Title
Prio+: Private Aggregate Statistics via Boolean Shares

Author
Jaffe, Eli Aaron

Publication Date
2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Prio+: Privacy Preserving Aggregate Statistics via Boolean Shares

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Computer Science

by

Eli Aaron Jaffe

2023

ABSTRACT OF THE DISSERTATION

Prio+: Privacy Preserving Aggregate Statistics via Boolean Shares

by

Eli Aaron Jaffe
Doctor of Philosophy in Computer Science
University of California, Los Angeles, 2023
Professor Rafail Ostrovsky, Chair

This paper introduces Prio+, a privacy-preserving system for the collection of aggregate statistics, with the same model and goals in mind as the original and highly influential Prio paper by Henry Corrigan-Gibbs and Dan Boneh (NSDI 2017). As in the original Prio, each client holds a private data value (e.g. number of visits to a particular website) and a small set of servers privately compute statistical functions over the set of client values (e.g. the average number of visits). To achieve security against faulty or malicious clients, unlike Prio, Prio+ clients use Boolean secret-sharing instead of zero-knowledge proofs to convince servers that their data is of the correct form and Prio+ servers execute a share conversion protocol as needed in order to properly compute over client data. This allows us to ensure that clients' data is properly formatted essentially for free, and the work shifts to novel share-conversion protocols between servers, where some care is needed to make it efficient. Our overall approach is simpler than Prio and our Prio+ strategy reduces the client's computational burden by at least two orders of magnitude (or more depending on the statistic) while keeping server costs comparable to Prio. Prio+ permits computation of exactly the same wide range of complex statistics as the original Prio protocol, including high-dimensional linear regression over private values held by clients. We report detailed benchmarks of our

Prio+ implementation and compare these to both the original Go implementation of Prio and the Mozilla implementation of Prio. Our Prio+ software is open-source and released with the same license as Prio.

The dissertation of Eli Aaron Jaffe is approved.

Alexander Sherstov

George Varghese

Leonard Kleinrock

Rafail Ostrovsky, Committee Chair

University of California, Los Angeles

2023

TABLE OF CONTENTS

1 Introduction **1**

2 Technical Overview **7**

3 Preliminaries **13**

4 Necessary Primitives **17**

5 Private Summation Without Robustness **21**

6 Protecting Correctness **24**

7 Complex Statistics **29**

8 Resolving Frequency Count **36**

9 Share Conversion **39**

10 Security **43**

11 Practical Evaluation **44**

11.1 Data: SUM 45

11.2 Data: MAX 51

11.3 Data: linReg 56

11.4 Data: Offline Pre-computation 61

12 Conclusions and Future Work **62**

A Full Protocol Descriptions **63**

B Security Definitions **74**

C Proofs of Security **77**

C.1 Privacy 77

C.2 Robustness 86

References **90**

LIST OF FIGURES

11.1 This chart shows the time necessary for a single client P_i holding private value x_i to encode that value, compute any necessary additional proofs, and secret-share both the encoding and the proof(s) when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in milliseconds and data is arranged according to the number of bits in x_i. 46

11.2 This chart shows the size of the message P_i (holding private value x_i) sends to each server when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in bytes and data is arranged according to the number of bits in x_i. 47

11.3 This chart shows the computation time for each server when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in microseconds and data is arranged according to the number of bits in the client's private value x_i. Prio+ (end-to-end) includes time to pre-compute daBits for share conversion. 48

11.4 This chart shows the communication required by each server when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in bytes and data is arranged according to the number of bits in x_i. Prio+ (end-to-end) includes time to pre-compute daBits for share conversion. 49

11.5 This chart shows the time necessary for a single client P_i holding private value x_i in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$ to encode that value, compute any necessary additional proofs, and secret-share both the encoding and the proof(s) when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$. Results are in milliseconds. 52

11.6 This chart shows the size of the message P_i (holding private value x_i in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$) sends to each server when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$. Results are in bytes. 53

11.7 This chart shows the average computation time per server when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$, where each x_i held by P_i lies in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$. Results are in milliseconds. 54

11.8 This chart shows the average bytes communicated by each server when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$, where each x_i held by P_i is in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$. Results are in bytes. 55

11.9 This chart shows the time necessary for a single client P_i holding private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$ to encode that value, compute any necessary additional proofs, and secret-share both the encoding and the proof(s) when executing a protocol to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$. Each $x_i^{(k)}$ is an 8-bit integer. Results are in milliseconds and data is arranged according to the degree d of each $\vec{x}_i$. 57

11.10 This chart shows the size of the message sent by client P_i (holding private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$) to each server when executing a protocol to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$. Each $x_i^{(k)}$ is an 8-bit integer. Results are in milliseconds and data is arranged according to the degree d of each $\vec{x}_i$. 58

11.11 This chart shows the total server time to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$, where each of the 50,000 clients P_i holds private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$ and each $x_i^{(k)}$ is an 8-bit integer. Results are in seconds and data is arranged according to the degree d of each $\vec{x}_i$. Prio+ (end-to-end) includes time to pre-compute daBits for share conversion. 59

11.12 This chart shows the average bytes communicated by each server when executing a protocol to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$, where client P_i holds private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$ and each $x_i^{(k)}$ is an 8-bit integer. Results are in bytes and data is arranged according to the degree d of each $\vec{x}_i$. Prio+ (end-to-end) includes time to pre-compute daBits for share conversion. 60

LIST OF TABLES

1.1 Table of asymptotic comparisons between Prio and the protocols Π comprising Prio+. M is the number of multiplication gates required to check if an input is properly encoded for the relevant statistic. λ is the bit-length of each client's input, $M' := M - \lambda$ is the number of multiplication gates not used for checking bit-length, and n is the total number of clients. Note $M \geq \lambda$ because λ multiplication gates are always used for checking bit-length in Prio. 'Serv Comm.' is the communication between servers in bits. Note that Prio and the Linear PCP extension implement a general protocol for all statistics, whereas Prio+ uses different protocols for each statistic. Thus the entries in the first two rows represent all statistics computed by Prio, where the value M varies depending on the particular statistic. Note that $\Pi_{\mathsf{sum}}, \Pi_{\mathsf{var/linReg}}, \Pi_{\mathsf{frq}}$ also use λ symmetric key operations (oblivious transfers) during pre-computation. $\Theta(\cdot)$ notation suppressed to improve readability. 5

ACKNOWLEDGMENTS

This work was made possible by my loving family and friends, especially my mom. I wouldn't be here without you. To those who loved me at my lowest (you know who you are), you are more a part of this achievement than I can express, thank you endlessly. Extra appreciation to Davina, Paul, and Alexis, I was a better student and person for having you all with me these past 5 years (an extra thank you to Alexis for helping me solve a tricky problem in this paper as well). To my advisors Rafi and Len, I am incredibly grateful for your support and guidance, thank you for everything.

VITA

2015-2018 Instructor, Olga Radko Math Circle, UCLA. Taught abstract math at various K-12 levels. Collaborated with UCLA faculty to produce curriculum.

2016 Grader, Mathematics Department, UCLA. Graded assignments for Math 61 (discrete math) under direction of Professor Assaf Shani.

2018 B.S. (Mathematics, Specialization in Computing), UCLA.

2018-2023 Research Assistant, Computer Science Department, UCLA.

2019-2020 Student Intern, Stealth Software Technologies Inc. Surveyed existing private set intersection protocols and analyzed real-world performance, implemented cryptographic primitives as part of large-scale multi-party computation platform.

2020 Teaching Assistant, Computer Science Department, UCLA. Taught sections of CS 183 (introduction to cryptography) under direction of Professor Rafail Ostrovsky.

2021 Research Intern, Microsoft. Improved efficiency of ElectionGuard, a Microsoft tool for election security, under the supervision of Josh Benaloh. Optimized zero-knowledge proof systems within the context of in-person voting.

2022 Research Intern, Microsoft. Worked under the supervision of Betul Durak to develop a novel zero-knowledge scheme for proving ownership of anonymous tokens as part of a study into the notion of non-transferable anonymous credentials.

2022 Teaching Assistant, Computer Science Department, UCLA.

PUBLICATIONS AND PRESENTATIONS

Prio+: Private Aggregate Statistics via Boolean Shares. Security and Cryptography for Networks, 2022.

Prio+: Private Aggregate Statistics via Boolean Shares. CRYPTO 2022 PPML Workshop. Presented our work and fielded questions.

Introduction to Cryptography, Zero-knowledge, and Blockchain. Blockchain Acceleration Foundation. Prepared and delivered introductory lectures to corporate leaders looking to adopt blockchain technology.

CHAPTER 1

Introduction

In recent decades, modern society has exploded with a wave of internet-enabled devices. Smart-watches, cell-phones, cars, and ATMs are constantly collecting data on their surroundings to improve performance. For many cloud services controlling such devices, collecting and computing statistics over such a large pool of data has become a hugely profitable endeavor. Navigation apps detect congestion with user location data [Jes13], fitness trackers collect average data for user comparison [HPK16]. Aggregate statistics are one of the principal currencies in the modern data-driven economy.

Although these services only wish to compute aggregate statistics, not collect individual data, their methods often involve storing users' personal data in the clear on their servers and then computing statistics directly on that data. Such a centralized cache of sensitive user data presents clear security risks. As described in [CB17], a motivated attacker may simply steal and disclose this sensitive information [KLP15, WWW16], cloud services could misuse or sell this information for profit [Smi14], and intelligence agencies may acquire the data for targeting or mass surveillance purposes [GLL14].

The specific problem to be solved is as follows: each of n clients P_i holds a private value $x_i \in \{0, 1\}$, and they wish to learn the sum $\sum_i x_i$. As described in the original Prio paper [CB17], some previous systems have also attempted to solve this problem of privately computing aggregate statistics. One such attempt involves using a randomized response system to provide differential privacy [EKP14, FPE15]. That is, some user data is replaced with

random data according to some fixed probability $p < 0.5$. By aggregating this "noisy" data, data collectors can get a somewhat accurate estimate of overall statistics. This technique scales well and provides robustness (each malicious client can at most affect the sum by ± 1), but the privacy guarantees are relatively weak. There is an inherent trade-off between the privacy guarantee to the client and the accuracy of the overall statistic. Another option is to have clients submit encryptions of their data to servers. Then, servers can sum up the ciphertexts and only decrypt the final sum [DFK13, EDG14, JL13, MDC15, enc09, PBB11]. This achieves stronger privacy guarantees but sacrifices robustness: a malicious client can affect the final sum arbitrarily because servers cannot tell the difference between an encryption of 0/1 and an encryption of some large integer. Such attacks are often incentivized: if used for a voting scheme, this would allow any client to submit as many votes as they like. These attacks can be mitigated using zero-knowledge proofs [EKO19], but such approaches heavily impact scalability. Servers require expensive public-key operations to verify these proofs, and clients are burdened with the computationally difficult task of generating the proofs.

Prio is a brilliant and highly influential private aggregation system which successfully resolves this discrepancy between privacy, robustness, and scalability. Prio works within the client-server model in which the n clients rely on a small number of computationally powerful servers in order to compute aggregate statistics. Prio provides relatively strong privacy guarantees: it guarantees privacy so long as at least one of the computation servers is honest. It also provides robustness: any malicious client can affect the protocol no more than misreporting their private data value as another valid value. That is, if the client is supposed to submit a value in the range $[0, 64]$, Prio servers can syntactically reject submissions of any value outside that range, but clients can of course choose to submit a different value in that range besides their true private value. To achieve this, Prio utilizes a new technique called SNIPs (secret-shared non-interactive proofs), which allows servers to collaboratively check a shared proof of correctness at a low communicational cost. In particular, the bandwidth

used by servers during verification remains constant as the size of user inputs increases. The Prio protocol has been widely adopted, and has even been re-implemented by Mozilla for use in privately collecting web usage statistics. It is being run as a service for other web-based organizations by the Internet Services Research Group (ISRG), a non-profit focused on reducing barriers to secure communication over the internet. Both Apple and Google have also begun using it to perform analytics in their exposure notifications express (ENX) system for measuring health data.

Prio achieves highly desirable security guarantees, and their solution achieves significant efficiency gains over other comparably secure data-collection systems. However, the client-side computation and client-to-server communication of their solution, though better than other comparable systems, each increase at a superlinear rate as the size of user data increases. In the client-server model, client computation and communication costs are most often the bottlenecks for overall efficiency, since clients are on low-power devices and high-latency connections whereas servers are usually collocated high-power machines. For example, clients usually run on either web browsers or cell phones when using the Mozilla Firefox browser. It is somewhat inefficient then to require clients on such devices to use SNIPs to verify simple properties like the size of a user's input, as it places an unnecessarily large computational burden on these weak devices.

In this paper, we present Prio+, a new and improved version of Prio that is optimized to reduce overall burden on the client, not the servers. Prio+ utilizes a Boolean secret-sharing scheme so that clients can prove their data falls within the correct range at essentially zero computational cost. Then, servers simply execute a Boolean-to-arithmetic share conversion protocol (if necessary) to continue computing statistics as in the original Prio protocol. We utilize known methods for share conversion using daBits (double-authenticated bits) [RW19] since their pre-computed nature helps to maximize the efficiency of our protocol's online

phase. For some of our protocols on particular statistics, Prio+ still uses SNIPs, as they are particularly well-suited for verifying certain multiplicative relationships. We do not, however, use them to verify everything about client inputs as is done in Prio. We implemented Prio+ in C++ and our code is publicly available at `https://github.com/KuraTheDog/Prio-plus`.

Our strategy significantly reduces both client computation and client-to-server communication, the two most expensive computational resources in our efficiency model. Even for the few statistics where SNIPs are still necessary (variance and linear regression), the size of the SNIPs and the work necessary for clients to create them decreases dramatically. The result is a system which computes the same set of complex statistics as Prio with identical privacy and robustness guarantees but with reduced client computation and client-to-server communication. For example, when collecting the distribution of tens of thousands of client responses to a simple true/false question, Prio+ clients can encode their data over 350x faster than Prio clients, and the client's message size is nearly 5x smaller. In many cases, Prio+ also improves server efficiency: For the example given above, once servers have received all client inputs, Prio+ servers are able to process client submissions 85x faster with essentially the same server-to-server communication. As the size of inputs increases, Prio+ sees an increase in server communication whereas Prio's server communication stays constant. But, for practically sized inputs, Prio+ servers still communicate only a few hundred bytes per client submission.

Both Prio and Prio+ support computation of more complex statistics in addition to just summation. While Prio applies a general SNIP-based solution across the board, Prio+ applies a specialized approach for each complex statistic and uses SNIPs for a lighter relations as needed. For example, the Prio+ summation protocol, the protocol which is most heavily used by Mozilla, does not require SNIPs at all. This specialized approach means the efficiency of Prio+ varies depending on the statistic being computed. For some such statis-

	Client Mults	Proof Size	Server Mults	Serv Comm.
Prio[CB17]	$M \log M$	M	$(M \log M)n$	n
Linear PCPs [**BBC19**]	$M \log M$	$\log M$	$(M \log M)n$	n
Π_{sum}	None	None	λ	$n + \lambda^2$
$\Pi_{\mathsf{and/or/max/min}}$	None	None	None	n
$\Pi_{\mathsf{var/linReg}}$	$M' \log M'$	M'	$(M' \log M' + \lambda)n$	$n + \lambda^2$
Π_{frq}	None	None	$n\lambda$	$(n + \lambda) \log n$

Table 1.1: Table of asymptotic comparisons between Prio and the protocols Π comprising Prio+. M is the number of multiplication gates required to check if an input is properly encoded for the relevant statistic. λ is the bit-length of each client's input, $M' := M - \lambda$ is the number of multiplication gates not used for checking bit-length, and n is the total number of clients. Note $M \geq \lambda$ because λ multiplication gates are always used for checking bit-length in Prio. 'Serv Comm.' is the communication between servers in bits. Note that Prio and the Linear PCP extension implement a general protocol for all statistics, whereas Prio+ uses different protocols for each statistic. Thus the entries in the first two rows represent all statistics computed by Prio, where the value M varies depending on the particular statistic. Note that $\Pi_{\mathsf{sum}}, \Pi_{\mathsf{var/linReg}}, \Pi_{\mathsf{frq}}$ also use λ symmetric key operations (oblivious transfers) during pre-computation. $\Theta(\cdot)$ notation suppressed to improve readability.

tics, Prio+ not only improves client performance and server computation, but additionally achieves extremely low and constant server communication. As an example, when computing the maximum of client data in the range $[0, 128]$, Prio+ servers communicate a constant 16 bytes per client, compared to a constant 740 bytes for Prio servers. This is in addition to a nearly 750x faster client encode time, 5x smaller client message size, and a 43x improvement in server computation time. Even for the few statistics where Prio+ still utilizes SNIPs, we see significant improvements in client encode time, client message size, and server compute time at little-to-no server bandwidth cost.

Contributions: In what follows we summarize our contributions:

- Provide a detailed daBit-based semi-honest Boolean-to-arithmetic share conversion protocol whose output shares lie in a field $\mathbf{Z}_p$, which was not explicit in [RW19],
- demonstrate how to use Boolean-to-arithmetic share conversion in conjunction with Boolean secret-sharing and smaller-scope SNIPs for particular relations in order to provide robustness and privacy in a large-scale data collection system, and

- demonstrate that client usage of Boolean representation avoids expensive Zero-Knowledge proofs which leads to dramatic speed-ups of the system overall, and
- exhibit the effectiveness of our protocols with a full-scale and publicly available implementation allowing private and robust computation of a wide range of complex statistics.

We now give an informal theorem of our results for a single protocol (Π_{sum}). We establish similar results for the remaining protocols according to the asymptotics given in Table 1.1. Note that although our results are stated in a model using only two servers, all results generalize trivially to k servers and provide the same security guarantees as Prio (security against a coalition of $k-1$ semi-honest servers).

Theorem 1. *(Informal) Suppose n players $P_1, \ldots, P_n$ each hold a private L-bit integer x_i, and they wish to rely on two servers S_L and S_R to compute the sum $f(x_1, \ldots, x_n) = \sum_i x_i$. There exists a protocol Π_{sum}, returning the sum $f(\cdot)$ to each client and returning no output to either server, which is both private and robust against a coalition of up to n malicious clients and one semi-honest server and requires zero client-side multiplications, no zero-knowledge proofs, $O(L)$ multiplications per server (in some prime field $\mathbf{Z}_p$ for large p), $O(L)$ symmetric key operations per server, and $O(n)$ bits of communication between servers.*

With Prio+, we hope to provide the same benefits as Prio to systems and organizations whose clients cannot withstand the burden of generating and sending expensive zero-knowledge proofs.

CHAPTER 2

Technical Overview

In this section we briefly overview the remaining sections of the paper.

Arithmetic vs. Boolean Secret-Sharing: Secret-sharing (specifically, threshold secret-sharing) is a cryptographic tool which allows a user to "share" a private value x into a vector of values $[x] = ([x]_1, \ldots, [x]_n)$ in such a way that any strict subset of these values reveals nothing about x, but all values together can be used to reconstruct x completely. Prio is primarily built around arithmetic secret-sharing: $[x]_i$ are random values in a ring $\mathbf{Z}_M$ (often this is a field and $M = p$) subject to the constraint $\sum_i [x]_i = x \pmod M$. Clients share their private inputs and send one share to each server. Servers then sum them up locally and then return those local aggregations to clients. Clients them combine the local aggregations to learn the total sum. Servers force clients to submit private values within a particular range by requiring clients to also submit specialized zero-knowledge proofs attesting to that fact.

Another possible secret-sharing scheme is called Boolean secret-sharing. Here, the client holds a private L-bit integer x and secret-shares it as $([x]_1, \ldots, [x]_n)$, where each $[x]_i$ is a random L-bit integer subject to the constraint $\bigoplus_i [x]_i = x$, where the direct sum is done component-wise with each bit of the binary representation. This scheme obeys many of the same properties as arithmetic secret-sharing, particularly that each share appears random except when combined with all other shares. The crucial difference is the following: if each share $[x]_i$ is an L-bit integer, it is guaranteed that the private value x is also an L-bit integer.

This allows servers to verify the bit-length of a client's submitted private value via simple local checks on the shares themselves, without needing expensive zero-knowledge proofs. Note: to distinguish Boolean shares from arithmetic shares, arithmetic shares will be denoted $[x]^A$, and Boolean shares will be denoted $[x]^B$.

Boolean-to-Arithmetic Share Conversion: Using Boolean shares presents an issue: how do servers compute the sum over Boolean shares of client inputs? Prio's method of summing shares locally and then returning those aggregated values to clients only works with arithmetic shares, not Boolean shares. Prio+ servers use Boolean-to-arithmetic share conversion, which converts Boolean shares of x to arithmetic shares of the same x. Such protocols have been studied extensively [DSZ15, RW19, EGK20, CDI05], and the most efficient method based on oblivious transfer (OT) is due to [DSZ15] and outputs arithmetic shares in a ring $\mathbf{Z}_M$. For use with SNIPs, M should be a prime.

Semi-Honest Boolean-to-arithmetic Share Conversion into $\mathbf{Z}_p$ via daBits: For more efficient share conversion in this case, we use an offline phase to generate pre-computed daBits (double-authenticated bits) from [RW19], which allow share conversion with less communication and only use one OT to generate each beforehand. This cheap offline phase makes the resulting online protocol much more efficient than Prio. A daBit is a known primitive used for various functionalities, including share conversion. A daBit is a secret-shared pair $([b]^A, [b]^B)$ where $b \in \{0, 1\}$ is a random secret bit. It is known how to convert Boolean shares of an λ-bit integer to arithmetic shares using λ daBits [RW19] in the malicious setting, so we present an efficient semi-honest version of this protocol for quick parallelizable share conversion. Malicious security is not necessary since all servers are semi-honest (following the threat model of Prio), and they perform the conversion. Details of the generation and share conversion, as well as measures of the complexity of these procedures, are in Section 9. Although we would like to provide an analytical comparison between the efficiency of resulting

protocol with Prio, the Prio paper does not give analytic measures of their complexity to enable such a comparison. Thus we rely on a practical comparison of the two systems.

Complex Statistics: Prio+ supports computation of the exact same statistics as Prio. In particular, in addition to SUM, clients can compute Boolean AND / OR (where each client holds a single bit), MAX / MIN and frequency count FRQ (where each client holds a value in some small range $[0, K]$), integer variance VAR and standard deviation STDDEV (where each client holds an L-bit integer), and linear regression linReg (where each client holds a degree d feature vector of L-bit integers).

Many of these statistics (AND, OR, MAX, MIN) are even simpler than SUM, requiring no share conversion, no zero-knowledge proofs, and virtually no communication between servers. FRQ, similar to SUM, requires share conversion to allow summation on the server side, but does not require any zero-knowledge proofs. Instead, servers use some simple logical mechanisms to detect improperly encoded inputs, which we will discuss in the next section.

The only statistics which do require zero-knowledge proofs are VAR, STDDEV, and linReg. In these cases, clients encode their private values in such a way that SNIPs are the most efficient method for verifying that encoding. The key difference is that SNIPs in this case are only being used to verify one small part of the encoding, whereas in Prio they are used to verify every property of the encoded value. At a high level, we have removed the need for SNIPs to verify the length of client inputs, reducing the overall complexity. Since SNIPs operate on arithmetic shares of the input, we first apply daBit-based Boolean-to-arithmetic share conversion on the clients' submitted Boolean shares, which we need to do anyway to perform aggregation after the validation.

