SIGMA: Secure GPT Inference with Function Secret Sharing

Kanav Gupta*
Microsoft Research
kanav0610@gmail.com

Neha Jawalkar*†
Indian Institute of Science
jawalkarp@iisc.ac.in

Ananta Mukherjee
Microsoft Research
t-mukherjeea@microsoft.com

Nishanth Chandran
Microsoft Research
nichandr@microsoft.com

Divya Gupta
Microsoft Research
divya.gupta@microsoft.com

Ashish Panwar
Microsoft Research
ashishpanwar@microsoft.com

Rahul Sharma
Microsoft Research
rahsha@microsoft.com

***Abstract*—Secure 2-party computation (2PC) enables secure inference that offers protection for both proprietary machine learning (ML) models and sensitive inputs to them. However, the existing secure inference solutions suffer from high latency and communication overheads, particularly for transformers. Function secret sharing (FSS) is a recent paradigm for obtaining efficient 2PC protocols with a preprocessing phase. We provide SIGMA, the first end-to-end system for secure transformer inference based on FSS. By constructing new FSS-based protocols for complex machine learning functionalities, such as Softmax and GeLU, and also accelerating their computation on GPUs, SIGMA improves the latency of secure inference of transformers by $11-19\times$ over the state-of-the-art that uses preprocessing and GPUs. We present the first secure inference of generative pre-trained transformer (GPT) models. In particular, SIGMA executes GPT-Neo with 1.3 billion parameters in 7.4s and HuggingFace's GPT2 in 1.6s.**

1. Introduction

In the problem of secure inference, model providers own proprietary machine learning (ML) models that they want to offer as services and clients who want to learn the inference results on their sensitive data. The security requirement is that the client should learn nothing about the model beyond the inference output and the model provider should learn nothing about the client's input. This problem can be solved by the technique of secure 2-party computation (2PC) that provides cryptographic security guarantees.

In recent years, the applicability of 2PC-based solutions has scaled up from models with thousands of parameters [10], [39], [48], [50]–[52], [55], [58], [60], [62], [72], to models with millions of parameters [20], [32], [35], [38], [42], [59], [73], [76], to now BERT models with hundreds of millions of parameters [7], [19], [33], [41], [45]. In this paper, we take a step further in this direction by providing secure inference of Generative Pre-trained Transformer (GPT) models with billions of parameters.

. * Equal contribution.
. † Work done while at Microsoft Research.

Transformer-based generative language models have gained significant traction in recent times due to their remarkable performance on various natural language tasks e.g., question-answering, summarization, language translation, code generation [15], [16], [67]. Apart from ensuring model/input privacy, secure inference of such models opens up other interesting scenarios like "prompt privacy". AI companies are spending significant efforts building prompts that lead to good inference results and they want to keep the prompts hidden. Secure inference allows a company holding a proprietary prompt and a client holding sensitive data to generate inference results from a public language model without revealing their inputs to each other. However, the current state-of-the-art systems for secure inference deliver unsatisfactory results on transformers.

We posit that a system for secure ML inference must satisfy the following requirements: (1) *accuracy* - i.e., the accuracy under secure inference should match that of the plaintext, (2) *security* - i.e., the system should provide standard 2PC security, (3) *efficiency* - i.e., the latency and communication overheads of secure inference should be low, and (4) *scalability* - i.e., the system must scale to models with billions of parameters. We show that existing systems fail to meet (often more than one of) these requirements.

Existing secure transformer inference systems include THE-X [19], Iron [33], and CrypTen [41], [45], [75] (we discuss other works in Section 8). THE-X sacrifices both accuracy, by replacing *complex* non-linearities (based on elementary functions, e.g., e^x) with simple non-linearities ($\max(x, 0)$), and security, by revealing intermediate values. Iron maintains both accuracy and security, but has huge communication overheads, requiring over a hundred GB of communication even for BERT models. Although CrypTen leverages GPU acceleration and preprocessing to improve efficiency, its online latency and communication for secure inference are still significant. Moreover, it fails to provide standard 2PC security because it uses insecure[1] *local* trun-

1. Secure inference works like CrypTen [41] and many others [52], [65], [72], [73], [76] use cheap local truncations that have recently been established as insecure [46].

cations. Furthermore, because of GPU memory overflows, it fails to scale to larger models.

1.1. Our Contributions

In this paper, we propose SIGMA[2] - a system that advances the state-of-the-art for secure inference of transformer-based models along multiple dimensions. Like CrypTen, SIGMA works in 2PC with preprocessing model and uses GPU acceleration, but is an *order of magnitude* more efficient in latency and communication while providing standard 2PC security guarantees. SIGMA maintains the model accuracy under secure inference through precise approximations of complex non-linearities and scales efficiently to GPT models with billions of parameters.

