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

Table 1: Communication costs per simulation step in our RIPPLE instantiations.

(5) We demonstrate the practicality of RIPPLE by a detailed performance evaluation. Our findings indicate that our protocol 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.

Fig. 2 summarises the phases of the RIPPLE framework in the context of a single simulation setting. Note that multiple simulations can be executed in parallel.

3 EVALUATION

We evaluate the computation and communication efficiency of the simulation phase of our two RIPPLE protocols. The simulations can ideally be done overnight while mobile phones are charging and have access to a high-bandwidth WiFi connection. We run the benchmarks on the server-side with three servers 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. We instantiate our protocols in RIPPLE with $\kappa = 128$ bit security. A typical simulation step would be one day such that 14 steps simulate two weeks.

Communication Costs. We consider participant sizes, denoted by p, ranging from thousand (1K) to twenty million (20M) to analyse the scalability of our protocols. Tab. 1 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.

	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
$\text{PIR}^{\text{I}}_{\text{sum}}$	0.30	11.73	4.8e-2	4.26e-3	0.16	6.72e-4
$\text{PIR}^{\text{II}}_{\text{sum}}$	0.30	3.0e-3	4.8e-2	4.26e-3	4.2e-5	6.72e-4

Table 2: Average participant computation times per simulation step distributed across various tasks. Values are obtained using a mobile for a population of 500K.

Computation Costs. Tab. 2 summarizes the computation time with respect to a participant $\mathcal{P}_i$ for a two-week simulation over a half-million population. For a 14-day simulation with a population of half a 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 $\text{PIR}^{\text{II}}_{\text{sum}}$, while it increases to around 165 milliseconds for the case of $\text{PIR}^{\text{I}}_{\text{sum}}$.

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. RIPPLE is provably privacy-preserving by construction. We refer the reader to the full version [4] for more elaborate details of our work.

Code availability. Available at `DOI: 10.5281/zenodo.6595449`.

Acknowledgements. This project received funding from the ERC under the European Union's Horizon 2020 research and innovation program (grant agreement No. 850990 PSOTI). It was co-funded by the DFG within SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251805230, and by the BMBF and the HMWK within ATHENE.

REFERENCES

[1] David Adam. 2020. Special report: The simulations driving the world's response to COVID-19. *Nature* (2020).
[2] 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).
[3] 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).
[4] Daniel Günther, Marco Holz, Benjamin Judkewitz, Helen Möllering, Benny Pinkas, Thomas Schneider, and Ajith Suresh. 2022. Privacy-Preserving Epidemiological Modeling on Mobile Graphs. (2022). https://doi.org/10.48550/arXiv.2206.00539
[5] 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).
[6] Dyani Lewis. 2020. Where Covid contract-tracing went wrong. *Nature* (2020).
[7] 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).
[8] 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*.
[9] Robin N. Thompson. 2020. Epidemiological models are important tools for guiding COVID-19 interventions. *BMC Medicine* 18, 1 (2020), 152.
[10] Paul Tupper, Sarah P. Otto, and Caroline Colijn. 2021. Fundamental Limitations of Contact Tracing for COVID-19. *FACETS* (2021).
[11] 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).