Frequency Count: To compute FRQ, P_i is required to submit shares of an impulse

vector δ_{x_i} with value 1 at index x_i and value 0 at all other locations. Servers then sum up these impulse vectors component-wise to output a histogram $\vec{h}$ of the distribution of client values. Verifying that no client submitted a non-impulse vector is somewhat non-trivial: it requires more than simply verifying the length of individual shares, but checking the relation via SNIPs requires many multiplication gates since it is not an inherently multiplicative relation. However, servers can accomplish this task at relatively low cost using some simple logic. First, servers use the Boolean shares of client inputs to check the parity of the number of 1's in each client's submitted vector. This requires passing a single bit per client and reveals nothing about any honest client's input, since an honest input will always be an impulse and thus have an odd number of 1's. This step ensures that no client submitted a zero vector, since that would have an even number of 1's. Next, servers check whether the total number of 1's is equal to the total number of players. If so, this (in combination with the fact that no player submitted a zero vector) implies that all players submitted a single impulse. If not, they can locate the misbehaving player by repeating this check a logarithmic number of times on smaller subsets of the players, honing in on the misbehaving player via binary search. To check whether the total number of 1's is equal to the total number of players, servers simply apply Boolean-to-arithmetic share conversion on each component of each shared vector, locally sum all components of all shared vectors, recombine the total sum and check this against the total number of players. After this check identifies all misbehaving players and those inputs are discarded, servers sum the arithmetic shares of all well-behaving clients and return these summed vectors to the clients who sum them together to get the histogram $\vec{h}$.

Practical Comparison: We compared Prio+ to both the original Go implementation of Prio as well as another implementation by Mozilla which only computes SUM, no complex statistics. We did not compare to the updated construction given in [BBC19] because it has not yet been implemented and is focused on the case of extremely long inputs. In the case of practically-sized inputs, large constant terms overshadow the asymptotically smaller client

message sizes.

First, we compared all three implementations in evaluating the $\mathsf{SUM}(x_1, \ldots, x_n)$ where $n = 10,000$, $n = 50,000$, and $n = 100,000$. We recorded the encode time (milliseconds per client), client message size (bytes per client per server), server compute time (milliseconds per server per client), and server communication (bytes per server per client) and averaged the results from each value of n. We ran four separate experiments in which x_i was 1-bit, 8-bits, 16-bits, and 32-bits respectively. Prio+ clients were able to encode their data up to 540x faster than in the Go implementation of Prio, and up to 3000x faster than in the Mozilla implementation. Client messages in Prio+ were over 3x smaller than in the Mozilla implementation, and up to 23x smaller than in the Go implementation. Although reducing server computation time was not the primary goal of this project, Prio+ servers processed client submissions up to 116x faster than the Go implementation and up to 615x faster than the Mozilla implementation. The expected drawback of these savings was server communication, since conserving server bandwidth was the primary motivation for Prio's SNIPs. Prio+ does see increased server bandwidth in some cases, particularly as the size of user inputs increases. However, the practical bandwidth usage for 1-bit integers is essentially the same as in the Mozilla implementation and 18x less than the Go implementation. By comparison, for 32-bit inputs, Prio+ server communication is still 3x less than the Go implementation, and just 7x more than the Mozilla implementation. This tells us that for practically-sized user inputs, Prio+ achieves monumental improvements in client computation, client communication, and server computation with minimal impact to server communication.

We also ran similar experiments between Prio+ and the Go implementation of Prio for MAX and linReg. Since MAX requires no SNIPs and no share conversion, we saw improvements across the board: when client values lied in the range $[0, 128]$, Prio+ client encode time was 750x less, client message size was 5x smaller, server compute time was 43x less,

and server communication was 46x less. Even though linReg still requires some SNIPs (with reduced scope), we saw up to 30x lower client encode time, up to 4x smaller client message size, up to 3x less server compute time, and server communication varying between 5x less (for degree 2 inputs) and 1.5x more (for degree 8 inputs). Prio+ is clearly more efficient across the board when computing MAX and low-degree linReg. For higher-degree linReg, we see significant gains in encode time, client message size, and server compute time for a slight increase in server communication.

CHAPTER 3

Preliminaries

In this section we will describe our computational model. This includes a description of our ideal functionality, the client/server setup, our efficiency model, the set of adversaries we defend against, and the assumptions we rely upon to build our protocol.

"Two-Party" Setting: In this work we construct protocols for secure computation of a wide range of aggregate statistics in the client-server model. That is, a set of n clients with private data wish to compute statistics on that data with the help of two honest-but-curious servers. The basis of our system is a secure two-party protocol between these servers. That is to say that each client with an input, secret-shares his/her input between the two computation servers (which are assumed to not collude). Then, the two computation servers run the secure two-party computation protocol on the input shares which does not reveal to either server any information about client inputs. Finally, they send the output shares back to the clients who then reconstruct the output. This technique of using secure two-party computation in the client-server setting was first described in the ABY framework of [DSZ15] and has been used in many applications since, including [CB17].

Just as in the ABY framework of [DSZ15], this also allows for reactive computations in which the two computation servers maintain some secure state information between multiple executions. This could be used in the case where client data is being collected over time and the data is sent to the servers one point at a time.

Formally, our deployment consists of n clients $\{P_1, \ldots, P_n\}$ and 2 servers S_L and S_R. Each client P_i holds some private input x_i. Each client can communicate with each server and servers can communicate with clients and each other via private channels, but clients do not communicate with each other. Side note: although Prio+ is described as a two-server protocol, it can be easily generalized to a k server protocol for any positive integer k by simply using k-wise instances of each primitive (secret-sharing, daBits, SNIPs as necessary).

Efficiency Model: We assume clients have low computational power and servers have high computational power. Similarly, we assume a low-bandwidth connection from clients to servers, and a high-bandwidth connection between servers. Thus we seek to minimize client computation and communication as our highest priority, and server costs as an afterthought. In general, we assume network latency is the greatest computational bottleneck and are concerned more with optimizing communication than computation for both clients and servers.

Security Against Semi-Honest Servers, Malicious Clients: Our deployment protects client *privacy* as long as at most one server is passively corrupted (regardless of malicious client misbehavior). Our system cannot tolerate malicious, misbehaving servers as this comes at a direct cost of functionality, as discussed in [CB17]. Our deployment always provides *robustness* (correctness) so long as neither server maliciously misbehaves. We summarize our security definitions here, please refer to Appendix B for details.

Privacy: Intuitively, our deployment provides f-privacy for an aggregation function (statistic) f if an adversary controlling any number of clients and all but one server learns nothing about the honest clients' inputs besides what is revealed by the output of f. More formally, any such adversary can simulate its view of the protocol run given the output of f. For some aggregation functions, we weaken our protocol to provide $\hat{f}$-privacy where $\hat{f}$ leaks slightly more information than the statistic itself (for example, servers may leak the number

of clients who provided invalid inputs).

Robustness: A protocol is t-robust if a coalition of t malicious clients cannot affect the output of the protocol beyond misreporting their private data values. This is the strictest notion of correctness in the malicious security model, since a client's private input is known to nobody but themselves, meaning we cannot prevent them from misreporting it. If this data value is meant to come from a specific domain, however, malicious clients should *not* be able to submit data from outside of that domain. This is particularly relevant in our setting, where client data must be encoded correctly to permit efficient computation of the aggregate statistic. Clients should absolutely not be able to submit improperly encoded data, as this allows a single client to affect the output arbitrarily without detection. Our deployment is robust against malicious clients, but not against malicious servers. Though robustness against malicious servers may seem desirable, it comes at a direct cost to performance, as argued in [CB17]. Since the number of clients is much larger than the number of servers, it is much more reasonable to prevent and/or replace faulty servers than faulty clients.

Analogously to [CB17], we assume cryptographic primitives for the establishment of pairwise authenticated channels (CCA-secure public key encryption [CS98], digital signatures [Sho01a, Sho01b], etc.). We make no synchrony assumptions about our network and do not rely on external systems to provide users anonymity.

Notation: We write $x \oplus y$ to denote the XOR operation (addition modulo 2), $x +^l y$ for addition within the ring $\mathbf{Z}_{2^l}$, and $x +^p y$ for addition in the field $\mathbf{Z}_p$. When $\vec{x}, \vec{y} \in \mathbf{Z}_{2^m}$ are vectors of bits, we will write $\vec{z} = \vec{x} \oplus \vec{y}$ to denote the bitwise-XOR operation. That is, $(\vec{z})_i = (\vec{x})_i \oplus (\vec{y})_i$ for each $0 \leq i < m$. We assume a maximum bit-length l on all integer data and thus treat all integer-valued data as elements of the ring $\mathbf{Z}_{2^l}$, except in the case where we perform share conversion on client data. In that case, the resulting shares will lie

within the field $\mathbf{Z}_p$, as this is required for our polynomial identity testing procedure. This of course requires that $p > 2^l$, which will always be the case. We denote an arithmetically secret-shared variable x by $[x]^A$. A variable x shared in the Boolean secret-sharing scheme is denoted $[x]^B$. We will exclusively use two-party secret-sharing, and thus shares held by server S_L will be written $[x]_L^t$, $t \in \{A, B\}$. Shares held by server S_R will similarly be written $[x]_R^t$. For an integer x, we refer to the i'th least significant bit of the binary representation of x as $(x)_i$. We will say that a function $f : \mathbf{N} \longrightarrow \mathbf{R}$ is negligible if for every positive polynomial poly there exists an integer N_{poly} such that for $x > N_{\mathsf{poly}}$, $|f(x)| < \frac{1}{\mathsf{poly}(x)}$.

CHAPTER 4

Necessary Primitives

For our purposes, we focus on having $N = 2$ servers with secrets shared between them. Clients hold a secret value x, and want to split it into two shares $\mathsf{Share}(x) = [x]_L, [x]_R$ for servers L and R. This can also be reversed, where $\mathsf{Rec}([x]_L, [x]_R) = x$. Privacy here is straightforward, where any one server can't recover the secret, but both together can. Correctness means that Rec succeeds in the presence of both shares.

Definition 4.1. Arithmetic Secret-Sharing: Given an integer $x \in \mathbf{Z}_M$, an arithmetic secret-sharing of x is a random pair $a, b \in \mathbf{Z}_M$ subject to the condition $a + b = x \pmod{M}$.

Semantics: The two-party arithmetic secret-sharing scheme consists of the following pair of functions:

- $\mathsf{Share}_{+,M} : \mathbf{Z}_M \longrightarrow (\mathbf{Z}_M)^2$, $\mathsf{Share}_{+,M}(x) = ([x]_L^A, [x]_R^A)$, which are random elements of $\mathbf{Z}_M$ subject to the constraint $[x]_L^A + [x]_R^A = x \pmod{M}$.
- $\mathsf{Rec}_{+,M} : (\mathbf{Z}_M)^2 \longrightarrow \mathbf{Z}_M$, $\mathsf{Rec}_{+,M}([x]_L^A, [x]_R^A) = [x]_L^A + [x]_R^A \pmod{M}$.

Addition/Scalar Multiplication: Addition and scalar multiplication over arithmetic secret shares are trivial. To compute a share of $z = x + y$ given shares $[x]^A$ and $[y]^A$, each server $i \in \{L, R\}$ locally computes $[z]_i^A = [x]_i^A + [y]_i^A$. Similarly, to compute scalar multiplication $[w]^A = c \cdot [x]^A$ for public $c \in \mathbf{Z}_M$, each server locally computes $[w]_i^A = c \cdot [x]_i^A$.

Definition 4.2. Boolean Secret-Sharing Given an integer $x \in \mathbf{Z}_2$, a Boolean secret-sharing of x is a random pair $c, d \in \mathbf{Z}_{2^\lambda}$ subject to the condition $c \oplus d = x$.

Semantics: The two-party λ-bit Boolean secret-sharing scheme consists of the following pair of functions:

- $\mathsf{Share}_{\oplus,\lambda} : \mathbf{Z}_{2^\lambda} \longrightarrow (\mathbf{Z}_{2^\lambda})^2$, $\mathsf{Share}_{\oplus,\lambda}(x) = ([x]^B_L, [x]^B_R)$, which are random elements of $\mathbf{Z}_{2^\lambda}$ subject to the constraint $[x]^A_L \oplus [x]^A_R = x$.
- $\mathsf{Rec}_{\oplus,\lambda} : (\mathbf{Z}_{2^\lambda})^2 \longrightarrow \mathbf{Z}_{2^\lambda}$, $\mathsf{Rec}_{\oplus,\lambda}([x]^A_L, [x]^A_R) = [x]^A_L \oplus [x]^A_R$.

XOR: Computing XOR over Boolean shares is trivial. $[z]^B = [x]^B \oplus [y]^B$. Each server $s \in \{L, R\}$ locally computes $[z]^B_s = [x]^B_s \oplus [y]^B_s$.

These two schemes each have strengths and weaknesses. If client data is arithmetically secret-shared under a large modulus M, servers can efficiently compute the sum of shared values via associativity of addition by locally summing their shares modulo M, as in Prio [CB17]. This is not efficient with Boolean shares as they are built using bitwise XOR instead. On the other hand, Boolean shares of x are the same bit-length as x itself, meaning servers can trivially verify the size of client inputs. Prio+ leverages both of these advantages to privately compute complex aggregate statistics efficiently. Prio+ clients submit their data in the Boolean scheme (so that servers can verify the bit-length efficiently) and then servers convert these shares back to the arithmetic scheme in order to sum the data together and compute the given statistic.

Boolean to arithmetic share conversion is a well-studied technique. The current most efficient protocol in the semi-honest two-party setting is due to [DSZ15] and is based on Oblivious Transfer (OT). In particular, to convert a pair of λ-bit Boolean shares to arithmetic shares in $\mathbf{Z}_{2^\lambda}$, they use λ independent instances of OT where each OT transfers on

average a string of length $(\lambda+1)/2$. The total communication cost is $\lambda(\kappa+(\lambda+1)/2) = O(\lambda^2)$ [DSZ15].

To achieve share conversion with more efficient online work, we utilize precomputed pairs called daBits (doubly-authenticated bits), discussed in [RW19]. Although they are primarily used in the malicious setting, we are able to use them very efficiently in the semi-honest setting, which to our knowledge hasn't been detailed explicitly. Compared to OT share conversion above, for the same number of bits converted this uses the same number of OTs for generating the precomputed daBits, and then only communicates a single bit between servers, per bit converted using daBits. They require the same OTs to generate as the share conversion protocol in [DSZ15], and only require a single bit communicated per converted bit to perform the computation. See Section **??** for further details.

The final piece of Prio+, used only for a few statistics, is the secret-shared non-interactive zero-knowledge proof (SNIP) which underpins the Prio protocol of [CB17]. Although we claim that Prio overuses SNIPs in unnecessary situations, SNIPs are an incredibly efficient method for verifying multiplicative relationships on secret-shared inputs. Below we review how SNIPs allow servers to efficiently verify that some client input x is valid without learning any additional information. The following description comes directly from [CB17].

A secret-shared non-interactive proof (SNIP) protocol consists of an interaction between a client (the prover) and multiple servers (the verifiers). At the start of the protocol:

- Each server i holds a share $[x]_i^A \in \mathbf{F}^\lambda$ for some field $\mathbf{F}$.
- The client holds the secret input (vector) $x = \sum_i [x]_i^A \in \mathbf{F}^\lambda$.
- All parties hold an arithmetic circuit representing $\mathsf{Valid} : \mathbf{F}^\lambda \longrightarrow \mathbf{F}$.

The client's goal is to convince the servers that $\mathsf{Valid}(x) = 1$ without revealing any additional information about x. To do so, the client sends a proof string to each server. After receiving these proof strings, the servers gossip amongst themselves and then conclude either that $\mathsf{Valid}(x) = 1$ (accept x) or $\mathsf{Valid}(x) \neq 1$ (reject x).

A valid SNIP must satisfy correctness, soundness, and zero-knowledge.

Correctness. If all parties are honest, the servers will accept x.

Soundness. If all servers are honest, and if $\mathsf{Valid}(x) \neq 1$, then for all malicious clients, even ones running in super-polynomial time, the servers will reject x with overwhelming probability. In other words, no matter how the client cheats, the servers will almost always reject.

Zero-knowledge. If the client and at least one server are honest, then the servers learn nothing about x, except that $\mathsf{Valid}(x) = 1$. More precisely, there exists a simulator (that does not take x as input) that accurately reproduces the view of any proper subset of malicious servers executing the SNIP protocol.

The construction in [CB17], based on a generalized version of the polynomial-based batched multiplication verification technique of Ben-Sasson et al. [BFO12], satisfies each of these properties as proven in their Appendix D.

CHAPTER 5

Private Summation Without Robustness

Our protocols are all based on the following simple scheme for computing the sum of clients' private bits. This is also the basis of Prio's solution and is described in [CB17]. We reiterate that scheme here for convenience.

Each client P_i, $i \in \{1, \ldots, n\}$, holds a private bit $x_i \in \{0, 1\}$. They wish to learn the sum $f(x_1, \ldots, x_n) = \sum_{i=1}^{n} x_i$. Even this basic functionality has wide-ranging applications for data collectors, since it allows one to survey clients on any yes/no question and learn the distribution of responses. Consider the following protocol for computing the sum.

1. **Upload**: Each P_i computes $\mathsf{Share}_{+,M}(x_i) = ([x_i]_L^A, [x_i]_R^A)$. The client then sends one additive share to each server over secure pairwise-authenticated channels. Note: although x_i is a single bit, we treat it here as an element of $\mathbf{Z}_M$ for $M > n$.

2. **Aggregate**: S_L and S_R hold accumulator values $A_L, A_R \in \mathbf{Z}_p$ respectively, initially set to zero. For each i, when S_L receives $[x_i]_L^A$ from P_i, computes $A_L \leftarrow A_L + [x_i]_L^A \pmod{M}$. S_R does the same with its accumulator A_R upon receiving $[x_i]_R^A$ from P_i.

3. **Publish**: Once data is collected, servers publish their accumulator values A_L, A_R to every client.

4. **Client Computation**: Clients compute the sum of the accumulator values $A_L + A_R \pmod{M}$.

Note that if all players behave, each client's output is

$$A_L + A_R \pmod{M} = \sum_i [x_i]_L^A +^p \sum_i [x_i]_R^A \pmod{M} = \sum_i x_i \pmod{M}$$

Since each x_i is at most 1, requiring $M > n$ ensures that $\sum_i x_i$ never overflows the modulus 2^l. This means that the output, interpreted as an integer, is indeed the sum $\sum_i x_i$.

The authors of [CB17] make two crucial observations about this simple scheme. First, it provides privacy as long as one server is honest. The adversary's view only includes a single share, say $[x]_L^A$, of an honest client's input x, which appears totally random without $[x]_R^A$. Second, the scheme does *not* provide robustness against malicious clients. A single malicious client can completely corrupt the protocol output by submitting (for example) arithmetic shares of a random integer $r \in \mathbf{Z}_M$ to each server.

The authors of Prio [CB17] solve the client robustness issue by forcing clients to construct and submit SNIPs (secret-shared non-interactive proofs), a novel type of zero-knowledge proof which allows the servers to non-interactively verify the form of client inputs. This is effective, but requires somewhat expensive computation and communication on the clients' part to construct and send these SNIPs to the servers. Since clients are presumed to be computationally weak compared to servers in our model, latency between clients and servers is high, this is not ideal. We would rather invoke some extra communication and computation among the servers if it would allow us to reduce the communication and computation on the client side.

We observe that clients can solve the robustness issue by instead submitting their data using the 1-bit Boolean secret-sharing scheme. This allows servers to verify that $x \in \{0, 1\}$

by simply confirming that each share is a single bit, forcing the client to sharing a value which is only a single bit in length. Then, servers can execute a share conversion protocol to privately convert these into arithmetic shares of the same secret values in the larger ring $\mathbf{Z}_M$ and sum them up as before. This observation is the intellectual core of our deployment.

In the next section, we give further details of how we protect against malicious clients using share conversion. In particular, we show how to apply our observation described above to prevent clients from submitting data outside the range $\{0, 1\}$. This can be trivially generalized to verifying that client data falls within the range $[0, 2^l - 1]$. In future sections, we show how servers can verify other aspects of user inputs besides the bit-length. These tools will be necessary for computing more complex statistics. In that case, clients will encode their inputs before sharing with the servers, and servers will need tools to verify that submitted data is properly encoded. The authors of Prio [CB17] once again accomplish this via SNIPs, but we will show that such machinery is, in most cases, wasteful and unnecessary for computing the desired statistics.

CHAPTER 6

Protecting Correctness

In this section we give further details about share conversion and how we use it to make Prio+ robust against malicious clients.

The reason a malicious client can cheat and submit $x_i \notin \{0, 1\}$ is that when x_i is shared arithmetically in $\mathbf{Z}_M$, a single share reveals nothing about the size of x_i. From the servers' perspective, the underlying data x_i could be any element of $\mathbf{Z}_M$. Thus, server S_L holding a single share $[x_i]_L^A$ cannot tell whether $x_i \in \{0, 1\}$.

Imagine, however, that P_i shares x_i via the 1-bit Boolean scheme as $x_i = [x_i]_L^B \oplus [x_i]_R^B$. If $[x_i]_L^B \in \{0, 1\}$ and $[x_i]_R^B \in \{0, 1\}$, then it is guaranteed that $x_i \in \{0, 1\}$. Using Boolean to arithmetic share conversion, servers can then compute arithmetic shares $[x_i]^A$ of the value x_i in the extended ring $\mathbf{Z}_M$ and continue computation according to the simple scheme. If this conversion is done securely, these shares will have the same amount of entropy as the original 1-bit Boolean shares, effectively hiding the underlying secret value. Thus, all we need in order to make the simple scheme robust against malicious clients is a Boolean to arithmetic share conversion protocol achieving the following ideal functionality $\mathcal{F}_{B2A}$ described below.

Definition 6.1. The two-player l-bit ideal functionality $\mathcal{F}_{B2A}$ (with output in $\mathbf{Z}_M$) behaves as follows:

- $\mathcal{F}_{B2A}$ receives $[x]_L^B, [x]_R^B \in \mathbf{Z}_{2^l}$ as inputs from S_L, S_R respectively.
- $\mathcal{F}_{B2A}$ computes $x = [x]_L^B \oplus [x]_R^B$

- $\mathcal{F}_{B2A}$ computes $\mathsf{Share}_{+,M}(x) = ([x]_L^A, [x]_R^A)$ satisfying $[x]_L^A + [x]_R^A = x \pmod{M}$.
- $\mathcal{F}_{B2A}$ returns $[x]_L^A, [x]_R^A$ to S_L, S_R respectively as outputs.

Share conversion has been well-studied, both in its theoretical limitations [CDI05] and its practical performance [DSZ15]. In this case we use a daBit (double-authenticated bit) based Boolean-to-arithmetic share conversion (see [RW19] for discussion of daBits). Share conversion is further detailed in Section 9.

From now on, we will use the notation $\mathsf{B2A}_{l,M}([x]_L^B, [x]_R^B)$ to represent an evaluation of a Boolean to arithmetic share conversion protocol with l-bit inputs and whose output lies in $\mathbf{Z}_M$. When the bitlength of the input shares and the output range are clear by context, we may simply write $\mathsf{B2A}([x]_L^B, [x]_R^B)$

With $\mathsf{B2A}$ in our toolbox, we can now strengthen the simple scheme from the previous section to prevent malicious clients from corrupting the output. In particular, we force clients to submit single bits by making them submit data under the 1-bit Boolean secret-sharing scheme. Then, servers convert the Boolean shares into the arithmetic scheme using $\mathsf{B2A}$. Servers then compute the sum of the shared data according to the simple scheme. Servers only accept shares consisting of a single bit, which means that if both servers accept shares $[x_i]_L^B, [x_i]_R^B$ from P_i, we are guaranteed that $[x_i]_L^B \oplus [x_i]_R^B \in \{0, 1\}$. This precisely guarantees robustness in the sense that a malicious client cannot affect the output of the protocol beyond misreporting their private data. A detailed description of this strengthened protocol can be found below. Detailed proofs of its privacy and robustness can be found in Appendix C.

