distinctness of the queries made by the participant \mathcal{P}_i . We note that allowing a small amount of leakage to S_0 could improve this communication and is discussed next.

Improving Verification Costs in $\text{PIR}_{\text{sum}}^{II}$. Consider the following modification to the $\text{PIR}_{\text{sum}}^{II}$ protocol. Instead of sampling θ_q for each query $q \in Q$ (cf. Step 2 in Fig. 10), \mathcal{P}_i, S_1 , and S_2 sample only one random shift value θ and uses it for all τ instances. Since the queries must be distinct, \mathcal{P}_i is forced to send distinct q' values to S_0 in Step 3 of Fig. 10. If not, S_0 can send abort to S_1 and S_2 at this step, eliminating the need for communication intensive verification. The relative distance between the queried indices would be leaked to S_0 as a result of this optimization. In concrete terms, if we use the same θ value for any two queries $q_m, q_j \in Q$, then $q_m - q_n = q'_m - q'_n$. Because S_0 sees all q' values in clear, it can deduce the relative positioning of \mathcal{P}_i 's actual queries. However, since S_0 has no information about the underlying database D , this leakage may be acceptable for some applications.

D ADDITIONAL BENCHMARKING

This section provides additional benchmarking details of RIPPLE framework, specifically $\text{RIPPLE}_{\text{TEE}}$ in §4.1 and $\text{RIPPLE}_{\text{PIR}}$ in §4.2.

Tab. 6 and Tab. 7 detail the computation costs per simulation and communication costs per simulation step at various stages in our instantiations of RIPPLE’s. As shown in Tab. 6, data transfer time as part of anonymous communication through servers accounts for the majority of computation time and begins to affect overall performance as population grows. Our system crashed due to memory constraints after a population of 500K while running the experiments. This will not be the case in a real-world deployment

of powerful servers linked by high-speed bandwidth. We find that a participant’s computation and communication costs are very low in comparison to the overall costs. In $\text{RIPPLE}_{\text{TEE}}$, a participant communicates at most 268 KB and incurs a runtime of 92 seconds over a two-week simulation over a population of one million. In $\text{PIR}_{\text{sum}}^{\text{II}}$, the cost is reduced to 100 KB and 40 seconds of runtime. While the runtime remains same as that of $\text{PIR}_{\text{sum}}^{\text{II}}$, communication increases to 1.2 MB in $\text{PIR}_{\text{sum}}^{\text{I}}$ due to the participant’s handling of DPF keys. Finally, Tab. 7 does not include costs for verification against malicious participants since they can be eliminated using server S_0 (cf. §C.2.3) or sketching algorithms similar to those in [17].

Stages of RIPPLE	Protocol	Population (p)					
		1K	10K	50K	100K	500K	1M
Message Generation by $\mathcal{P}_i \in \mathcal{P}$ (in sec)	$\text{RIPPLE}_{\text{TEE}}$ (§4.1)	1.12	1.12	1.12	1.12	1.12	1.12
	$\text{RIPPLE}_{\text{PIR}}$: (§4.2)	4.26e-3	4.26e-3	4.26e-3	4.26e-3	4.26e-3	4.26e-3
Secure Shuffle by C (in sec)	$\text{RIPPLE}_{\text{TEE}}$ (§4.1)	0.70	5.20	25.38	60.77	211.47	493.33 ^{*a}
	$\text{RIPPLE}_{\text{PIR}}$ (§4.2)	0.78	6.65	32.36	71.17	386.68	1542.30 [*]
Output Computation ^b (in sec)	$\text{RIPPLE}_{\text{TEE}}$ (§4.1)	44.66	44.66	44.66	44.66	44.66	44.66
	$\text{PIR}_{\text{sum}}^{\text{I}}$ (§5.2.1)	32.31	32.33	32.34	32.35	32.36	32.37
	$\text{PIR}_{\text{sum}}^{\text{II}}$ (§5.2.2)	32.20	32.20	32.20	32.20	32.20	32.20

^{a*} denotes system crash due to memory. ^b includes message download, decryption/PIR queries, summation.

Table 6: Detailed computation costs per simulation ($N_{\text{step}} = 14$, i.e., 14 days) in RIPPLE.

Stages of RIPPLE	Protocol	Population (p)									
		1K	10K	50K	100K	500K	1M	2M	5M	10M	20M
Message Generation by $\mathcal{P}_i \in \mathcal{P}$ (in KB)	$\text{RIPPLE}_{\text{TEE}}$ (§4.1) ^a	12.80	12.80	12.80	12.80	12.80	12.80	12.80	12.80	12.80	12.80
	$\text{RIPPLE}_{\text{PIR}}$: ● (§4.2)	3.20	3.20	3.20	3.20	3.20	3.20	3.20	3.20	3.20	3.20
	$\text{RIPPLE}_{\text{PIR}}$: ● (§4.2)	2.30	2.34	2.38	2.39	2.41	2.43	2.44	2.45	2.46	2.48
Secure Shuffle by C (in GB)	$\text{RIPPLE}_{\text{TEE}}$ (§4.1)	0.02	0.19	0.96	1.92	9.60	19.20	38.40	96.00	192.00	384.00
	$\text{RIPPLE}_{\text{PIR}}$ - ● (§4.2)	0.01	0.10	0.48	0.96	4.80	9.60	19.20	48.00	96.00	192.00
	$\text{RIPPLE}_{\text{PIR}}$ - ● (§4.2)	0.01	0.07	0.36	0.72	3.62	7.28	14.63	36.75	73.88	148.50
Output Computation by $\mathcal{P}_i \in \mathcal{P}$ (in KB) ^b	$\text{RIPPLE}_{\text{TEE}}$ (§4.1)	6.40	6.40	6.40	6.40	6.40	6.40	6.40	6.40	6.40	6.40
	$\text{PIR}_{\text{sum}}^{\text{I}}$ - ● (§5.2.1)	51.36	62.42	69.97	73.22	80.77	84.02	87.27	91.56	94.81	98.06
	$\text{PIR}_{\text{sum}}^{\text{I}}$ - ● (§5.2.1)	26.48	32.64	37.69	39.82	44.77	47.05	49.38	52.33	54.77	57.26
	$\text{PIR}_{\text{sum}}^{\text{II}}$ (§5.2.2)	3.45	3.49	3.52	3.53	3.56	3.57	3.59	3.60	3.62	3.63

● - 128-bit address for $\text{RIPPLE}_{\text{PIR}}$ and ● - $40 - 1 + \log_2(p \cdot E^{\text{avg}})$ bit address for $\text{RIPPLE}_{\text{PIR}}$.

^aIncludes registration of public keys with the exit node N_{exit} . ^b includes message download, decryption/PIR queries, summation.

Table 7: Detailed communication costs per simulation step in RIPPLE.