Π_{bitSum}

Inputs: $x_i \in \{0,1\}$ for $i \in [n]$.
Output: $\sum_{i=1}^{n} x_i$.

1. **Upload**:
   (a) Each client P_i computes $\mathsf{Share}_{\oplus,1}(x_i) \longrightarrow [x_i]_L^B, [x_i]_R^B$ via Definition 4.2
   (b) Each P_i sends $[x_i]_L, [x_i]_R$ to S_L, S_R respectively.
2. **Verify Bit-Length**: Initially, $n' = n$. If a server receives a share which is not 1 bit in length from P_i (assume S_L w.l.o.g.):
   (a) S_L sends the index i to S_R.
   (b) Both servers discard $[x_i]^B$.
   (c) Both servers set $n' \leftarrow n' - 1$
3. **Convert Shares**: S_L and S_R jointly evaluate $\mathsf{B2A}_{1,2^\lambda}(\{[x_i]_L^B, [x_i]_R^B\})$ on each of the n' valid pairs of Boolean shares. S_L receives as output $\{[x_i]_L^A\}_i$ and S_R receives as output $\{[x_i]_R^A\}_i$.
4. **Aggregate**: S_L locally adds all arithmetic shares into an accumulator A_L, initially zero. That is: $A_L \longleftarrow A_L + \sum_i [x_i]_L^A$. S_R analogously accumulates its arithmetic shares into $A_R \longleftarrow A_R + \sum_i [x_i]_R^A$.
5. **Publish**: Once all n' shares have been accumulated, S_L and S_R publish A_L and A_R to every client.
6. **Client Computation**: Clients output $A_L + A_R$.

As we can see, this minor modification of the simple scheme guarantees both privacy and robustness without any heavy client-side computation. It also easily generalizes to compute the sum of l-bit integers for $l > 1$ by simply sharing data in the l-bit Boolean scheme and then using Boolean to arithmetic share conversion. If servers intentionally leak the number of honest clients n', then this protocol for computing the sum is sufficient for computing the arithmetic mean as well. The authors of Prio [CB17] extend their scheme beyond just computing the sum and arithmetic mean. Prio supports computation of a wide array of aggregation functions including: variance (VAR), standard deviation (STDDEV), Boolean OR and AND, integer MIN and MAX, frequency count (FRQ), and linear regression (linReg). They accomplish this by having users encode their input in particular ways such that the sum of the encoded inputs reveals the desired statistic. These are discussed in the literature as affine aggregatable encodings (AFEs), and we refer the reader to [CB17] for more information regarding these encodings.

The only obstacle in computing these statistics analogously using our system is to design a way for servers to successfully verify that clients' inputs are properly encoded. If this can be done, our Prio+ servers can verify that client data is properly encoded (and is within the proper range), perform share conversion, sum the resulting arithmetic shares and return the sum of the properly encoded inputs using the exact same technique as our Π_{bitSum} protocol. This sum of encoded inputs will be precisely the desired statistic according to the underlying AFE. Instead of universally relying on SNIPs for verifying client encodings, as Prio does [CB17], we use SNIPs sparingly. We use them exclusively for verifying multiplicative relationships within the encoding, and all other properties of the encoded inputs are verified via alternative, novel methods. We believe that this is a more appropriate application of SNIPs since they are, at their core, designed around verifying the outputs of multiplication gates within a validity circuit.

In the following section, we outline each of the more complex statistics that our scheme computes. For each statistic we describe the corresponding encoding for which summing the encoded inputs produces the desired statistic.

CHAPTER 7

Complex Statistics

In this section, we describe each of the statistics our deployment computes. For each statistic, we describe the corresponding encoding which we use to enable computation. That is, how can we compute an encoding of x_i (written $\mathsf{en}(x_i)$) so that the statistic f we wish to compute is given by $f(x_1, \ldots, x_n) = \sum_i \mathsf{en}(x_i)$ (or some locally computable function of this sum). These encodings are referred to as "affine aggregatable encodings," or AFEs. As most of the machinery of AFEs is irrelevant to our applications, we omit a detailed discussion and instead refer readers to [CB17] for more information.

As a reminder, each client P_i holds input x_i from some secret-space $\mathcal{D}$ which will be encoded as $\mathsf{en}(x_i)$. They wish to compute $f(x_1, \ldots, x_n)$ using servers S_L, S_R. The servers are responsible for verifying that $\mathsf{en}(x_i)$ is a proper encoding of some $x_i \in \mathcal{D}$, as well as for summing these encodings and returning them to clients so they can reconstruct the value $f(x_1, \ldots, x_n)$. For each statistic f, we give the domain $\mathcal{D}$ of the input x_i (also referred to as the "secret-space") as well as the encoding we will use for computing f. We will also give a brief intuition of why this encoding is sufficient and how the servers will use it to compute the statistic.

SUM, MEAN:

$\mathsf{SUM}(x_1, \ldots, x_n) = \sum_{i=1}^{n} x_i$

$\mathsf{MEAN}(x_1, \ldots, x_n) = \frac{1}{n} \cdot \mathsf{SUM}(x_1, \ldots, x_n)$

Secret-Space: $\mathcal{D} = \mathbf{Z}_{2^l}$

Encoding: $\mathsf{en}_{\mathsf{int}}(x_i) = x_i$

Intuition: Clients submit their data secret-shared via the l-bit Boolean scheme. Servers use $\mathsf{B2A}$ to convert valid l-bit Boolean shares to arithmetic shares of the same secret x_i in $\mathbf{Z}_M$ for $M > n$ and then locally sum the resulting arithmetic shares. In order to compute the mean, we allow the servers to modestly leak the number of players $n - c$ whose valid shares are included in the aggregate. Then clients can locally compute the mean. Note: this means our integer mean protocol achieves only $\hat{f}$-privacy, where $\hat{f}$ leaks $n - c$.

AND, OR:

$\mathsf{AND}(x_1, \ldots, x_n) = 1 \iff \forall i, x_i = 1$

$\mathsf{OR}(x_1, \ldots, x_n) = 0 \iff \forall i, x_i = 0$

Secret Space: $\mathcal{D} = \{0, 1\}$

Encoding: $\mathsf{en}_{\mathsf{and}}(x_i) = (1 - x_i)\vec{r} \in \mathbf{F}_2^\lambda$ for some security parameter λ and random $\vec{r} \in \mathbf{F}_2^\lambda$. That is, if $x_i = 1$, $\mathsf{en}_{\mathsf{and}}(x_i) = \vec{0}$, and if $x_i = 0$, $\mathsf{en}_{\mathsf{and}}(x_i) = \vec{r}$.

$\mathsf{en}_{\mathsf{or}}(x_i) = x_i \cdot \vec{r} \in \mathbf{F}_2^\lambda$ for some security parameter λ and random $\vec{r} \in \mathbf{F}_2^\lambda$. That is, if $x_i = 0, \mathsf{en}_{\mathsf{or}}(x_i) = \vec{0}$, and if $x_i = 1, \mathsf{en}_{\mathsf{or}}(x_i) = \vec{r}$.

Intuition: Here, share conversion is unnecessary because the aggregation operator and the reconstruction operator for the Boolean secret-sharing scheme are both XOR. Thus, servers

can simply locally XOR their valid shares and publish these aggregated values to clients, who will then XOR the aggregated values together to produce the output. When computing AND, if every client has $x_i = 1$, then every $\mathsf{en}(x_i) = \vec{0}$, and so the XOR of these encodings will certainly be $\vec{0}$. In this case, clients can conclude $\mathsf{AND}(x_1, \ldots, x_n) = 1$. Otherwise, if some client has $x_i = 0$, then $\mathsf{en}(x_i) = \vec{r}$ and the XOR of the encodings will be non-zero with probability $1 - \frac{1}{2^\lambda}$. In this case, they conclude that $\mathsf{AND}(x_1, \ldots, x_n) = 0$. The argument is analogous in the case of OR.

MAX, MIN:

$\mathsf{MAX}(x_1, \ldots, x_n) = \max_i x_i$

$\mathsf{MIN}(x_1, \ldots, x_n) = \min_i x_i$

Secret Space: $\mathcal{D} = \{0, \ldots, M\}$ for small $M \in \mathbf{Z}$

Encoding: $\mathsf{en}_{\mathsf{max}}(x_i) = (\vec{r}_1, \ldots, \vec{r}_i, \vec{0}, \ldots, \vec{0}) \in \mathbf{F}_2^{\lambda \times M}$, where each $\vec{r}_j \in \mathbf{F}_2^\lambda$ is independently random. This is equivalent to applying the $\mathsf{en}_{\mathsf{or}}()$ function to each component of the vector $(1, \ldots, 1, 0, \ldots, 0) \in \mathbf{F}_2^M$ where the first i components are 1.

Intuition: To compute the maximum, servers run the OR protocol M times in parallel on each component of the encoded input. That is, they analogously XOR their shares locally, and return them to clients to XOR and reconstruct the output. The clients parse this $(\lambda \times M)$-bit string in λ-bit chunks, reading each chunk as a 0 if and only if every bit of that chunk is 0. The clients compute the largest index k for which the corresponding OR protocol gave output 1, and conclude the maximum is k. That is, they compute the largest value k such that the k'th substring of λ consecutive bits contains a 1. This is certainly bounded above by the maximum, and the probability that it undershoots the maximum by Δ is $\frac{1}{2^{\lambda \times \Delta}}$,

which is negligible in the security parameter λ.

To compute the minimum, clients first represent their input x_i as $\chi_i = M - x_i$, and compute the maximum of these χ_i values, $y_i = \mathsf{MAX}(\chi_1, \ldots, \chi_n)$, as above. The desired minimum will be precisely $\mathsf{MIN}(x_1, \ldots, x_n) = M - y_i$, with the same error bounds as the MAX protocol.

VAR, STDDEV:

$\mathsf{VAR}(x_1, \ldots, x_n) = \frac{1}{n} \sum_{i=1}^{n} (x_i - \mathsf{MEAN}(x_1, \ldots, x_n))^2$

$\mathsf{STDDEV}(x_1, \ldots, x_n) = \sqrt{\mathsf{VAR}(x_1, \ldots, x_n)}$

Secret Space: $\mathcal{D} = \mathbf{Z}_{2^l}$

Encoding: $\mathsf{en}_{\mathsf{var}}(x_i) = (x_i, x_i^2)$

Intuition: Servers parse the encoded input into its two parts and compute shares of $(\sum_i x_i, \sum_i x_i^2)$ using two parallel instances of our protocol for SUM, which they return to clients. The clients divide these values by n, which is a public parameter, to compute $\mathbf{E}[X]$ and $\mathbf{E}[X^2]$, where X is a random variable taking on each value x_i with equal probability. From this, clients can locally compute $\mathsf{VAR}(x_1, \ldots, x_n) = \mathbf{E}[X^2] - (\mathbf{E}[X])^2$. Note: in the case where clients may misbehave, this protocol is only $\hat{f}$-private, where $\hat{f}$ leaks $\mathbf{E}[X]$ and the remaining number of behaving players n' in addition to the output. Clients who wish to compute the standard deviation simply add a local square root operation to the end of the protocol.

linReg:

$\mathsf{linReg}((x_1, y_1), \ldots, (x_n, y_n)) = (c_0, c_1)$, where $\hat{y}(x) = c_0 + c_1 x$ is the unique line which minimizes the sum of squares loss $\sum_i (y_i - \hat{y}(x_i))^2$.

Secret Space: $\mathcal{D} = \mathbf{Z}_{2^l} \times \mathbf{Z}_{2^l}$

Encoding: $\mathsf{en}_{\mathsf{reg}}(x_i, y_i) = (x_i, x_i^2, y_i, x_i y_i)$

Intuition: Analogously to the variance computation, servers compute the sum of the various parts of the encoding in parallel via our protocol for SUM and return shares of $(\sum_i x_i, \sum_i x_i^2, \sum_i y_i, \sum_i x_i y_i)$ to all clients. The clients can solve for the desired real regression coefficients c_0 and c_1 locally using the following linear system:

$$\begin{pmatrix} n & \sum x_i \\ \sum x_i & \sum x_i^2 \end{pmatrix} \cdot \begin{pmatrix} c_0 \\ c_1 \end{pmatrix} = \begin{pmatrix} \sum y_i \\ \sum x_i y_i \end{pmatrix} \tag{7.1}$$

Note that again, in the case of misbehaving clients, this implies servers must also reveal the value $n - c$ of remaining players in the protocol, introducing a modest leakage. Thus this protocol will also be $\hat{f}$-private, where $\hat{f}$ leaks $n - c$ in addition to the output. This technique also trivially generalizes to d-dimensional client inputs $(x^{(0)}, x^{(1)}, \ldots, x^{(d)})$ for $d > 2$ as described by [CB17].

FRQ:

$\mathsf{FRQ}(x_1, \ldots, x_n) = \vec{h} = (f_1, \ldots, f_k) \in \mathbf{Z}_{n+1}^k$, where $f_j = |\{x_i : x_i = j\}| \leq n$ is the frequency of input $j \in \mathbf{Z}_k$

Secret Space: $\mathcal{D} = \{0, \ldots, k-1\}$ for small $k \in \mathbf{Z}$

Encoding: $\mathsf{en}_{\mathsf{frq}}(x_i) = (\delta_{x_i}) \in \mathbf{Z}_{2^k}$, where the x_i'th component of (δ_{x_i}) is 1 and all other components are 0. That is, (δ_{x_i}) is an impulse at x_i.

Intuition: If all players behave, taking the sum of these encodings yields the desired vector $\vec{h}$. Thus, servers evaluate k independent instances of $\mathsf{B2A}$ on each vector (one for each com-

ponent) and then locally sum the resulting vectors of arithmetic shares. They then publish these aggregated vectors to clients who add them together to get $\vec{h}$ in the same manner as our SUM protocol.

We have now intuitively argued that, conditioned on the fact that all inputs are properly encoded, servers can efficiently and privately compute each of these desired statistics. All that remains is to describe the method by which servers can privately verify that shares $[x_i]_L, [x_i]_R$ reconstruct proper encoding of an element $x_i \in \mathcal{D}$ for each of the encodings described above.

In the case of SUM, MEAN, AND, OR, MAX, and MIN, we simply observe that any Boolean vector of the correct length is a valid encoding of some input in the relevant secret space. This makes our job simple, since servers can trivially verify the bit-length of the input by checking the bit-length of its shares.

For VAR, STDDEV, and linReg, servers must verify multiplicative relationships among different parts of the secret-shared encoded inputs. We accomplish this using SNIPs analogously to the methods of [CB17] described in Section 2. For a more detailed description of SNIPs construction and usage, see [CB17].

The trickiest encoding to verify is $\mathsf{en}_{\mathsf{frq}}$. Here, servers are given Boolean shares of v and they must verify that $\sum_i (v)_i = 1$. That is, v has exactly one component equal to 1. In other words, v is an impulse. This is a difficult encoding to verify because the number of valid encodings is quite small compared to the total number of proper-length Boolean vectors. This principle applies to all of our constructions: verifying an encoding is easiest when the set of proper encodings is "dense" within the set of all Boolean vectors of proper length. That is, if most or all Boolean vectors are proper encodings, little to no verification is necessary. In

hindsight, this justifies the fact that verifying $\mathsf{en}_{\mathsf{VAR}}$ (and similarly linReg) requires so much additional machinery, since the density of valid encodings in that case is $\frac{2^l}{2^{3l}} = \frac{1}{2^l}$.

In the case of FRQ, the proportion of valid encodings is very low, $\frac{l}{2^l}$. Furthermore, there does not seem to be any straightforward way to verify that a shared vector is a single impulse using only addition and multiplication of shares. Thus, our servers will have a bit more verification work to do. It would be desirable to construct some alternative encoding for computing frequency count which still permits computation but whose density of proper encodings in the set of Boolean vectors is higher. Despite these challenges, we give a novel and practically efficient technique for verifying this encoding in the next section.

CHAPTER 8

Resolving Frequency Count

In this section we give a unique method for servers S_L, S_R, given $[x]^B \in \mathbf{Z}_{2^l}$, to verify that x is an impulse. In particular, we detail how to do this efficiently in parallel for $[x_1]^B, \ldots, [x_n]^B$. When we say x is an impulse, we mean $(x)_j = 1 \iff j = j^*$ for some unique $j^* \in \{0, \ldots, k-1\}$. Crucially, they must do this without leaking anything about honest players' input x_i besides whether or not it is an impulse.

Intuitively, we break this process down into two parts:

1. Make sure no player submitted a zero vector.
2. Make sure no player submitted multiple impulses by verifying that the sum of components in the final histogram equals n', the number of valid shares received.

We can perform the first check on all n inputs using just $2n$ bits of communication and $n \cdot l$ XOR operations. This is accomplished by checking the parity of the number of ones in each vector and throwing out any vector with an even number of ones. Servers can perform this parity check in the clear since it will never reveal any information about an honest player's input besides the fact that there are an odd number of ones, which is already implied by the fact that they submitted an honest impulse vector. Thus, each server simply computes the parity of their share and they combine these parities in the clear to get the parity of the client's encoded input. This uses exactly two bits of communication and l XOR operations per server for a single parity check. More precisely:

- S_L, S_R initially hold $[x]_L^B$, $[x]_R^B$ respectively.
- S_L computes $b_L = \bigoplus_j ([x]_L^B)_j$ and sends b_L to S_R.
- S_R computes $b_R = \bigoplus_j ([x]_R^B)_j$ and sends b_R to S_L.
- Both servers compute $b = b_L \oplus b_R$.
- If $b = 0$, both servers discard $[x]^B$.

Now, assume the first condition holds. Since no player submitted a zero vector, we must now detect players who submitted "bad" vectors with multiple ones in them. We know from the first condition that if the sum of the components in the final histogram, which we call $\mathsf{sum}(\vec{h})$, is precisely n', the number of players who submitted non-zero vectors, we can conclude that no player submitted a bad vector of multiple ones.

We would ideally like to allow servers to simply sum the components of their local shares, producing a single arithmetic share of $\mathsf{sum}(\vec{h})$. Then, they could reconstruct this value and compare it to n'. This would leak no information about honest players' inputs, since it is already known that honest players submitted a single impulse, and this sum reveals nothing about the location of any particular impulse, only the total number of such impulses. Since components of the input vectors are originally shared using the Boolean scheme, however, we must first convert each component into the arithmetic secret-sharing scheme using $\mathsf{B2A}$. Once we make this conversion, servers can sum these vectors of Arithmetic shares locally to produce arithmetic shares of $\mathsf{sum}(\vec{h})$, reconstruct this in the clear, and compare it to n', as detailed above.

In the best case, if this check succeeds the first time we attempt it, we know with certainty that no player submitted a non-impulse vector (conditioned on the fact that no player submitted a zero vector, as confirmed by the initial verification step). If it fails, we split the

set of behaving players in half and recursively repeat this check on each half, locating misbehaving players via binary search. In the worst case, this brings the total communication to $2nl \cdot \log(n)$ bits of communication and $O(n \log n)$ local additions/multiplications.

As we mentioned, the rest of the verification is simple: servers compute $s_L = \sum_i ([x]^*_L)_i$ and $s_R = \sum_i ([x]^*_R)_i$. We will then have that $s_L + s_R = \mathsf{sum}(\vec{h})$. To compare this value to n', S_L sends s_L to S_R and S_R computes $s_L + s_R - n'$. If this value is zero, then all inputs are single impulses and S_R replies with the bit 1. If not, S_R replies with the bit 0, indicating that some input has failed, and they continue the recursive search for misbehaving clients.

Once both verification steps succeed, we know that every player submitted a valid impulse except with negligible probability. From this point, servers proceed to compute the frequency count by locally summing their valid vectors of arithmetic shares into local accumulators $A_L, A_R \in \mathbf{Z}^k_{2^l}$. If each such vector is valid, then $A_L + A_R = \vec{h}$, the desired histogram. Servers then return A_L, A_R to clients who reconstruct $\vec{h}$ in precisely this manner.

CHAPTER 9

Share Conversion

In this section, we detail our daBit-based Boolean to arithmetic share conversion protocol. The core of the semi-honest share conversion is to efficiently convert Boolean shares of a single bit $[b]^B$ to arithmetic shares $[b]^A$, i.e. $b = [b]^B_L \oplus [b]^B_R = [b]^A_L + [b]^A_R$ for secret bit b. Then for an arbitrary λ-bit value x shared using Boolean shares, we can convert the shares to arithmetic in parallel and combine them. Namely, given $[x]^B = ([x_0]^B, \ldots, [x_{\lambda-1}]^B)$ where $x_i \in \{0, 1\}$ is the i^{th} bit of x, we have $[x]^A = \sum_{i=0}^{\lambda-1} 2^i [x_i]^A$, noting that $x = \sum_{i=0}^{\lambda-1} 2^i([x_i]^B_L \oplus [x_i]^B_R) = \sum_{i=0}^{\lambda-1} 2^i([x_i]^A_L + [x_i]^A_R)$. This means that converting a λ-bit Boolean share requires L parallel conversions of a single bit, which requires the same number of rounds. This fact is used by both the OT-based protocol of [DSZ15], and the daBit based protocol here in Prio+. Recall that the OT-based protocol of [DSZ15] uses OT in the online phase to accomplish the same goal of semi-honest Boolean to arithmetic share conversion.

A daBit is merely a shared correlated pair $([b]^B, [b]^A)$ for some random bit b, where the Boolean share is a single bit, and each server has one share of $[b]^B$ and one share of $[b]^A$. To convert some single bit Boolean share $[x_i]^B$ to $[x_i]^A$, the servers compute their respective shares of $[x_i]^B \oplus [b]^B$ and swap them to get $v = x_i \oplus b$ in the clear. Then, they locally compute $[x_i]^A_L = v + [b]^A_L - 2v[b]^A_L$ and $[x_i]^A_R = [b]^A_R - 2v[b]^A_R$. This requires communication of only a single Boolean value v. If both players behave honestly, we get $[x_i]^A_L + [x_i]^A_R = v + b - 2vb = v \oplus b$ since v and b are single bits, and $v \oplus b = x_i$ by definition of v. To convert a λ-bit integer x, servers convert each bit x_i in parallel and then locally compute $[x]^A = \sum_{i=0}^{\lambda-1} 2^i [x_i]^A \pmod p$ to get arithmetic shares of x. Because we are working in the semi-honest case, as

these are used only by the semi-honest servers, this is more efficient than in [RW19], where they needed to use an arithmetic Beaver triple to generate each daBit in the malicious setting.

The servers are able to generate each daBit in parallel offline using a single OT each. To convert a single bit, the daBit B2A share conversion only needs to communicate a single bit in the online phase (and consume a daBit), while OT share conversion in [DSZ15] requires an online OT (OT where inputs depend on client data). Hence, our daBits share conversion is much faster in the online phase. End to end, daBits B2A requires an OT to generate a daBit in the offline phase to precompute the correlated daBit, so end-to-end daBit share conversion including the offline phase only requires a single bit more than the OT protocol in [DSZ15]. Formal protocols for daBit generation and share conversion are given below.

$\mathsf{daBitGen}_p$

Inputs: one OT

Output: A random daBit $([b]^A, [b]^B)$ per server.

1. **Sample** Both servers $i \in \{L, R\}$ samples a random bit $b_i \in \{0, 1\}$. S_L also samples a random integer $x \bmod p$.

2. **Use OT**

   - S_L acts as the OT sender, sending $(x, x + b_L)$. S_L also sets $y_L = -x \pmod p$.
   - S_R acts as the OT receiver, using b_R as the choice bit. S_R receives $y_R = x + b_L b_R \pmod p$

3. **Compute**

   (a) Both servers set $[b]_i^A = b_i - 2y_i \pmod p$.

   (b) They also set $[b]_i^B = b_i$.

4. **Output** Server S_i outputs $([b]_i^A, [b]_i^B)$.

$\mathsf{B2A}_p$

Inputs: Boolean shares of a single bit $[x]_L^B, [x]_R^B \in \mathbf{Z}_2$. A single dabit $([b]^A, [b]^B)$

Output: Arithmetic shares of the same bit $[x]_L^A, [x]_R^A \in \mathbf{Z}_p$.

1. **Compute** $v = x \oplus b$

   (a) Both servers $i \in \{L, R\}$ compute $[v]_i^B = [x]_i^B \oplus [b]_i^B$.

   (b) Servers send their share $[v]_i^B$ to each other.

   (c) Servers now have $v = x \oplus b = [v]_L^B \oplus [v]_R^B$ in the clear.

2. **Convert**

   (a) Since v is in the clear, specifically server L computes $[x]_L^A = v + [b]_L^A - 2v[b]_L^A \pmod p$, and server R computes $[x]_R^A = [b]_R^A - 2v[b]_R^A \pmod p$.

3. **Output** Server S_i outputs $[x]_i^A$.

CHAPTER 10

Security

In this section we briefly describe the security properties of the protocols in our system. For formal statements and proofs, see Appendix C.

Let $0 \leq c^* \leq n$ be the number of corrupted clients who submit invalid input shares. For every protocol, up to n malicious players colluding with one semi-honest server learn nothing but the output except with negligible probability, as well as some modest leakage in some cases based on the specific statistic and/or AFE construction. In particular, to compute MEAN servers must leak the number of players $n - c^*$ whose inputs were included in the aggregate. Servers must also give this value when computing linReg. Due to the specific AFE construction for VAR, the output necessarily leaks $\mathbf{E}[X]$ in addition to the variance. All of these modest leakages are analogous to the results of [CB17].

In terms of robustness, all protocols provably prevent any coalition of up to n malicious players from corrupting the output beyond misreporting their private values, except with negligible probability. All statements given above are true, as proven in Appendix C, except with negligible probability.

CHAPTER 11

Practical Evaluation

In this section we describe the practical performance of our implementation of Prio+.

We implemented our scheme in 7,000 lines of C/C++. We utilized the libOTe toolkit [Rin] for OT and silent OT-extension. Our scheme uses semi-honest OTs, since our servers are assumed to be semi-honest. Similar to Prio, clients use NaCl's "box" primitive to encrypt and sign messages. This means that TLS is not required to secure client-server communication.

Our implementation supports secure computation of SUM, AND, OR, MAX, MIN, VAR, FRQ and linReg. Two implementations of the original Prio protocol exist: the original implementation, written in Go, supports secure computation of SUM, AND, OR, MAX, MIN, and linReg. Since the original paper's publication, another implementation was written by Mozilla in C. The Mozilla implementation only supports SUM.

We provide comparison data for three statistics: SUM, MAX, and linReg. These represent our three categories of protocols: SUM requires share conversion but no SNIPs, MAX requires neither share conversion nor SNIPs, and linReg requires both share conversion and SNIPs.

We collected four types of data for comparison: client encode time (milliseconds per client), client message size (bytes per client per server), server compute time (milliseconds

per server per client), and server communication (bytes per server per client). Most data we collected shows a direct linear relationship with the total number of clients, and so we ran multiple trials with 10k, 50k, and 100k clients and then averaged the result. The only exception to this is server compute time for VAR and linReg. Since they utilize SNIPs which are easily batched, server compute time increases sublinearly with number of clients. In that case only, we performed trials with 50k clients and averaged those results. In our end-to-end implementation we used the exact same servers as the original Prio paper, two c4.2x large AWS servers. All protocols were implemented using just two servers. Both servers are located in the us-east-1f zone as to mimic a low-latency, high bandwidth connection. The client code was run from a separate instance of the same c4.2x large AWS server, and all client data was randomly generated. All three implementations (Go, Mozilla, Prio+) use this same 2-server setup.

11.1 Data: SUM

For SUM, we compared Prio+ to both the original Go implementation and the Mozilla implementation. We ran four separate experiments with clients holding 1-bit, 8-bit, 16-bit, and 32-bit integers. Our results can be found in Figures 11.1, 11.2, 11.3, and 11.4. We also measured end-to-end runtime of our system. In total, our Prio+ implementation computes the sum of 100,000 16-bit integers in 0.47 seconds. This excludes offline pre-computation time but includes client encode time, communication time, and server compute time.

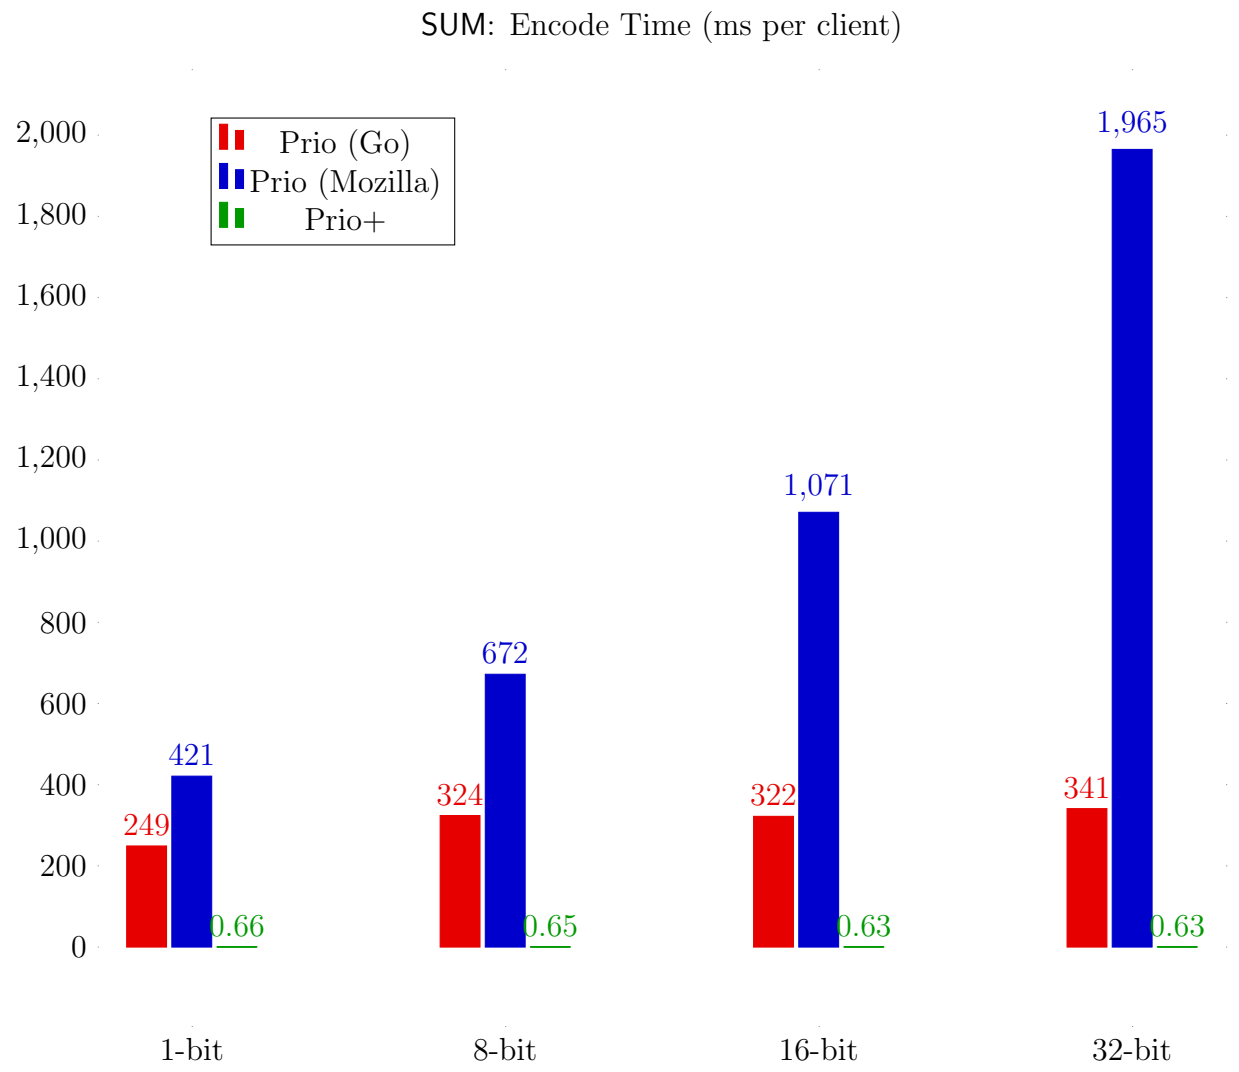

Figure 11.1: This chart shows the time necessary for a single client P_i holding private value x_i to encode that value, compute any necessary additional proofs, and secret-share both the encoding and the proof(s) when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in milliseconds and data is arranged according to the number of bits in x_i.

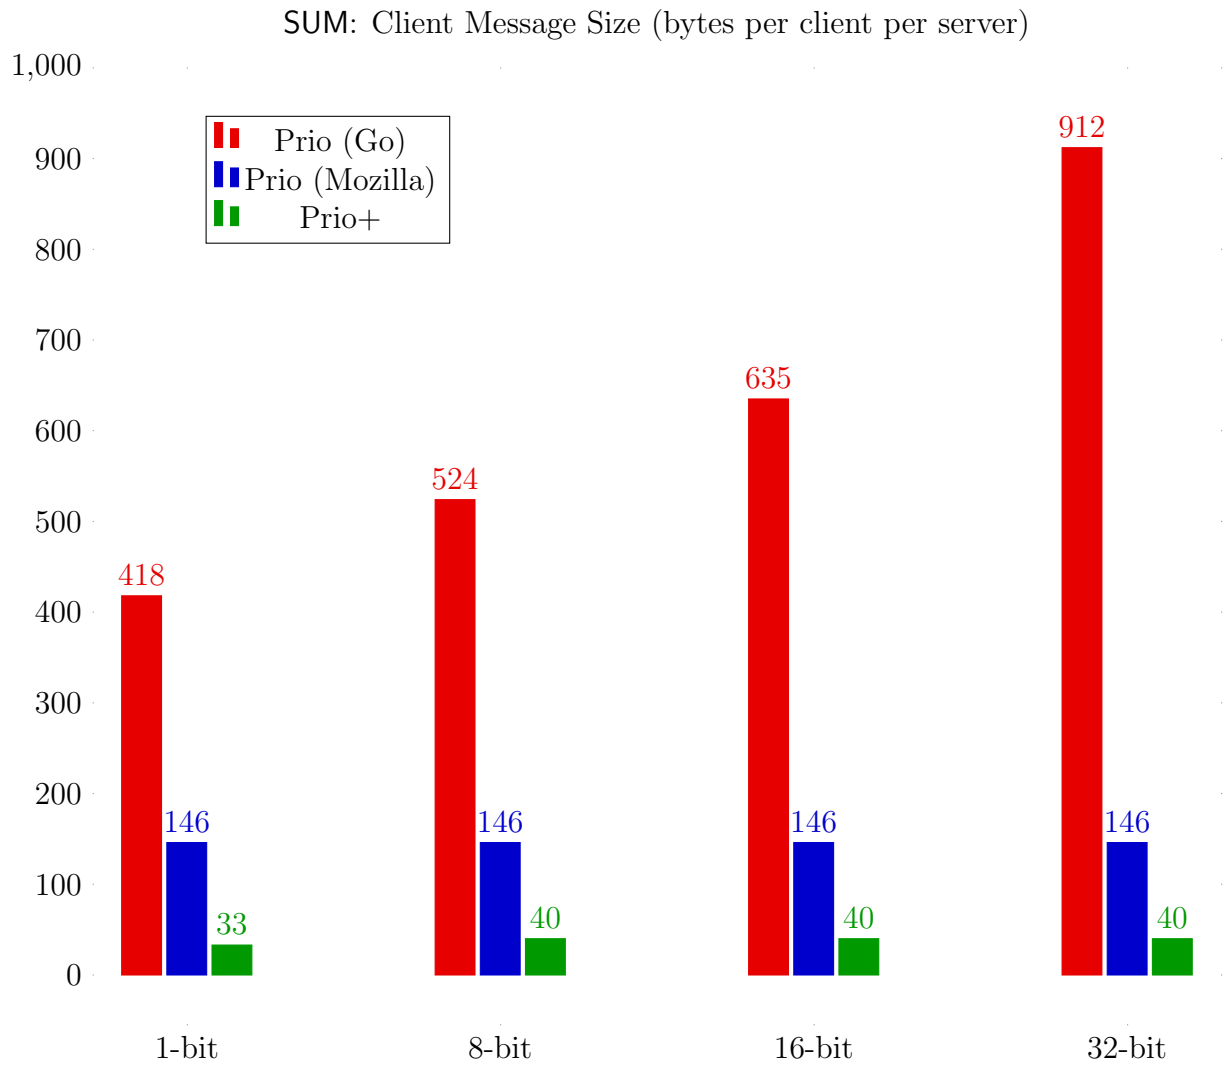

Figure 11.2: This chart shows the size of the message P_i (holding private value x_i) sends to each server when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in bytes and data is arranged according to the number of bits in x_i.

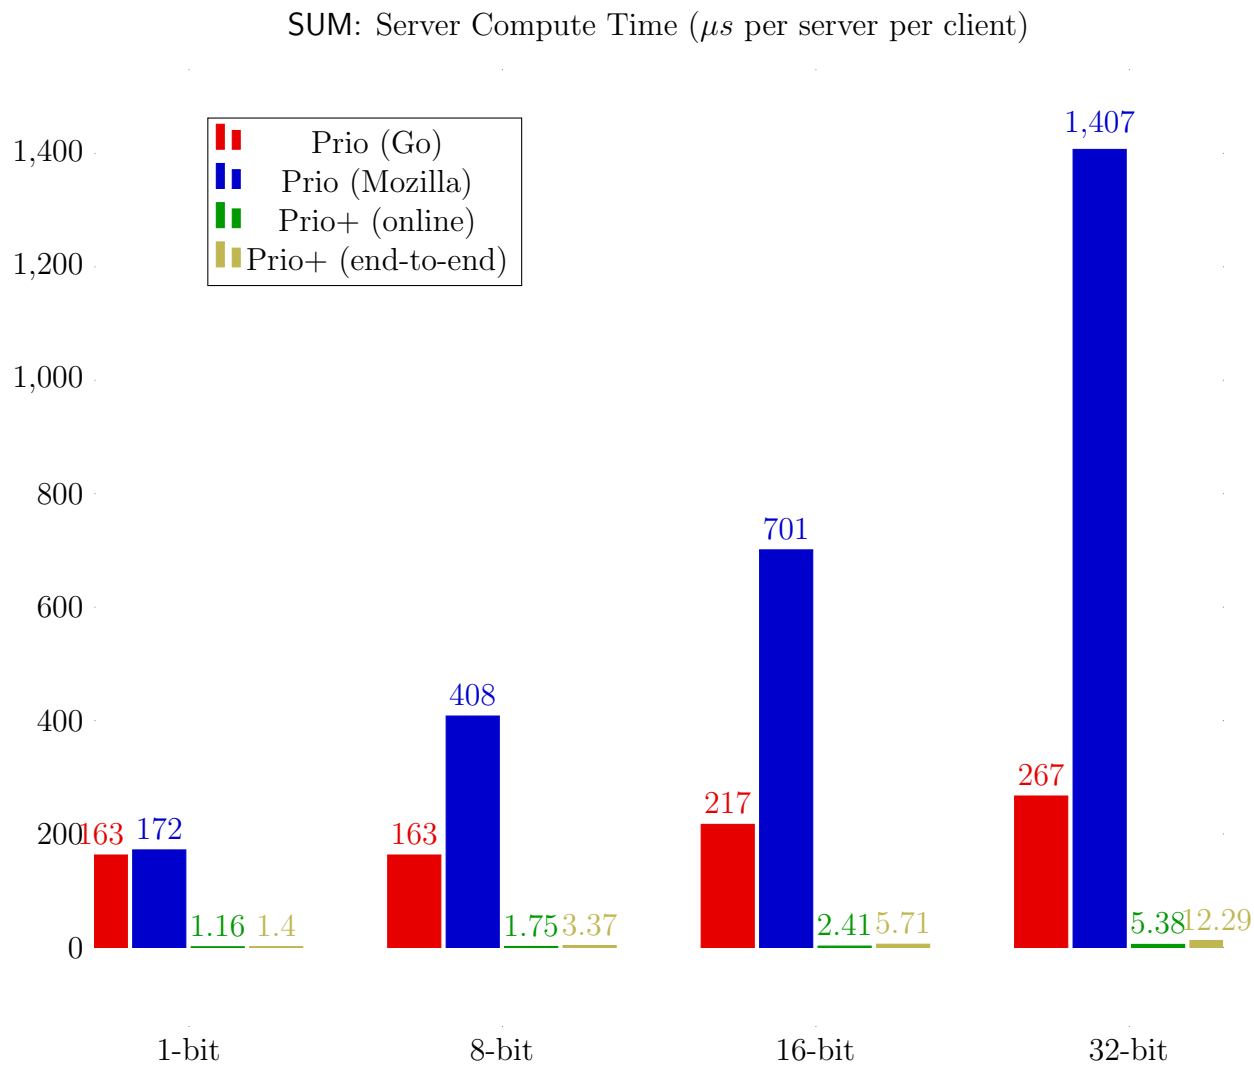

Figure 11.3: This chart shows the computation time for each server when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in microseconds and data is arranged according to the number of bits in the client's private value x_i. Prio+ (end-to-end) includes time to pre-compute daBits for share conversion.

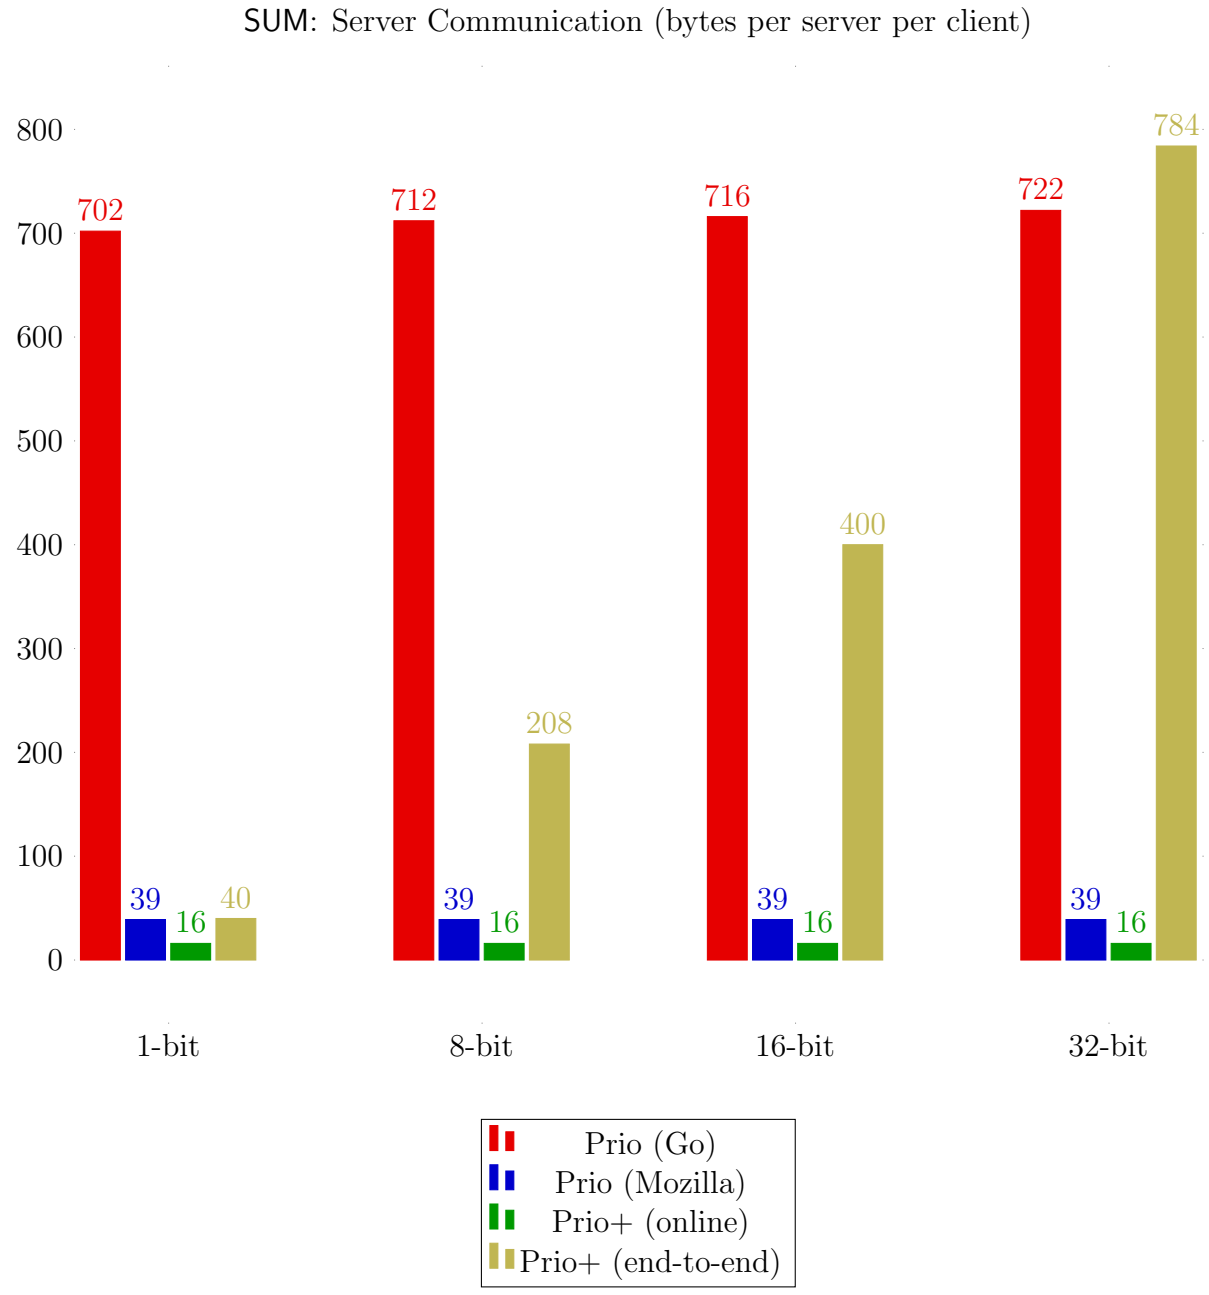

Figure 11.4: This chart shows the communication required by each server when executing a protocol to compute $\mathsf{SUM}(x_1, \ldots, x_n)$. Results are in bytes and data is arranged according to the number of bits in x_i. Prio+ (end-to-end) includes time to pre-compute daBits for share conversion.

Analysis: As expected, our implementation of Prio+ overwhelmingly outperforms both implementations of Prio in terms of client encode time and client message size. An additional pleasant result is that Prio+ heavily outperforms the Prio implementations in terms of server compute time as well. We expected to see performance drawbacks in server communication, since our server communication is no longer constant with respect to input size. Prio+ has comparable server communication to the Mozilla implementation for summing single-bit integers, but as the size of the integers increases, so does the communication. The Go implementation has high, but still constant, server communication. This is due to it being designed for a larger number of servers, whereas our experiments were run on a two-server implementation. Prio+ servers still communicate less than those in the original Go implementation except in the 32-bit case. The Mozilla implementation achieves lower server communication than Prio+, and this is the expected drawback of using share conversion as a replacement for SNIPs, which were designed to minimize server communication only.

11.2 Data: MAX

Since the Mozilla implementation does not support MAX, we compared Prio+ to the original Go implementation only. In our experiment, clients held integers in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$. In total, our Prio+ implementation computes the maximum of 100,000 4-bit integers in 5.25 seconds. This excludes offline pre-computation time but includes client encode time, communication time, and server compute time.

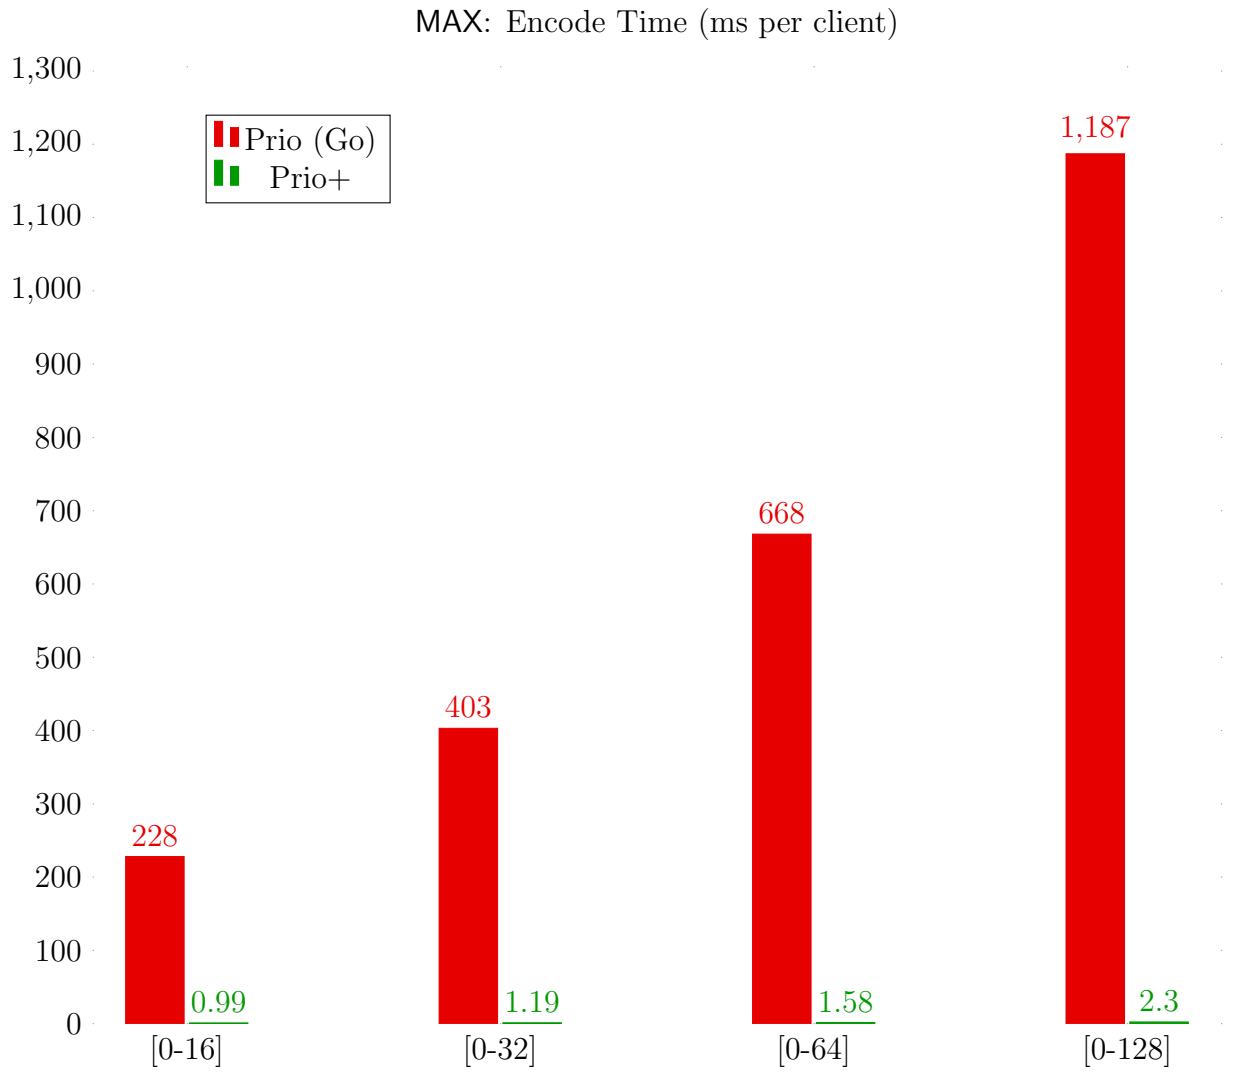

Figure 11.5: This chart shows the time necessary for a single client P_i holding private value x_i in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$ to encode that value, compute any necessary additional proofs, and secret-share both the encoding and the proof(s) when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$. Results are in milliseconds.

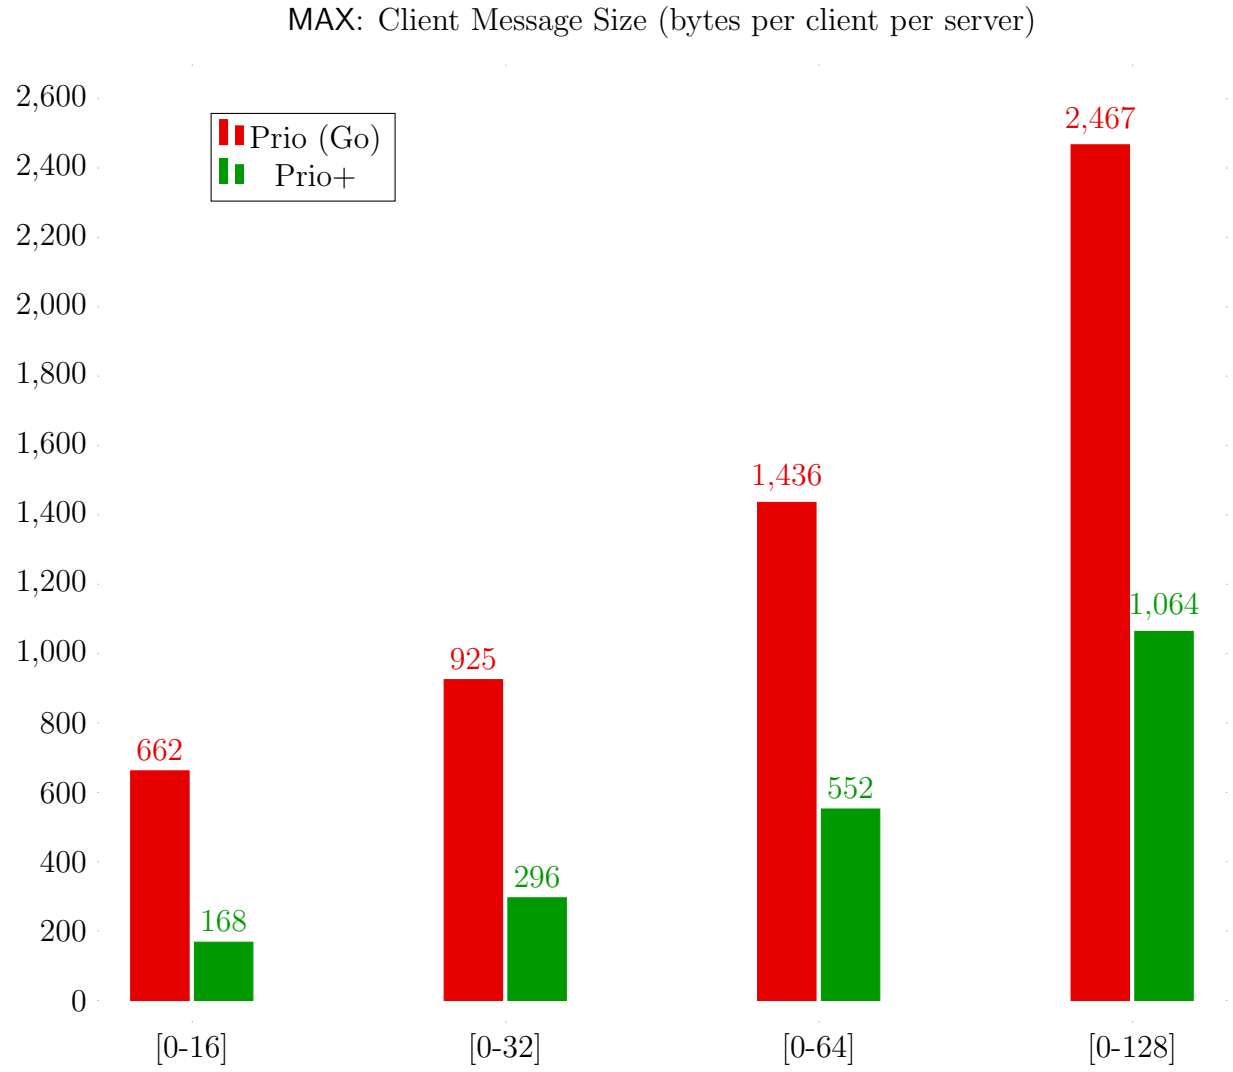

Figure 11.6: This chart shows the size of the message P_i (holding private value x_i in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$) sends to each server when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$. Results are in bytes.

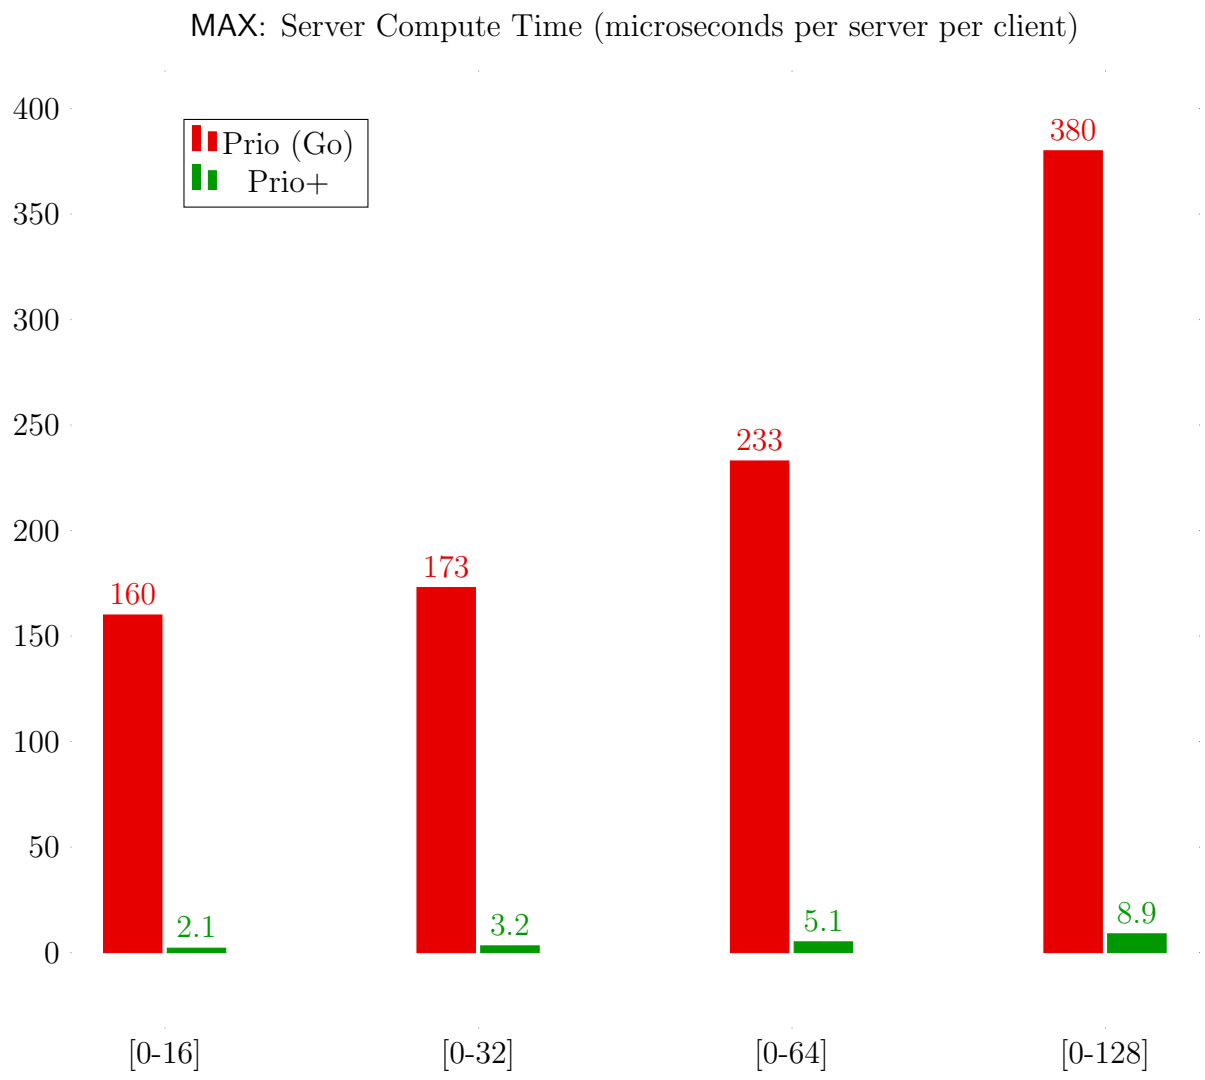

Figure 11.7: This chart shows the average computation time per server when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$, where each x_i held by P_i lies in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$. Results are in milliseconds.

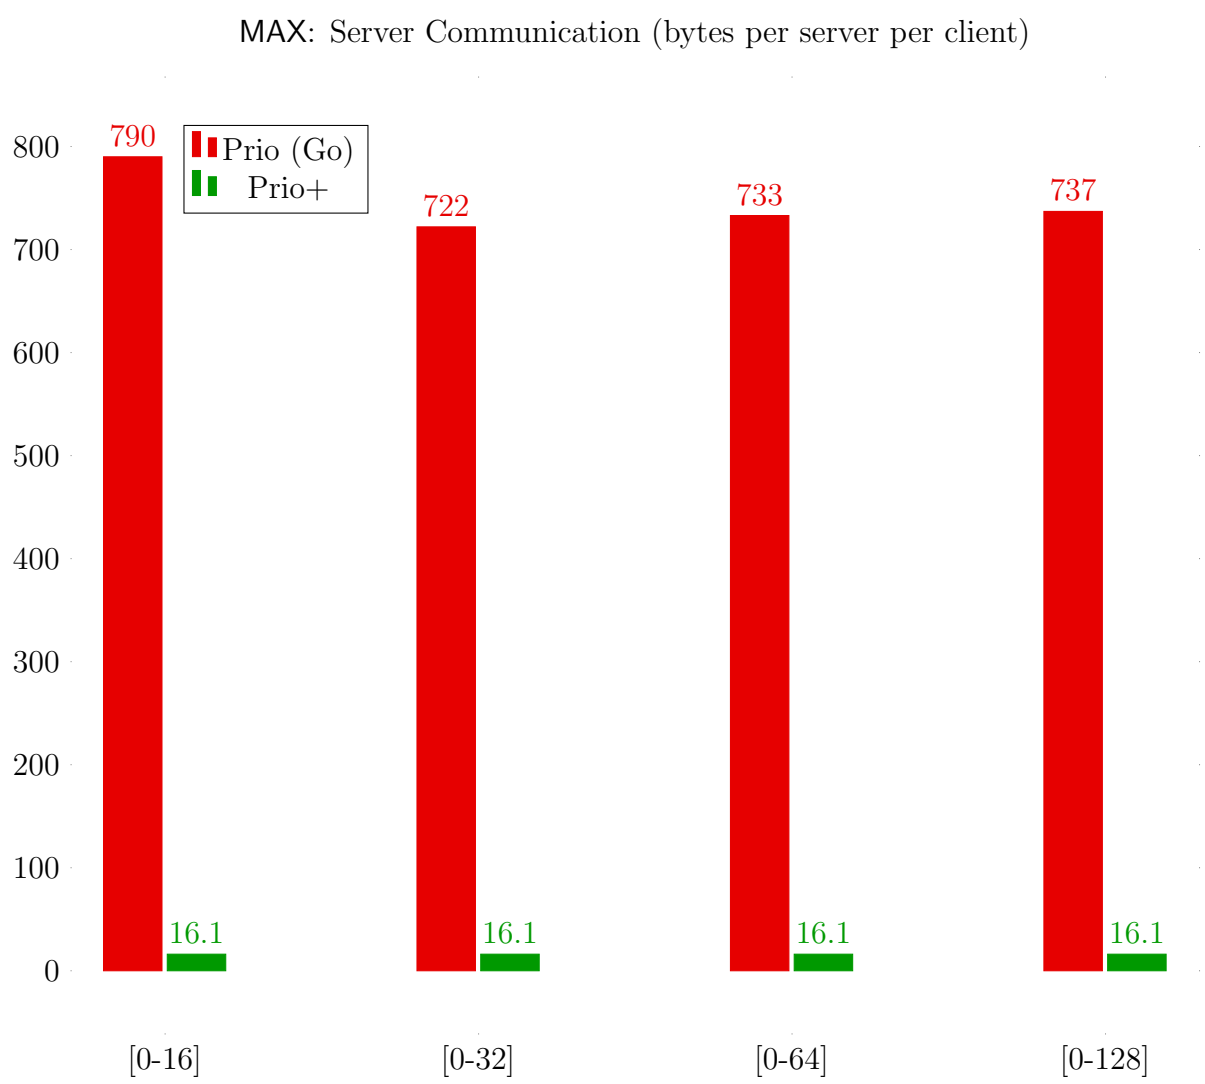

Figure 11.8: This chart shows the average bytes communicated by each server when executing a protocol to compute $\mathsf{MAX}(x_1, \ldots, x_n)$, where each x_i held by P_i is in the range $[0, x]$ for $x \in \{16, 32, 64, 128\}$. Results are in bytes.

Analysis: In the case of MAX, Prio+ requires no share conversion and no SNIPs, resulting in dramatic performance benefits for both clients and servers. Particularly relevant is the nearly 1,000x decrease in client encode time. Since servers do not have to perform share conversion, we get an added benefit of sharply decreased server communication.

11.3 Data: linReg

Similar to MAX, linReg is only supported by Prio+ and the original Go implementation. We performed four separate experiments in which the feature vectors held by clients are of degree 2, 4, 6, and 8 respectively. All results are averaged between experiments with 10k, 50k and 100k clients except for server compute time where all experiments were performed with 50k clients. In total, our Prio+ implementation computes a line-of-best-fit over 100,000 degree-2 vectors of 8-bit integers in 7.41 seconds. This excludes offline pre-computation time but includes client encode time, communication time, and server compute time.

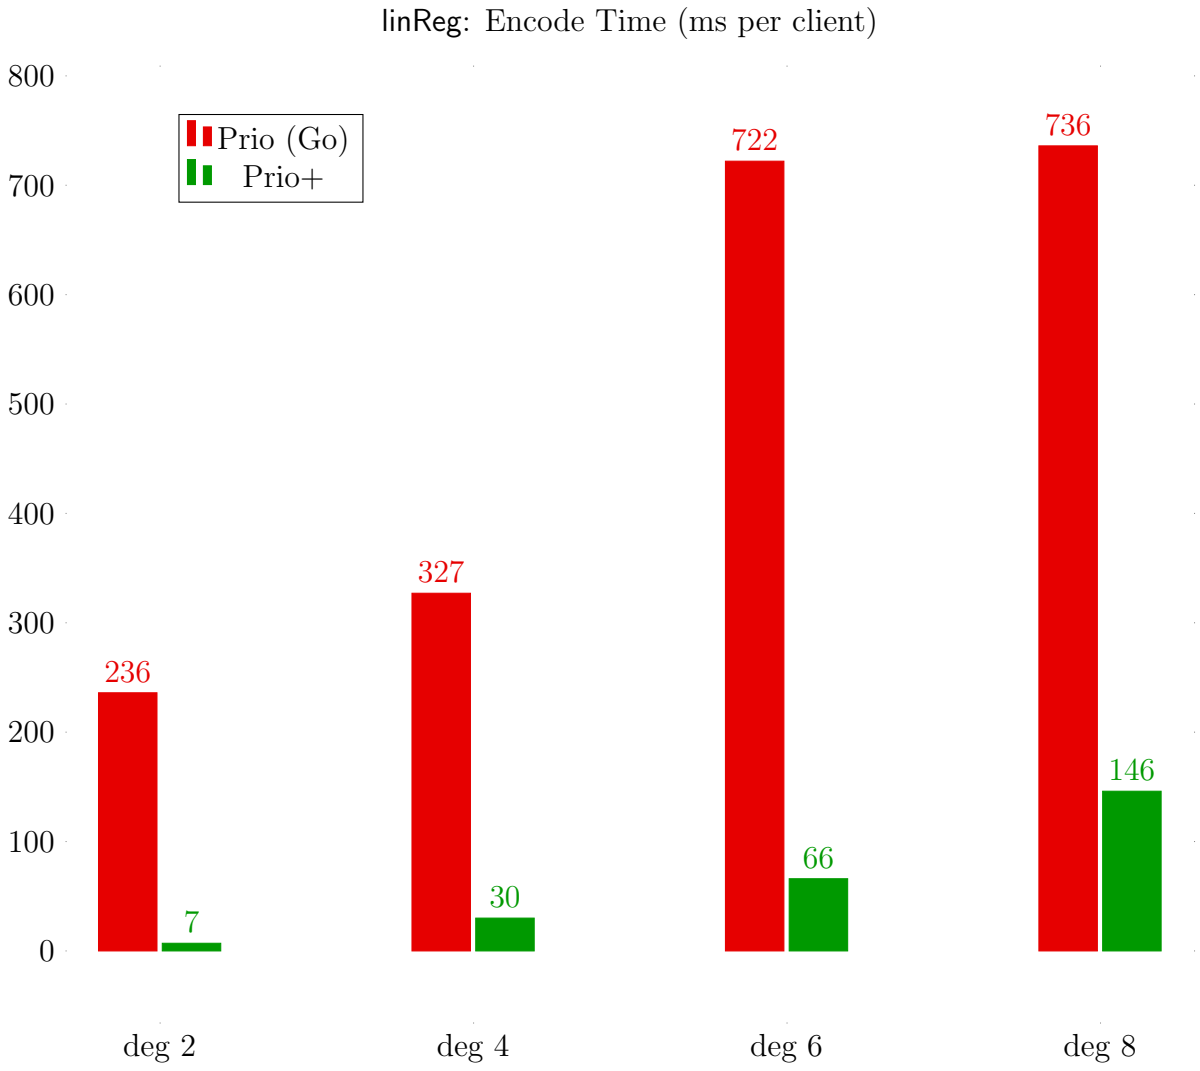

Figure 11.9: This chart shows the time necessary for a single client P_i holding private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$ to encode that value, compute any necessary additional proofs, and secret-share both the encoding and the proof(s) when executing a protocol to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$. Each $x_i^{(k)}$ is an 8-bit integer. Results are in milliseconds and data is arranged according to the degree d of each $\vec{x}_i$.

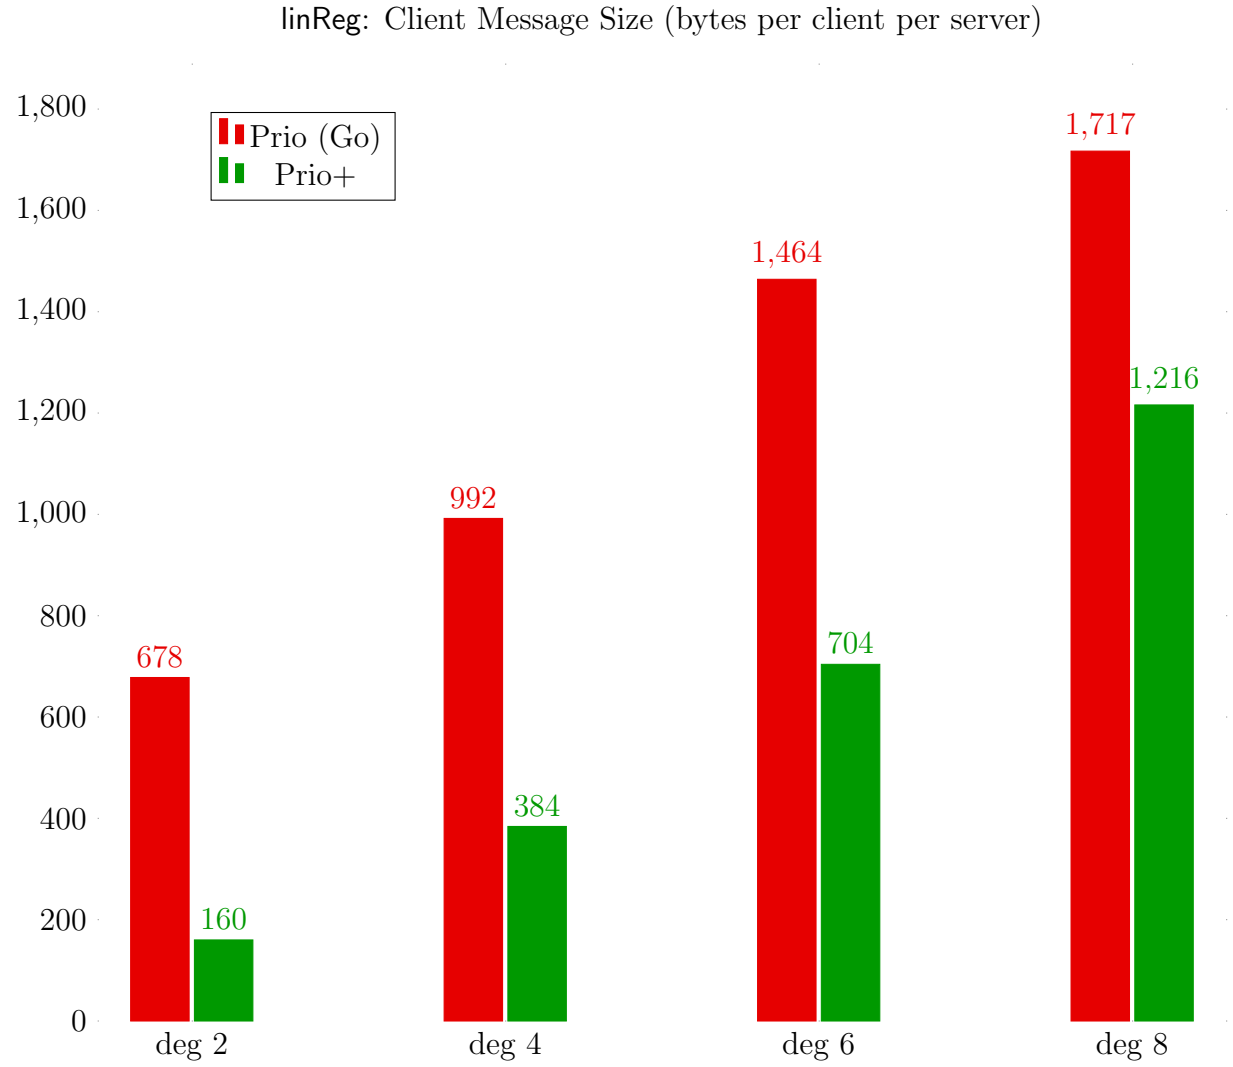

Figure 11.10: This chart shows the size of the message sent by client P_i (holding private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$) to each server when executing a protocol to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$. Each $x_i^{(k)}$ is an 8-bit integer. Results are in milliseconds and data is arranged according to the degree d of each $\vec{x}_i$.

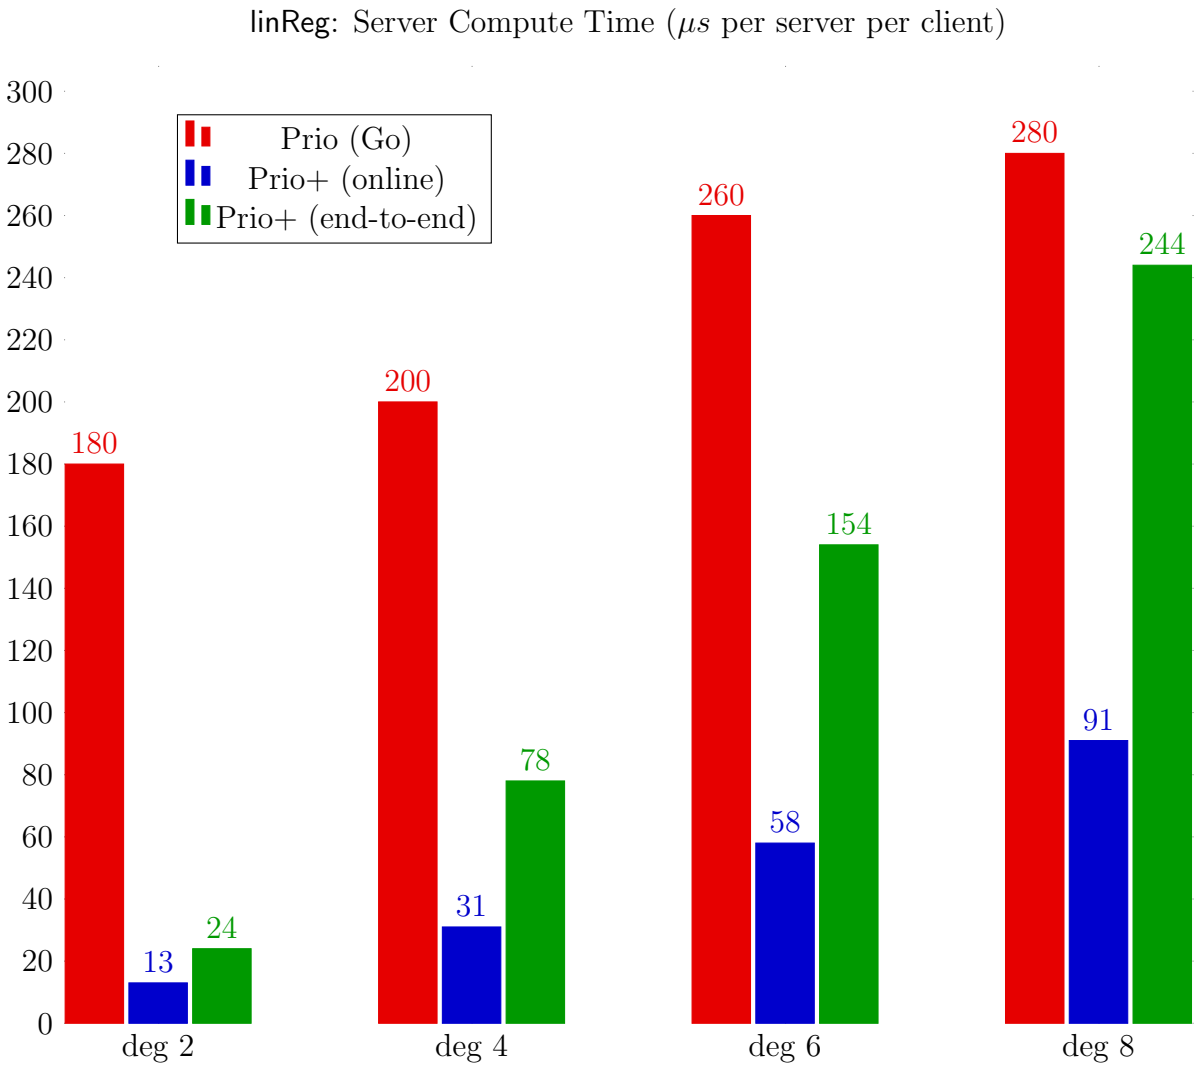

Figure 11.11: This chart shows the total server time to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$, where each of the 50,000 clients P_i holds private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$ and each $x_i^{(k)}$ is an 8-bit integer. Results are in seconds and data is arranged according to the degree d of each $\vec{x}_i$. Prio+ (end-to-end) includes time to pre-compute daBits for share conversion.

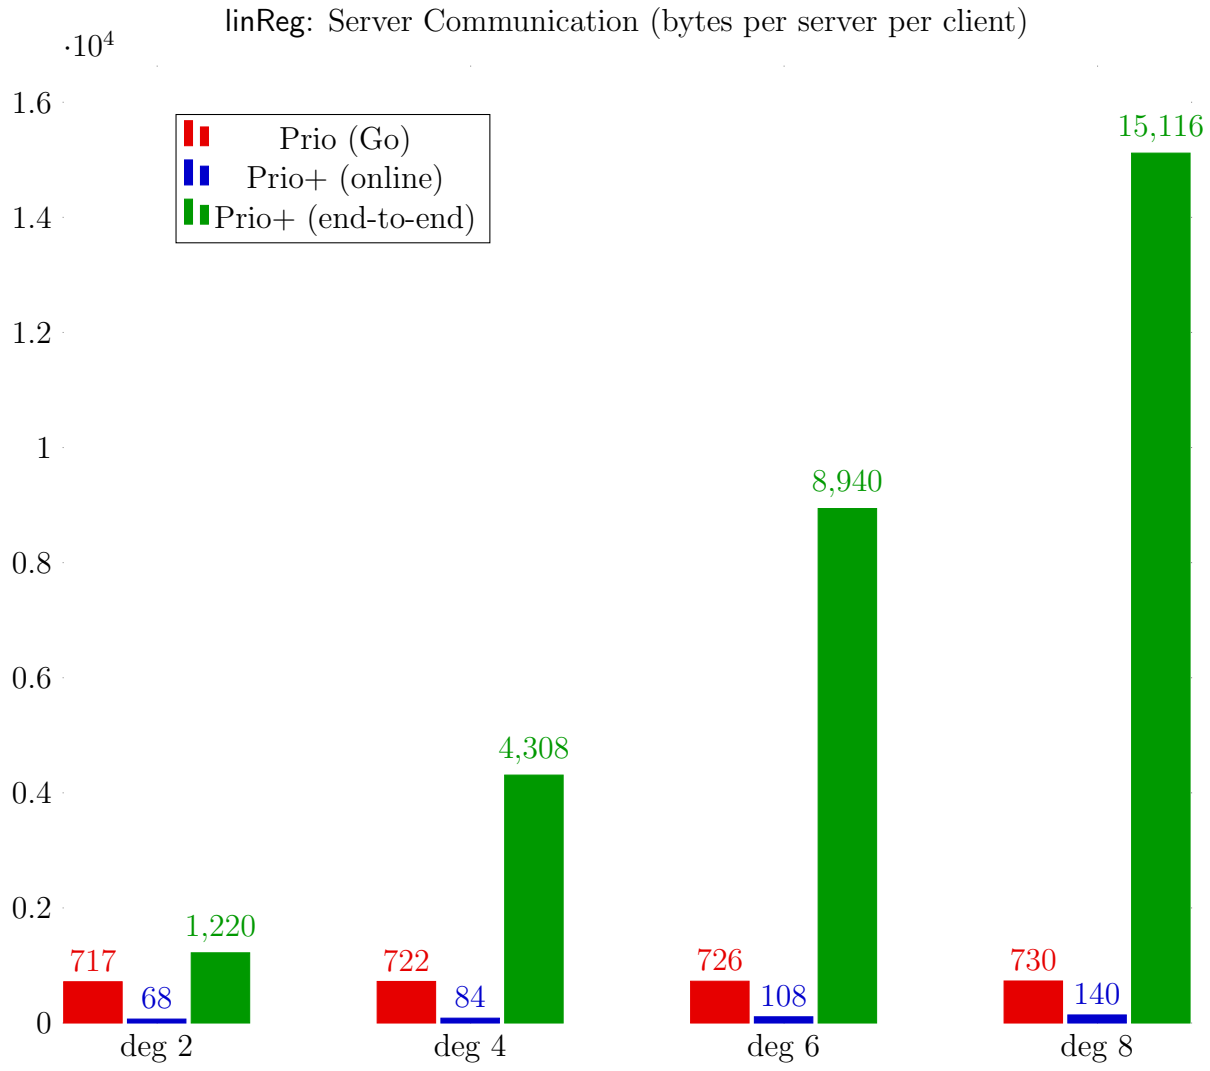

Figure 11.12: This chart shows the average bytes communicated by each server when executing a protocol to compute $\mathsf{linReg}(\vec{x}_1, \ldots, \vec{x}_n)$, where client P_i holds private value $\vec{x}_i = (x_i^{(0)}, \ldots, x_i^{(d)})$ and each $x_i^{(k)}$ is an 8-bit integer. Results are in bytes and data is arranged according to the degree d of each $\vec{x}_i$. Prio+ (end-to-end) includes time to precompute daBits for share conversion.

Analysis: Prio+ provides significant benefits to the client in computing linReg at a minimal server-side cost. Client encode time and message size are dramatically decreased, although the difference in client message size decreases as the degree increases. This is because as the degree increases, bit-length verification represents a smaller proportion of the multiplication gates in the VALID circuit. Once again we have an added benefit of decreased server computation, particularly for lower degree regression. Based on this trend, we expect the server computation in Prio+ to exceed that of Prio beyond degree 8 linear regression. We see a similar trend in terms of server communication, where Prio+ servers communicate less for lower degree computation, but that communication grows with the degree whereas Prio's server communication remains constant.

11.4 Data: Offline Pre-computation

Prio+ requires pre-computed data which is independent of client data. In our semi-honest setting, we are able to use a single OT to generate a daBit, leading to very efficient generation. On average, our daBit generation can produce around 4,000,000 daBits per second. Thus it takes, for example, 1.25 seconds to compute enough daBits to perform 8-bit degree 2 linear regression of 100,000 submissions.

CHAPTER 12

Conclusions and Future Work

By leveraging properties of Boolean secret-sharing, Prio+ effectively restructures the original Prio protocol to privately compute aggregate statistics with minimal burden on the client. Often the computational burden on the servers is also reduced, occasionally at the cost of increased server communication. The scale of these performance improvements depends heavily on the statistic being computed. For some statistics ($\mathsf{AND}, \mathsf{OR}, \mathsf{MAX}, \mathsf{MIN}$), Prio+ significantly reduces the burden on both clients and servers across the board. For others, particularly linear regression, the improvements in client performance and server computation come with a moderate increase in server bandwidth usage. These costs are only apparent, however, as the size of client inputs becomes large. For small and often practical client values, Prio+ still outperforms Prio across all examined metrics. Prio+ does require an additional offline pre-computation phase which enables efficient share conversion. However, this computation is relatively inexpensive and can be done during times when data is not being collected and servers are otherwise idle.

In the future, we wish to expand Prio+ to compute additional aggregate statistics. We would particularly like to identify statistics similar to $\mathsf{AND}, \mathsf{OR}, \mathsf{MAX}$, and MIN for which neither SNIPs nor share conversion are required. This would ideally involve a formal classification of the set of statistics for which AFE's exist (especially AFE's where aggregation is done via bitwise XOR). We are also interested in generalizing the process of computing new statistics, so that a novel AFE is not necessary each time a new statistic needs to be computed.

APPENDIX A

Full Protocol Descriptions

In this section we give detailed descriptions of the protocols which up until now have been described at a more informal level. We omit the protocol for MIN as it consists of a single call to MAX plus a single pre-processing step and a single post-processing step. We also omit the protocols for integer mean and standard deviation, as these too require a single post-processing step on top of the integer sum and variance protocols respectively.

Π_{sum}

Inputs: $x_i \in \mathbf{Z}_{2^{l'}}$ for $i \in [n]$.
Output: $\sum_i x_i$.

1. **Upload**:
   (a) Each client P_i computes $\mathsf{Share}_{\oplus,l'}(x_i) \longrightarrow [x_i]^B_L, [x_i]^B_R$ via Definition 4.2
   (b) Each P_i sends $[x_i]_L, [x_i]_R$ to S_L, S_R respectively.
2. **Verify Bit-Length**: Initially, $n' = n$. If a server receives a share which is not l bits in length from P_i (assume S_L w.l.o.g.):
   (a) S_L sends the index i to S_R.
   (b) Both servers discard $[x_i]^B$.

(c) Both servers set $n' \leftarrow n' - 1$

3. **Convert Shares**: S_L and S_R jointly evaluate $\mathsf{B2A}_{l',2^l}(\{[x_i]_L^B, [x_i]_R^B\})$ on each of the n' valid pairs of Boolean shares as described in [DSZ15]. S_L receives as output $\{[x_i]_L^A\}_i$ and S_R receives as output $\{[x_i]_R^A\}_i$.

4. **Aggregate**: S_L locally adds all arithmetic shares into an accumulator A_L, initially zero. That is: $A_L \longleftarrow A_L + \sum_i [x_i]_L^A$. S_R analogously accumulates its arithmetic shares into $A_R \longleftarrow A_R + \sum_i [x_i]_R^A$.

5. **Publish**: Once all n' shares have been accumulated, S_L and S_R publish A_L and A_R to every client.

6. **Client Computation**: Clients output $A_L + A_R$.

Π_{and}

Inputs: $x_i \in \{0,1\}$ for $i \in [n]$.

Output: 1 if and only if $x_i = 1$ for all $i \in \{1, \ldots, n\}$.

1. **Upload**: Each P_i encodes their as input as:

   $\hat{x}_i = 0 \in \mathbf{Z}_{2^\lambda}$ if $x_i = 1$

   $\hat{x}_i = r \in \mathbf{Z}_{2^\lambda}$ if $x_i = 0$, where r is uniformly random

   P_i then computes $\mathsf{Share}_{\oplus,\lambda}(\hat{x}_i) = ([\hat{x}_i]^B_{L,\lambda}, \hat{x}_i]^B_{R,\lambda})$ as in Definition 4.2 and sends one share to each server.

2. **Verify Bit-Length**: Initially, $n' = n$. If some server, say S_L, receives from P_i $[x_i]^B_{L,\lambda}$ which is an m-bit integer, $m \neq \lambda$:

   (a) S_L sends the index i to S_R.

   (b) Both servers discard $[x_i]^B_\lambda$ (removing from accumulator if necessary).

   (c) Both servers set $n' \leftarrow n' - 1$

3. **Aggregate**: S_L and S_R hold accumulator values $A_L, A_R \in \mathbf{Z}_{2^\lambda}$, initially set to 0. Once a λ-bit share is sent to S_L by P_i, S_L *immediately* XORs it with A_L: $A_L \leftarrow A_L \oplus [\hat{x}_i]_L$. S_R does the same with its accumulator A_R upon receiving a valid share. If either server learns that a share already accumulated should be discarded, they simply XOR it with their accumulator again.

4. **Publish**: S_L publishes A_L to all clients, S_R publishes A_R to all clients.

5. **Client Computation**: Clients compute $A = A_L \oplus A_R \in \mathbf{Z}_{2^\lambda}$. If $A = 0$, clients output 1. Otherwise, they output 0.

Π_{or}

Inputs: $x_i \in \{0,1\}$ for $i \in [n]$.

Output: 0 if and only if $x_i = 0$ for all $i \in \{1, \ldots, n\}$.

1. **Upload**: Each P_i encodes their as input as:

   $\hat{x}_i = 0 \in \mathbf{Z}_{2^\lambda}$ if $x_i = 0$

   $\hat{x}_i = r \in \mathbf{Z}_{2^\lambda}$ if $x_i = 1$, where r is uniformly random.

   P_i then computes $\mathsf{Share}_{\oplus,\lambda}(\hat{x}_i) = ([\hat{x}_i]^B_{L,\lambda}, \hat{x}_i]^B_{R,\lambda})$ as in Definition 4.2 and sends one share to each server.

2. **Verify Bit-Length**: Initially, $n' = n$. If some server, say S_L, receives from P_i $[x_i]^B_{L,\lambda}$ which is an m-bit integer, $m \neq \lambda$:

   (a) S_L sends the index i to S_R.

   (b) Both servers discard $[x_i]^B_\lambda$ (removing from accumulator if necessary).

   (c) Both servers set $n' \leftarrow n' - 1$

3. **Aggregate**: S_L and S_R hold accumulator values $A_L, A_R \in \mathbf{Z}_{2^\lambda}$, initially set to 0. Once a λ-bit share is sent to S_L by P_i, S_L *immediately* XORs it with A_L: $A_L \leftarrow A_L \oplus [\hat{x}_i]_L$. S_R does the same with its accumulator A_R upon receiving a valid share. If either server learns that a share already accumulated should be discarded, they simply XOR it with their accumulator again.

4. **Publish**: S_L publishes A_L to all clients, S_R publishes A_R to all clients.

5. **Client Computation**: Clients compute $A = A_L \oplus A_R \in \mathbf{Z}_{2^\lambda}$. If $A = 0$, clients output 0. Otherwise, clients output 1.

Π_{max}

Inputs: $x_i \in \{0, \ldots, K-1\}$ for $i \in [n]$.

Output: $\mathsf{max}_i x_i$

1. **Upload**: Each P_i encodes their private x_i as $\hat{x}_i \in \mathbf{Z}_{2^K}^{\lambda}$, where:

   - $(\hat{x}_i)_j = r_j \in \mathbf{Z}_{2^\lambda}$ for $0 \leq j \leq x_i$, where r_j is a uniformly random λ-bit integer.
   - $(\hat{x}_i)_j = 0 \in \mathbf{Z}_{2^\lambda}$ for $x_i < j \leq K-1$

   We will use $(\hat{x}_i)_{j,k}$ to refer to the k'th component of $(\hat{x}_i)_{j,k}$). Each P_i then computes $\mathsf{Share}_{\oplus,\lambda K}(\hat{x}_i) = [\hat{x}_i]^B_{L,\lambda K}, [\hat{x}_i]^B_{R,\lambda K} \in \mathbf{Z}_{2^\lambda}^K$ as described in [DSZ15] and sends one share of $\hat{x}_i$ to each server.

2. **Verify Bit-Length**: Initially, $n' = n$. If some server, say S_L, receives from P_i $[x_i]^B_{L,\lambda K}$ which is an m-bit integer, $m \neq \lambda K$:

   (a) S_L sends the index i to S_R.

   (b) Both servers discard $[x_i]^B_{\lambda K}$ (removing from accumulator if necessary).

   (c) Both servers set $n' \leftarrow n' - 1$

3. **Aggregate**: S_L computes $A_L = \bigoplus_i [\hat{x}_i]^B_{L,\lambda K}$ and S_R computes $A_R = \bigoplus_i [\hat{x}_i]^B_{R,\lambda K}$.

4. **Publish**: S_L and S_R publish their accumulator values A_L and A_R to every client.

5. **Client Computation**: Clients compute

$$((s_1, \ldots, s_\lambda), \ldots, (s_{(K-1)\lambda+1}, \ldots, s_{K\lambda})) := A_L \oplus A_R$$

   Clients output the largest index $i \in \{1, \ldots, K\}$ such that the substring $(s_{(K-i)\lambda+1}, \ldots, s_{K-i+1}\lambda)$ is not identically zero.

Π_{var}

Inputs: $x_i \in \mathbf{Z}_{2^l}$ for $i \in [n]$.

Output: $\mathsf{VAR}(x_1, \ldots, x_n) = \mathbf{E}[X^2] - (\mathbf{E}(X))^2$.

1. **Upload**:

   (a) Each P_i encodes their input as $\hat{x}_i = (x_i, x_i^2) \in \mathbf{Z}_{2^{3l}}$, where x_i is written using l-bits and x_i^2 using $2l$ bits. We will call these two components $\hat{x}_i^{(1)}$ and $\hat{x}_i^{(2)}$ respectively.

   (b) Each P_i then computes $\mathsf{Share}_{\oplus,3l}(\hat{x}_i) = [\hat{x}_i]^B_{L,3l}, [\hat{x}_i]^B_{R,3l} \in \mathbf{Z}_{2^{3l}}$ via 4.2. We refer to the first l bits of these shares as $[\hat{x}_i]^{B,(1)}_l$ and the last $2l$ bits of these shares as $[\hat{x}_i]^{B,(2)}_{2l}$.

   (c) Each P_i computes the circuit $\mathsf{Valid}_{\mathsf{VAR}}(\hat{x}_i)$ and constructs polynomials f, g, h representing the values on the input and output wires of each multiplication gate.

   (d) P_i sends one share of $[\hat{x}_i]^B_{3l}$ to each server, as well as one share of $[f(0)]^B$, one share of $[g(0)]^B$, and one share of $[h]^B$ (that is, one share of each coefficient of h). P_i also computes and sends shares of a triple $([a]^B, [b]^B, [ab]^B)$ for random $a, b \in \mathbf{Z}_{2^{3l}}$.

2. **Verify Bit-Length**: Initially, $n' = n$. If some server, say S_L, receives from P_i $[\hat{x}_i]^B_{L,3l}$ which is an m-bit integer, $m \neq 3l$:

   (a) S_L sends the index i to S_R.

   (b) Both servers discard $[x_i]^B_{3l}$ (removing from accumulator if necessary).

   (c) Both servers set $n' \leftarrow n' - 1$

3. **Convert Shares**: S_L and S_R jointly evaluate $\mathsf{B2A}_{l,p}(\{[\hat{x}_i]^{B,(1)}_{L,l}, [\hat{x}_i]^{B,(2)}_{R,l}\})$ as well as $\mathsf{B2A}_{2l,p}(\{[\hat{x}_i]^{B,(1)}_{L,2l}, [\hat{x}_i]^{B,(2)}_{R,2l}\})$ on each of the n' valid sets of Boolean shares as described

in Section 9. S_L receives as output $\{[\hat{x}_i]_L^{A,(1)}\}$ and $\{[\hat{x}_i]_L^{A,(2)}\}$ and S_R receives as output $\{[\hat{x}_i]_R^{A,(1)}\}$ and $\{[\hat{x}_i]_R^{A,(2)}\}$.

4. **Verify Encoding**: For each valid $\hat{x}_i$ received, servers verify it is truly of the form (x, x^2) as follows:

   (a) Servers compute values on all other wires of the $\mathsf{Valid}_{\mathsf{VAR}}(\hat{x}_i)$ circuit via affine operations on the values they have already (input wires, input/output wires of all multiplication gates). For input shares, servers use arithmetic shares from the output of share conversion.

   (b) Servers use polynomial interpolation on these shares to compute secret-shares of $[f]$ and $[g]$.

   (c) Servers choose a random $r \in \mathbf{Z}_{2^l}$ and each server locally computes $[f(r)]^B, [r \cdot g(r)]^B$, and $[r \cdot h(r)]^B$.

   (d) Servers use the Boolean multiplication triple $([a]^B, [b]^B, [ab]^B)$ to compute $[r \cdot f(r) \cdot g(r)]^B$. From this and $[r \cdot h(r)]^B$, servers use affine operations to compute $[r \cdot (f(r) \cdot g(r) - h(r))]^B$.

   (e) Servers reconstruct the value $r \cdot (f(r) \cdot g(r) - h(r))$. If this value is non-zero, servers reject the input $\hat{x}_i$.

5. **Aggregate**: S_L computes $A_L^{(1)} = \sum_i [\hat{x}_i]_L^{A,(1)}$ and $A_L^{(2)} = \sum_i [\hat{x}_i]_L^{A,(2)}$, and S_R computes $A_R^{(1)} = \sum_i [\hat{x}_i]_R^{A,(1)}$ and $A_R^{(2)} = \sum_i [\hat{x}_i]_R^{A,(2)}$.

6. **Publish**: S_L and S_R publish their accumulator values $A_L^{(1)}$, $A_L^{(2)}$, $A_R^{(1)}$, $A_R^{(2)}$ to every client.

7. **Client Computation**: P_i outputs $\frac{1}{n}(A_L^{(2)} + A_R^{(2)}) - \frac{1}{n^2}(A_L^{(1)} + A_R^{(1)})^2$.

Π_{linReg}

Inputs: $(x_i, y_i) \in \mathbf{Z}_{2^l} \times \mathbf{Z}_{2^l}$ for $i \in [n]$.

Output: $\mathsf{linReg}((x_1, y_1), \ldots, (x_n, y_n)) = (c_0, c_1)$, where $\hat{y}(x) = c_0 + c_1 x$ is the unique line which minimizes the sum of squares loss $\sum_i (y_i - \hat{y}(x_i))^2$.

1. **Upload**:
   (a) Each P_i encodes their input as $\hat{x}_i = (x_i, x_i^2, y_i, x_i y_i) \in \mathbf{Z}_{2^{6l}}$, where x_i, y_i are each written using l bits and $x_i^2, x_i y_i$ are written using $2l$ bits. We will call these four components (from left to right) $\hat{x}_i^{(j)}$, where $j \in \{1, 2, 3, 4\}$.
   (b) Each P_i then computes $\mathsf{Share}_{\oplus,6l}(\hat{x}_i) = [\hat{x}_i]^B_{L,6l}, [\hat{x}_i]^B_{R,6l} \in \mathbf{Z}_{2^{6l}}$ via Definition 4.2. We refer to the j'th component of these shares as $[\hat{x}_i]^{B,(j)}$.
   (c) P_i sends one share of $[\hat{x}_i]^B_{6l}$ to each server.
2. **Verify Bit-Length**: Initially, $n' = n$. If some server, say S_L, receives from P_i $[\hat{x}_i]^B_{L,6l}$ which is an m-bit integer, $m \neq 6l$:
   (a) S_L sends the index i to S_R.
   (b) Both servers discard $[x_i]^B_{6l}$ (removing from accumulator if necessary).
   (c) Both servers set $n' \leftarrow n' - 1$
3. **Convert Shares**: S_L and S_R jointly evaluate $\mathsf{B2A}_{l,p}(\{[\hat{x}_i]^{B,(1)}_L, [\hat{x}_i]^{B,(1)}_R\})$ as described in Section 9 on each of the n' valid pairs of Boolean shares, which returns $\{[\hat{x}_i]^{A,(1)}_L\}$ to S_L and $\{[\hat{x}_i]^{A,(1)}_R\}$ to S_R. They similarly compute $\mathsf{B2A}_{2l,p}(\{[\hat{x}_i]^{B,(2)}_L, [\hat{x}_i]^{B,(2)}_R\})$, $\mathsf{B2A}_{l,p}(\{[\hat{x}_i]^{B,(3)}_L, [\hat{x}_i]^{B,(3)}_R\})$, and $\mathsf{B2A}_{2l,p}(\{[\hat{x}_i]^{B,(4)}_L, [\hat{x}_i]^{B,(4)}_R\})$. S_L receives as output $\{[\hat{x}_i]^A_L\}$ and S_R receives as output $\{[\hat{x}_i]^A_R\}$, returning $\{[\hat{x}_i]^{A,(j)}_L\}$ to S_L and $\{[\hat{x}_i]^{A,(j)}_R\}$ to S_R for $j \in \{2, 3, 4\}$.
4. **Verify Encoding**: For each valid $\hat{x}_i$ received, servers verify it is truly of the form $(x_i, x_i^2, y_i, x_i y_i)$ analogously to the VAR protocol.

5. **Aggregate**: For each $j \in \{1,2,3,4\}$, S_L computes $A_L^{(j)} = \sum_i [\hat{x}_i]_L^{A,(j)}$ and S_R computes $A_R^{(j)} = \sum_i [\hat{x}_i]_R^{A,(j)}$.

6. **Publish**: S_L and S_R publish their accumulator values $A_L^{(j)}$, $A_R^{(j)}$ for $j \in \{1,2,3,4\}$ to every client as well as the value n'.

7. **Client Computation**: P_i computes $A^{(j)} = A_L^{(j)} + A_R^{(j)}$ for $j \in \{1,2,3,4\}$. P_i computes the output via Equation 7.1, using the values $A^{(1)} = \sum_i x_i$, $A^{(2)} = \sum_i x_i^2$, $A^{(3)} = \sum_i y_i$, $A^{(4)} = \sum_i x_i y_i$, and $n' = n$.

Π_{frq}

Inputs: $x_i \in \{0, \ldots, K-1\}$ for $i \in [n]$.

Output: $\mathsf{FRQ}(x_1, \ldots, x_n) = \vec{h} \in \mathbf{Z}^K$

1. **Upload**:
   (a) Each P_i encodes their input as $\hat{x}_i = (\delta_{x_i}) \in \mathbf{Z}_{2^K}$, the impulse vector at x_i.
   (b) P_i computes $\mathsf{Share}_{\oplus,K}(\hat{x}_i) = [\hat{x}_i]^B_{L,K}, [\hat{x}_i]^B_{R,K}$ via Definition 4.2.
   (c) P_i sends one share of $[\hat{x}_i]^B_K$ to each server.
2. **Verify Bit-Length**: Initially, $n' = n$. If some server, say S_L, receives from P_i $[\hat{x}_i]^B_{L,K}$ which is m-bits, $m \neq K$:
   (a) S_L sends the index i to S_R.
   (b) Both servers discard $[x_i]^B_K$ (removing from accumulator if necessary).
   (c) Both servers set $n' \leftarrow n' - 1$
3. **Convert Shares**: For each index $0 \leq j < K$ and for each of the n' valid pairs of K-bit Boolean shares, S_L and S_R jointly evaluate $\mathsf{B2A}_{1,2^l}(\{([x_i]^B_{L,K})_j, ([x_i]^B_{R,K})_j\})$ as described in [DSZ15]. S_L receives as output $\{[x_i]^*_L\}_i$ and S_R receives as output $\{[x_i]^*_R\}_i$, length-K vectors of Arithmetic shares satisfying $[x_i]^*_L + [x_i]^*_R = x_i \in \mathbf{Z}_{2^K}$.
4. **Verify Encoding**:
   (a) *Verify Odd Parity*:
      i. For each P_i, S_L computes the parity bit $\rho_i = \bigoplus_j ([\hat{x}_i]^B_{L,K})_j$ and sends ρ_i to S_R. S_R similarly computes its own parity bit $\omega_i = \bigoplus_j ([\hat{x}_i]^B_{R,K})_j$.

ii. S_R computes $\sigma_i = \rho_i \oplus \omega_i$, the parity of $\hat{x}_i$.

iii. If $\sigma_i = 1$, S_R does not respond and the verification passes. If $\sigma_i = 0$, S_R responds with the index i and both servers discard $[\hat{x}_i]_K^B$.

(b) *Verify Total Impulse Count*:

i. S_L computes $\sigma_L = \sum_i \sum_j ([\hat{x}_j]_L^*)_i \in \mathbf{Z}_{2^K}$. S_R similarly computes $\sigma_R = \sum_i \sum_j ([\hat{x}_j]_R^*)_i$.

ii. S_L sends σ_L to S_R.

iii. S_R computes $\sigma = \sigma_R + \sigma_L$

iv. If $\sigma = n'$, where n' is the number of non-discarded inputs currently being considered, the verification passes and S_R does not respond.

v. If $\sigma \neq n'$, S_R responds with '0' indicating failure. If $n' > 1$, the servers partition the set of n' remaining players in half lexicographically and repeat this check recursively in parallel on inputs from players $\{P_1, \ldots, P_{n'/2}\}$ and $\{P_{n'/2+1}, \ldots, P_{n'}\}$. If $n' = 1$, and this one remaining player is P_i, both servers discard their shares $[\hat{x}_i]$ once S_R responds with '0' and set $n' \leftarrow n' - 1$.

5. **Aggregate**: For all n' remaining clients, S_L computes $A_L = \sum_i [\hat{x}_i]_L^*$ and S_R computes $A_R = \sum_i [\hat{x}_i]_R^*$.

6. **Publish**: S_L and S_R publish their accumulator values A_L and A_R to every client.

7. **Client Computation**: Clients output $A = A_L + A_R \in \mathbf{Z}_{2^K}$.

APPENDIX B

Security Definitions

We use identical definitions of f-privacy, anonymity, and robustness as used in [CB17]. We give an informal definition of f-privacy, for the sake of readability, which captures the proper security properties. We use the standard notions of negligible functions and computational indistinguishability (see [Gol06]). We often leave the security parameter implicit.

Definition B.1 (f-privacy)**.** Suppose the final aggregate includes data from n different clients. We say that the protocol is f-private if, for:

- any malicious server S_i
- any number of malicious clients $m \leq n$

there exists an efficient simulator that, for every choice of the honest clients' inputs $(x_1, \ldots, x_{n-m})$, takes as input:

- the public parameters to the protocol run (all participants' public keys, the description of the aggregation function f, the cryptographic parameters, etc.),
- the indices of the adversarial clients and server,
- oracle access to the adversarial participants, and
- the value $f(x_1, \ldots, x_{n-m})$,

and outputs a simulation of the adversarial participants' view of the protocol run whose distribution is computationally indistinguishable from the distribution of the adversary's view of the real protocol run.

As discussed in [CB17], the adversary learns the value $f(x_1, \ldots, x_{n-m})$ exactly. If the number of honest players is not sufficiently large, this can leak significant information to the adversary about honest players' inputs. It is therefore up to the servers to ensure that sufficiently many honest players' inputs are included in the aggregate. Our system is subject to the same denial of service and intersection attacks as Prio for the same reasons they give. For a full description, see [CB17].

Definition B.2 (Anonymity)**.** We say that a data-collection scheme provides *anonymity* if it provides f-privacy (by Definition B.1) for $f = \mathsf{SORT}$, where SORT is the function that takes n inputs and outputs them in lexicographically increasing order.

A scheme achieving anonymity by this definition leaks the entire list of honest clients' inputs $(x_1, \ldots, x_{n-m})$ to the adversary. However, the adversary still learns nothing about which client submitted which value x_i. We also have the following Lemma from [CB17] which tells us that providing f-privacy for any symmetric function is necessary and sufficient for anonymity by this definition (where a symmetric function is one that is independent of the order of its inputs).

Lemma 2. *Let $\mathcal{D}$ be a data-collection scheme that provides f-privacy (by Definition B.1) for a symmetric function f. That is, $f(x_1, \ldots, x_n) = f(x_{\pi(1)}, \ldots, x_{\pi(n)}$ for all permutations π on n elements. Then $\mathcal{D}$ provides anonymity.*

Since all aggregate statistics we compute are symmetric functions, once we prove a particular data-collection scheme is f-private, we will have that it also provides anonymity by this result.

Before we can define robustness, recall that each client in the system holds a value $x_i \in \mathcal{D}$, where $\mathcal{D}$ is some set of valid data items (single bits, m-bit integers, impulse vectors). The definition of robustness tells us that when all servers are honest, a set of malicious clients cannot influence the output beyond their ability to choose *valid* inputs.

Definition B.3 (Robustness)**.** Fix a security parameter $\lambda > 0$. We say that a protocol in our scheme provides *robustness* if, when both servers execute the protocol faithfully, for every number m of malicious clients (with $0 \leq m \leq n$), and for every choice of honest clients' inputs $(x_1, \ldots, x_{n-m}) \in \mathcal{D}^{n-m}$, the servers, with all but negligible probability in λ, output a value in the set:

$$\Big\{ f(x_1, \ldots, x_n) | (x_{n-m+1}, \ldots, x_n) \in \mathcal{D}^m \Big\}.$$

APPENDIX C

Proofs of Security

In this section we give proofs of privacy and robustness for each protocol described in Section 8. As discussed in the previous section, [CB17] shows that privacy of any symmetric functionality also implies anonymity. Since each statistic given here is symmetric (does not depend on order of inputs), we get anonymity as a free corollary of privacy in each case.

C.1 Privacy

Our protocols amend the Prio protocol in a few ways. These amendments, however, do not significantly affect the privacy proofs. In particular, our shift from Arithmetic to Boolean secret-sharing affects nothing about the hiding property of the scheme, so any singular shares are still indistinguishable from random. The introduction of the OT-based B2A share conversion reveals no additional information besides the output based on the results of [DSZ15]. The output, from the view of a single corrupted server, is once again a single share of a secret value but now in the arithmetic scheme, again indistinguishable from random. Since these additional interactions between players are all efficiently simulatable, our protocols remain secure. We do have to, however, prove that our daBit-based share conversion protocol is also private in this regard.

Our protocols fall into three categories.

- No share conversion, no SNIPs ($\Pi_{\mathsf{and}}, \Pi_{\mathsf{or}}, \Pi_{\mathsf{max}}, \Pi_{\mathsf{min}}$)
- Share conversion, no SNIPs ($\Pi_{\mathsf{sum}}, \Pi_{\mathsf{frq}}$)
- Share conversion and SNIPs ($\Pi_{\mathsf{var}}, \Pi_{\mathsf{linReg}}$)

Intuitively, protocols in the first category have the easiest proofs of privacy. In these protocols, inputs are encoded in a manner such that any Boolean vector of the proper length is a valid encoding, meaning once client data is split into Boolean secret-shares, all validation work by servers can be done locally by simply verifying the length of the shares. Thus the view of an adversary controlling a single server and arbitrarily many clients contains only a single share of any honest player's input, which is indistinguishable from random according to the hiding property of our secret-sharing scheme. The only other message that the adversary sees is the honest server's aggregated value A_R. This value, however, can be computed deterministically from the output of the function and A_L, which are known to the adversary. Each proof in this category will follow this structure.

The second category is almost identical to the first, except that the protocol calls $\mathsf{B2A}$ on the set of clients' Boolean shares as a subroutine. The privacy of this protocol is proven in [DSZ15], and concludes that neither server, except with negligible probability, learns anything besides the corresponding set of arithmetic shares. This, in combination with the above argument, proves privacy of this set of protocols.

Protocols in the third category rely on the SNIP construction of [CB17] due to the fact that their encoding involves a multiplicative relationship. This is the ideal application of SNIPs and thus we utilize them in this situation. However, we still use $\mathsf{B2A}$ and our bit-length verification technique independently of the SNIP procedure. Thus, the only additional piece necessary to prove privacy for these protocols is the proof of zero-knowledge for SNIPs given in Appendix D.2 of [CB17]. This guarantees that a single semi-honest server

(and any number of misbehaving clients) learn nothing from the SNIP verification of an honest player's input besides the fact that the multiplicative relationship(s) hold. This, in combination with the above arguments, completes the proof of privacy for this third category.

We now give formal proofs of privacy for each protocol.

Theorem 3. *The protocol* Π_{and} *is* AND*-private, where* $x_i \in \{0,1\}$ *and* $\mathsf{AND}(x_1, \ldots, x_n) = 1 \iff \forall i \; x_i = 1$.

Proof. Suppose, without loss of generality, that the adversary A corrupts S_L as well as m clients $P_{n-m+1}, \ldots, P_n$ where $m \leq n$. Suppose honest players $P_1, \ldots, P_{n-m}$ hold inputs $x_1, \ldots, x_{n-m} \in \{0,1\}$. We must construct an efficient simulator which takes the value $\mathsf{AND}(x_1, \ldots, x_n)$ and plays the part of the honest players in the protocol to construct a simulated view V^* which is computationally indistinguishable from the adversary's real view in the protocol $\mathsf{view}(A) = \mathsf{view}(S_L, P_{n-m+1}, \ldots, P_n)$. A's actual view (excluding adversarial inputs, which can be perfectly simulated given oracle access to the adversarial clients) is precisely:

$$\{[\hat{x}_1]_L^B, \ldots, [\hat{x}_{n-m}]_L^B, A_R\}$$

The hiding property of the Boolean and Arithmetic secret-sharing schemes guarantees that a single share reveals nothing about the underlying encoded secret $\hat{x}_i$ and is thus indistinguishable from a random ring element. Thus, we simulate the honest players' shares seen by S_L by sampling random elements from the proper ring, in this case $\mathbf{Z}_2^\lambda$.

To simulate the value A_R, we examine two possible cases. If the output is $\mathsf{AND}(\cdot) = 1$, then the simulator chooses the simulated value $\hat{A}_R$ to be equal to A_L. This is the only possible value of A_R which would result in the proper output, $A_L \oplus A_R = 0$. Otherwise,

if $\mathsf{AND}(\cdot) = 0$ the simulator chooses A_R uniformly at random from $\mathbf{Z}_2^\lambda$, resulting in the proper output $A_L \oplus A_R = \vec{r}$. We know $\hat{A}_R$ and A_R come from the same distribution because they are both uniformly random subject to the constraint of producing the correct output. In conclusion, the simulated adversarial view is precisely $V^* = \{r_1, \ldots, r_{n-m}, \hat{A}_R\}$, where $r_i \in \mathbf{Z}_2^\lambda$ is chosen at random and $\hat{A}_R$ is defined as above. □

Corollary 4. *The protocol* Π_{or} *is* OR*-private, where* $x_i \in \{0, 1\}$ *and* $\mathsf{OR}(x_1, \ldots, x_n) = 0 \iff \forall i \; x_i = 0$.

Proof. The proof is identical except for the simulation of A_R. In this case, if the output is 0, the simulator chooses $\hat{A}_R = A_L$, and if the output is 1 it samples $\hat{A}_R$ randomly from $\mathbf{Z}_2^\lambda$. □

Theorem 5. *The protocol* Π_{max} *is* MAX*-private, where* $x_i \in \{0, \ldots, k-1\}$ *and* $MAX(x_1, \ldots, x_n) = \max\{x_1, \ldots, x_n\}$.

Proof. The proof is nearly identical to the previous proof for AND, but we provide a detailed proof regardless.

Again we suppose, without loss of generality, that the adversary A corrupts S_L as well as m clients $P_{n-m+1}, \ldots, P_n$ where $m \leq n$. Suppose honest players $P_1, \ldots, P_{n-m}$ hold inputs $x_1, \ldots, x_{n-m} \in \{0, \ldots, k-1\}$. We must again construct an efficient simulator which takes the value $\mathsf{MAX}(x_1, \ldots, x_n)$ and plays the part of the honest players in the protocol to construct a simulated view V^* which is computationally indistinguishable from the adversary's real view in the protocol $\mathsf{view}(A) = \mathsf{view}(S_L, P_{n-m+1}, \ldots, P_n)$. A's actual view (excluding adversarial inputs, which can be perfectly simulated given oracle access to the adversarial clients) is once again:

$$\{[\hat{x}_1]_L^B, \ldots, [\hat{x}_{n-m}]_L^B, A_R\}$$

Once again, the hiding property of our secret-sharing schemes guarantees that a single share reveals nothing about the underlying encoded secret $\hat{x}_i$ and is thus indistinguishable from a random ring element, this time in the ring $\mathbf{Z}_2^{\lambda \times k}$. Thus, we simulate the honest players' shares seen by S_L by sampling random elements from $\mathbf{Z}_2^{\lambda \times k}$.

To simulate the value A_R, we again utilize the output of the protocol. Suppose without loss of generality that $\mathsf{MAX}(x_1, \ldots, x_n) = t$, where $t \in \{0, \ldots, M-1\}$. The adversary simply chooses $\hat{A}_R$ at random from $\mathbf{Z}_2^{\lambda \times k}$ subject to the following constraint: $(A_L \oplus A_R)_j = 0$ for $j < t\lambda$ and $(A_L \oplus A_R)_k = 1$ for some $k \in \{t\lambda, \ldots, (t+1)\lambda - 1\}$. We once again know $\hat{A}_R$ and A_R come from the same distribution because they are both uniformly random subject to the constraint of producing the correct output. In conclusion, the simulated adversarial view is precisely $V^* = \{r_1, \ldots, r_{n-m}, \hat{A}_R\}$, where $r_i \in \mathbf{Z}_2^{\lambda \times k}$ is chosen at random and $\hat{A}_R$ is defined as above.

□

Corollary 6. *The protocol* Π_{min} *is* MIN*-private, where* $\mathsf{MIN}(x_1, \ldots, x_n) = \min\{x_1, \ldots, x_n\}$ *and* $x_i \in \{0, 1\}$.

Proof. The only interaction in Π_{min} is a single call to the Π_{max} protocol, whose privacy was established via the previous theorem. □

This concludes proofs of privacy for our first class of protocols which do not require share conversion. Although the next few proofs involve share conversion, they only require OT-based share conversion, so we simply rely on the proof of privacy from [DSZ15].

Theorem 7. *The protocol* Π_{sum} *is* SUM*-private, where* $x_i \in \mathbf{Z}_2^{l'}$ *and* $\mathsf{SUM}(x_1, \ldots, x_n) = \sum x_i$ *(as integers in* $\mathbf{Z}$*).*

Proof. Suppose, without loss of generality, that the adversary A corrupts S_L as well as m clients $P_{n-m+1}, \ldots, P_n$ where $m \leq n$. Suppose honest players $P_1, \ldots, P_{n-m}$ hold inputs $x_1, \ldots, x_{n-m} \in \mathbf{Z}_{2^{l'}}$. We must construct an efficient simulator which takes the value

$\mathsf{SUM}(x_1, \ldots, x_n) = \sum_i x_i \in \mathbf{Z}$ and plays the part of the honest players in the protocol to construct a simulated view V^* which is computationally indistinguishable from the adversary's real view in the protocol $\mathsf{view}(A) = \mathsf{view}(S_L, P_{n-m+1}, \ldots, P_n)$. The underlying $\mathsf{B2A}$ protocol is $\mathcal{F}_{b2a}$-private (see [DSZ15]), so the adversary only learns the output of $\mathsf{B2A}$ and nothing else. Thus A's actual view (excluding dishonest players' inputs) is precisely:

$$\{[x_1]_L^B, \ldots, [x_{n-m}]_L^B, [x_1]_L^A, \ldots, [x_{n-m}]_L^A, A_R\}$$

Based on the hiding property of the Boolean and Arithmetic secret-sharing schemes, a single share in either scheme reveals nothing about the underlying secret and is thus indistinguishable from a random element of the corresponding ring. Thus, we simulate the honest players' shares seen by S_L by sampling random elements of the proper length, l' bits for the Boolean shares and l bits for the Arithmetic shares. $\hat{A}_R$ can be computed as $\mathsf{SUM}(x_1, \ldots, x_n) - A_L$, since $A_L + A_R = \mathsf{SUM}(\cdot)$. Thus the simulated view is $V^* = \{r_1, \ldots, r_n - m, r'_1, \ldots, r'_{n-m}, \hat{A}_R\}$, where $r_i \in \mathbf{Z}_{2^{l'}}$ and $r'_i \in \mathbf{Z}_{2^l}$ are chosen at random. □

Note that this proves privacy of our bit sum protocol as a corollary by simply substituting $l' = 1$.

Corollary 8. *The protocol* Π_{mean} *is* MEAN*-private, where* $x_i \in \mathbf{Z}_{2^{l'}}$ *and* $\mathsf{MEAN}(x_1, \ldots, x_n) = (\frac{1}{n}\sum_i x_i, n)$.

Proof. The Π_{mean} protocol is identical to the Π_{SUM} protocol except that servers also release to clients in the clear the number of players n' who submitted valid inputs so that they can properly compute the mean. The adversary's view and the simulated view are identical. □

Theorem 9. *The protocol* $\Pi_{\mathsf{frq}}(x_1, \ldots, x_n)$ *is* FRQ*-private, where* $x_i \in \{0, \ldots, k-1\}$ *and* $\mathsf{FRQ}(x_1, \ldots, x_n = \vec{h} - (f_1, \ldots, f_k) \in \mathbf{Z}_{n+1}^k$, *where* $f_j = |\{x_i : x_i = j\} \leq n$ *is the frequency of input* $j \in \mathbf{Z}_k$.

Proof. We once again assume A corrupts S_L as well as the last m clients. If we ignore the two additional verification steps ensuring that each input is an impulse vector, the adversary's view is once again

$$\mathsf{view}(A) = \{[x_1]_L^B, \ldots, [x_{n-m}]_L^B, [x_1]_L^A, \ldots, [x_{n-m}]_L^A, A_R\}$$

just as in the SUM protocol. Simulating the S_L shares is once again done via random sampling in the rings $\mathbf{Z}_{2^{l'}}$ and $\mathbf{Z}_{2^l}$ for the Boolean and Arithmetic shares respectively. Simulating A_R is once again done using the output of the protocol and A_L. Specifically, $\hat{A}_R = \vec{h} \oplus A_L$. Thus, to prove privacy, we need only prove the privacy of these two additional verification steps.

The first such step is a parity check on each encoded input $\hat{x}_i$ to make sure it contains an odd number of 1's. This is accomplished in the clear by having each server computing the parity of their share (call these parity bits $p_{i,L}$ and $p_{i,R}$) and taking the direct sum $p_i = p_{i,L} \oplus p_{i,R}$. Note that for any honest player P_i, we must have $p_i = 1$. Similarly, for any corrupted player P_j, our simulator knows the value p_j exactly. Thus, we can simulate all $p_{i,R}$ exactly by computing $p_{i,R} = p_i \oplus p_{i,L}$.

The second check is to ensure that no player submitted multiple impulses. To do this, servers compute B2A on each component of each secret-shared "impulse." Based on the privacy of B2A from [DSZ15], this does not reveal any information besides the output, the corresponding Arithmetic shares. The servers sum all vectors together and sum the components of the final vectors to get arithmetic shares of the total number of impulses submitted by all clients. Note that this value is is simply $n - c$, where c is the number of adversarial clients who submitted invalid inputs. Since our simulator knows the value c, this can be simulated exactly. If $c = 0$, we are done. Otherwise, the process repeats recursively on smaller subsets of players. Note that no matter how small these subsets become, the sum

$n' - c'$ of total impulses in any subset will never reveal anything about an honest player's input, since it is publicly known that all honest players submitted a single impulse. This step will only possibly reveal adversarial players' inputs. Thus this second check is again private.

Since every message seen by the adversary in these additional verification steps is efficiently simulatable, we conclude that the protocol is private. □

Before we can prove the privacy of Π_{var}, Π_{stddev}, and Π_{linReg}, we must prove the privacy of the share conversion protocol on which these protocols rely.

Theorem 10. $\mathsf{B2A}_p$*, as described in Section 9, is* $\mathcal{F}_{B2A}$*-private.*

Proof. First note that there is only a single communication round in the protocol, in which servers send their share $[v]_i^B$ to each other. Thus, any information about x that is leaked in the protocol must be leaked by this single message. However, $[v]_i^B = [x]_i^B \oplus [b]_i^B$, where b is a random bit whose value is not known in the clear to either party. Since neither x nor b is known in the clear to either party, even an adversary who knows both $[v]_L^B = [x]_L^B \oplus [b]_L^B$ and $[v]_R^B = [x]_R^B \oplus [b]_R^B$ does not have enough information to deduce x or b. For example, an adversary corrupting server L would receive from the honest server R a value $[v]_R^B$, which it knows is equal to $[x]_R^B \oplus [b]_R^B$, but since both of these values are unknown, neither can be deduced from the value $[v]_R^B$. An analogous argument applies when server R is corrupted. Since no information about x or b can be deduced by an adversary corrupting a single party, and these are the only inputs to the protocol, we can conclude that $\mathsf{B2A}_p$ is $\mathcal{F}_{B2A}$-private. □

We now have the tools to prove that Π_{var}, Π_{stddev}, and Π_{linReg} are private.

Theorem 11. *The protocol* Π_{var} *is* $\hat{\mathsf{VAR}}$*-private, where* $x_i \in \mathbf{Z}_{2^l}$ *and* $\hat{\mathsf{VAR}}(x_1, \ldots, x_n) = (\mathsf{VAR}\{x_1, \ldots, x_n\}, \mathbf{E}\{x_1, \ldots, x_n\})$.

Proof. Our proof is analogous to the previous examples except for three departures. First, our proof additionally relies on the privacy of the SNIP verification procedure. The full proof of security can be found in Appendix D of [CB17]. Their conclusion is that as long as one server is honest, an adversary controlling the remaining servers (and any coalition of clients) learns nothing from the SNIP verification procedure. This is analogous to how we rely on the security of the OT-based $\mathsf{B2A}$ protocol from [DSZ15] to conclude that the adversary's learns nothing besides the output of the protocol, except with negligible probability. This tells us that the view of our adversary remains the same before and after SNIP verification.

The second departure is that this protocol relies on our daBit-based share conversion protocol. Since we have already proven that this protocol is private, we simply use that result to ensure that the adversary learns nothing but the output during share conversion. The adversary's view (not including adversarial inputs) is thus:

$$\{\{[f_i(0)]_L\}, \{[g_i(0)]_L\}, \{[h_i]\}, \{[\hat{x}_1]_L^B, \ldots, [\hat{x}_{n-m}]_L^B, [\hat{x}_1]_L^A, \ldots, [\hat{x}_{n-m}]_L^A, A_R\}$$

To simulate the polynomial shares, we call the simulator used in the SNIPs security proof in Appendix D of [CB17] as a subroutine. For the rest of the shares, we simulate via random sampling once again based on the hiding property of the secret sharing schemes.

The third departure is that the value A_R, when combined with A_L, reveals $\mathbf{E}[X]$ in addition to $\mathsf{VAR}[X]$, where X is a uniform random variable over values $x_1, \ldots, x_n$. This is why our functionality $V\hat{A}R$ additionally reveals $\mathbf{E}[X]$, meaning our simulator receives this as input. From this and the variance, the simulator can reconstruct $\mathbf{E}[X^2]$ using the identity $\mathsf{VAR}[X] = \mathbf{E}[X^2] - (\mathbf{E}[X])^2$. Then, since $A_L + A_R = \mathbf{E}[X^2] - (\mathbf{E}[X])^2$ with overwhelming probability, the simulator can simulate A_R by simply computing $\hat{A}_R = \mathbf{E}[X^2] - (\mathbf{E}[X])^2 - A_L$.

□

Corollary 12. *The protocol* Π_{stddev} *is* $\hat{\mathsf{STDDEV}}$*-private, where* $\hat{\mathsf{STDDEV}}(X) = (\sqrt{\hat{\mathsf{VAR}}(X)}, \mathbf{E}[X])$.

Proof. The protocol is identical to Π_{var} with an additional local operation applied on the client side. Since there is no additional communication, the $\hat{\mathsf{VAR}}$-privacy of Π_{var} implies the corollary. □

Theorem 13. *The protocol* Π_{linReg} *is* $\hat{\mathsf{linReg}}$*-private, where* $\hat{\mathsf{linReg}}$ *additionally outputs* n'*, the number of players whose shares were counted in the aggregate.*

Proof. This proof is completely analogous to the previous proof for Π_{var}. The only difference between the protocols, besides the length of encoded shares, is that two multiplicative relationships are being verified within the encoded inputs. These are both done using SNIPs, however, and so with overwhelming probability neither server learns any additional information based on Appendix D in [CB17]. Simulation of A_R, once n', c_0, c_1, and A_L are known, can be done via simple matrix operations on the matrix equation given in Section 6. □

C.2 Robustness

In this section, we give proofs that each of our protocols are robust. That is, no client can affect the output of the protocol beyond misreporting their private value as some other valid input value. Our first and second categories of protocols (with the exception of Π_{frq}) will have the simplest robustness proofs. This is because all Boolean strings of the proper length are valid encodings, meaning the bit-length is the only property which servers must verify, and doing so is trivial when client data is shared via the Boolean scheme. In the protocols which utilize SNIPs, our robustness relies on the SNIPs' soundness property, proven in Appendix D of [CB17]. It guarantees that a malicious prover has only a negligible probability of successfully submitting an invalid proof to the verifiers. This, in combination with the trivial verification of bit-length, guarantees that each input lies in the correct range and they obey the proper multiplicative relationships. The most involved proof of robustness is

for the Π_{frq} protocol, since we must give an argument from scratch of the correctness of our verification procedure.

Theorem 14. *The protocol* $\Pi_{\mathsf{and}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \{0, 1\}$.

Proof. Note that any λ-bit Boolean vector is a valid encoded input. Any player who does not submit a λ-bit Boolean vector to each server is ignored by the protocol, which is equivalent to having an input of zero. Any player who does submit a λ-bit Boolean vector to each server, call them $v_L, v_R \in \mathbf{Z}_{2^\lambda}$, has submitted an encoding of a valid input, since $v_L \oplus v_R$ is a λ-bit Boolean vector. Thus any misbehaving client must either submit a valid encoded input or be treated as if they submitted input 0, which satisfies the definition of robustness. □

Theorem 15. *The protocol* $\Pi_{\mathsf{or}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \{0, 1\}$.

Proof. See proof of previous theorem. □

Theorem 16. *The protocol* $\Pi_{\mathsf{max}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \{0, \ldots, M-1\}$.

Proof. Same as previous proofs, except that valid encoded inputs are $M \times \lambda$-bit Boolean vectors. □

Theorem 17. *The protocol* $\Pi_{\mathsf{min}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \{0, \ldots, M-1\}$.

Proof. Same as previous proof, valid encoded inputs are $M \times \lambda$-bit Boolean vectors. □

Theorem 18. *The protocol* $\Pi_{\mathsf{sum}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \mathbf{Z}_{2^{l'}}$, $l' < l$.

Proof. Note that once again, any l'-bit Boolean vector is a valid encoded input. Thus, as long as each server received an l'-bit Boolean vector, the client submitted a valid input. For the client to submit anything besides a valid input, he must send to at least one server something besides an l'-bit Boolean vector, at which point the servers will detect it with certainty, discard that clients' shares, and continue as if that client submitted a zero. □

Corollary 19. *The protocol* $\Pi_{\mathsf{mean}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \mathbf{Z}_{2^{l'}}$, $l' < l$.

Proof. The protocol consists of one call to Π_{sum} and local operations, so robustness of Π_{sum} implies robustness of Π_{sum}. □

This concludes our trivial robustness proofs. The following proofs rely on the soundness property of SNIPs, for which we rely on the results of [CB17].

Theorem 20. *The protocol* $\Pi_{\mathsf{var}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \mathbf{Z}_{2^l}$.

Proof. Clients must submit shares of the proper length by the argument given in previous proofs. The only additional factor in this protocol is that the encoded input must be of the form (x, x^2). This is accomplished via the use of SNIPs, and our robustness property relies on the soundness of SNIPs proven in Appendix D of [CB17]. Their result guarantees that no client has more than $\frac{2\mu+1}{2^l}$ probability of submitting an input not of the form (x, x^2) for which the SNIP verification succeeds, where μ is the number of multiplication gates in the Valid circuit. Here, $\mu = \mathsf{poly}(l)$, so this probability is negligible. Since any input whose SNIP verification fails will be ignored (which is equivalent to submitting the value 0), this implies that no adversarial client can affect the output beyond misreporting their private value except with negligible probability. Thus, Π_{var} is robust. □

Corollary 21. *The protocol* $\Pi_{\mathsf{stddev}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \mathbf{Z}_{2^l}$.

Proof. The protocol consists of one call to Π_{var} and local operations, so robustness of Π_{var} implies robustness of Π_{stddev}. □

Theorem 22. *The protocol* $\Pi_{\mathsf{linReg}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \mathbf{Z}_{2^l}$.

Proof. In this case, we must verify that the encoded input is of the form (x, x^2, y, xy). The proof is identical to the proof of robustness for Π_{var} except that the Valid circuit verifies both multiplicative relationships at once. That is, according to the soundness of SNIPs proven in Appendix D of [CB17], no client has more than a negligible probability ($\frac{2\mu+1}{2^l}$) of submitting an improperly encoded input for which the SNIP verification algorithm succeeds. Thus, by the same argument as Π_{var}, Π_{linReg} is robust. □

Theorem 23. *The protocol* $\Pi_{\mathsf{frq}}(x_1, \ldots, x_n)$ *is robust, where* $x_i \in \{0, \ldots, k-1\}$.

Proof. If an adversarial client submits anything besides an k-bit Boolean vector to either server, their input will be discarded. Otherwise, the only ways to submit an invalid encoding are to submit a zero vector or to submit multiple impulses. If P_i submits a zero vector, servers will discover it has an even parity when they compute the parity in the clear. Thus, the input will be discarded. Since all zero vectors are detected with certainty in this step, assume from this point onwards that the servers do not hold shares of any zero vector and that there are n' remaining pairs of shares. If P_i submits a vector x such that $(x)_i = (x)_j = 1$, then the sum of all the components in the sum of all vectors will be greater than n', since there are no zero vectors. This means the players will be partitioned lexicographically and the check will repeat recursively. In the base case of this recursion, P_i will be the only member of his partitioned set, and this check will reveal that the sum of the components in his input is larger than 1, and his input will be discarded. This happens with certainty. Thus, no client can submit an invalid input without their input being ignored, satisfying the definition of robustness. □

This concludes our proofs of security. We have shown that each protocol described in this paper is private (with at most the same modest leakage as Prio), robust, and anonymous.

REFERENCES

[BBC19] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. "Zero-Knowledge Proofs on Secret-Shared Data via Fully Linear PCPs." Cryptology ePrint Archive, Report 2019/188, 2019. https://eprint.iacr.org/2019/188.

[BFO12] Eli Ben-Sasson, Serge Fehr, and Rafail Ostrovsky. "Near-linear unconditionally-secure multiparty computation with a dishonest minority." In *Annual Cryptology Conference*, pp. 663–680. Springer, 2012.

[CB17] Henry Corrigan-Gibbs and Dan Boneh. "Prio: Private, robust, and scalable computation of aggregate statistics." In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pp. 259–282, 2017. https://crypto.stanford.edu/prio/paper.pdf.

[CDI05] Ronald Cramer, Ivan Damgård, and Yuval Ishai. "Share Conversion, Pseudorandom Secret-Sharing and Applications to Secure Computation." In Joe Kilian, editor, *Theory of Cryptography*, pp. 342–362, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. https://link.springer.com/content/pdf/10.1007%2F978-3-540-30576-7_19.pdf.

[CS98] Ronald Cramer and Victor Shoup. "A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack." In Hugo Krawczyk, editor, *Advances in Cryptology — CRYPTO '98*, pp. 13–25, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[DFK13] George Danezis, Cédric Fournet, Markulf Kohlweiss, and Santiago Zanella-Béguelin. "Smart Meter Aggregation via Secret-Sharing." In *Proceedings of the First ACM Workshop on Smart Energy Grid Security*, SEGS '13, p. 75–80, New York, NY, USA, 2013. Association for Computing Machinery.

[DSZ15] Daniel Demmler, Thomas Schneider, and Michael Zohner. "ABY-A framework for efficient mixed-protocol secure two-party computation." In *NDSS*, 2015.

[EDG14] Tariq Elahi, George Danezis, and Ian Goldberg. "PrivEx: Private Collection of Traffic Statistics for Anonymous Communication Networks." In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, p. 1068–1079, New York, NY, USA, 2014. Association for Computing Machinery.

[EGK20] Daniel Escudero, Satrajit Ghosh, Marcel Keller, Rahul Rachuri, and Peter Scholl. "Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits." Springer-Verlag, 2020.

[EKO19] K. Emura, H. Kimura, T. Ohigashi, T. Suzuki, and L. Chen. "Privacy-Preserving Aggregation of Time-Series Data with Public Verifiability from Simple Assumptions and Its Implementations." *The Computer Journal*, **62**(4):614–630, 2019.

[EKP14] Úlfar Erlingsson, Aleksandra Korolova, and Vasyl Pihur. "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response." *CoRR*, **abs/1407.6981**, 2014.

[enc09] "VPriv: Protecting Privacy in Location-Based Vehicular Services." In *18th USENIX Security Symposium (USENIX Security 09)*, Montreal, Quebec, August 2009. USENIX Association.

[FPE15] Giulia C. Fanti, Vasyl Pihur, and Úlfar Erlingsson. "Building a RAPPOR with the Unknown: Privacy-Preserving Learning of Associations and Data Dictionaries." *CoRR*, **abs/1503.01214**, 2015.

[GLL14] James Glanz, Jeff Larson, and ANDREW W Lehren. "Spy agencies tap data streaming from phone apps." *New York Times*, 2014.

[Gol06] Oded Goldreich. *Foundations of Cryptography: Volume 1*. Cambridge University Press, USA, 2006.

[HPK16] Andrew Hilts, Christopher Parsons, and Jeffrey Knockel. "Every step you fake: A comparative analysis of fitness tracker privacy and security." *Open Effect Report*, **76**(24):31–33, 2016.

[Jes13] Tobias Jeske. "Floating car data from smartphones: What google and waze know about you and how hackers can control traffic." *Proc. of the BlackHat Europe*, pp. 1–12, 2013.

[JL13] Marc Joye and Benoît Libert. "A Scalable Scheme for Privacy-Preserving Aggregation of Time-Series Data." In Ahmad-Reza Sadeghi, editor, *Financial Cryptography and Data Security*, pp. 111–125, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[KLP15] Josh Keller, KR Lai, and Nicole Perlroth. "How many times has your personal information been exposed to hackers." *New York Times (July 29, 2015)*, 2015.

[MDC15] Luca Melis, George Danezis, and Emiliano De Cristofaro. "Efficient Private Statistics with Succinct Sketches." *CoRR*, **abs/1508.06110**, 2015.

[PBB11] Raluca Ada Popa, Andrew J. Blumberg, Hari Balakrishnan, and Frank H. Li. "Privacy and Accountability for Location-Based Aggregate Statistics." In *Proceedings of the 18th ACM Conference on Computer and Communications Security*, CCS '11, p. 653–666, New York, NY, USA, 2011. Association for Computing Machinery.

[Rin] Peter Rindal. "libOTe: an efficient, portable, and easy to use Oblivious Transfer Library." `https://github.com/osu-crypto/libOTe`.

[RW19] Dragos Rotaru and Tim Wood. "MArBled Circuits: Mixing Arithmetic and Boolean Circuits with Active Security." Cryptology ePrint Archive, Report 2019/207, 2019.

[Sho01a] Victor Shoup. "OAEP Reconsidered." In Joe Kilian, editor, *Advances in Cryptology — CRYPTO 2001*, pp. 239–259, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.

[Sho01b] Victor Shoup. "A Proposal for an ISO Standard for Public Key Encryption." *IACR Cryptology ePrint Archive*, **2001**:112, 01 2001.

[Smi14] Ben Smith. "Uber executive suggests digging up dirt on journalists." *BuzzFeed News*, **18**, 2014.

[WWW16] Gang Wang, Bolun Wang, Tianyi Wang, Ana Nika, Haitao Zheng, and Ben Y Zhao. "Defending against sybil devices in crowdsourced mapping services." In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, pp. 179–191, 2016.