SIGMA leverages Function Secret Sharing (FSS) based 2PC protocols [11], [14], [32], [38] and builds on Orca [38] that is the state-of-the-art in GPU-accelerated FSS-protocols. However, Orca focuses primarily on convolutional neural networks (CNNs) that use simple non-linearities like ReLU. We show that Orca's techniques pose unacceptable overheads for transformers because of their heavy use of complex non-linearities (Section 7.1.2).

Since the latency of secure inference in transformers is dominated by complex non-linearities - GeLU, Softmax, and layer normalization [33] - we propose new FSS-based protocols for these operations and accelerate them with GPUs. Realizing these operations requires accurate computation of various elementary functions, e.g., exponentiation, reciprocal square root, inverse, etc. The prior work of Pika [71] uses large look-up tables (LUTs) for these functions. Although this approach is general, Grotto [64] shows that large LUTs are inefficient and provides protocols based on custom splines (when they exist). SIGMA's protocols minimize the size of LUTs, to maintain accuracy, while being more efficient than Grotto (Section 7.1.1). For instance, for GeLU over 50-bit values, while Pika requires an LUT of size 2^{50}, SIGMA uses an LUT of size 2^8 and overall requires $9\times$ lower compute than Grotto in the same threat model.

We evaluate SIGMA on various models based on GPT and BERT [15], [23], which are widely used for next-word-predictions and classification tasks, respectively. Our novel protocols allow us to securely and accurately evaluate GPT-Neo with 1.3 billion parameters – "a transformer model designed using EleutherAI's replication of the GPT-3 architecture" [4] – in 7.4 seconds. SIGMA runs the smaller GPT2 model [3] from HuggingFace (tens of millions of downloads each month) in 1.6 seconds, and the BERT models in $0.1 - 4.7$ seconds. Overall, SIGMA improves the latency of secure inference by $11.5 - 19.4\times$ over the state-of-the-art. Finally, by evaluating on GPT models with up to 13 billion parameters, we show that SIGMA scales well to even bigger models.

To guarantee standard 2PC security, SIGMA does away with local truncations and instead uses secure faithful truncations. Truncations are used extensively in both linear layers, i.e., after matrix multiplications, and non-linear layers. We provide a new protocol for faithful truncation (Section 4.2) that is much more efficient than the prior work [11] (up to $30\times$). Even though our truncations are costlier than (almost free) local truncations in CrypTen, our massive performance gains in GeLU and Softmax make SIGMA more than $10\times$ faster than CrypTen for end-to-end inference.

Our large scale evaluations are made possible by SIGMA's frontend that allows users to succinctly express a transformer architecture of choice and run it with SIGMA's protocols optimized for CPUs or GPUs (Section 6). The protocol design for CPUs and GPUs differ, and we support both (Section 5.1). In fact, SIGMA running on CPUs is already faster than CrypTen running on GPUs. SIGMA will be made publicly available.

Organization. Section 2 discusses cryptography background. Section 3 provides an overview of transformers and GPUs. Section 4 provides our protocols for primitive operations like truncations and comparisons. Section 5 builds on these primitives to provide novel protocols for complex non-linearities. Section 6 provides implementation details, Section 7 describes our empirical results, Section 8 discusses related work, and Section 9 concludes.

2. Preliminaries

2.1. Notation

Let λ be the computational security parameter, $N = 2^n$ and $L = 2^\ell$. Let $\mathbb{R}$ denote the set of real numbers and $\mathbb{U}_{2^n}$ denote the set of n-bit unsigned integers. We use standard 2's complement representation to represent signed values in $\mathbb{U}_N$. For $x \in \mathbb{U}_N$, $\mathsf{int}_n(x)$ and $\mathsf{uint}_n(x)$ denote the corresponding signed and unsigned integers in $\mathbb{Z}$, respectively. We denote arrays using boldface and its i-th element (starting at 0) using the same symbol in normal typeface followed by $[i]$, e.g., $\boldsymbol{a} = \{a[0], a[1], a[2], \dots\}$.

Fixed-Point Representation. Fixed-point representation, parameterized by bitwidth n and precision f, encodes a real value $r \in \mathbb{R}$ into an n-bit integer $x \in \mathbb{U}_N$ such that $x = \lfloor r \cdot 2^f \rfloor \mod N$. Conversely, an n-bit fixed-point number x with precision f decodes into real number $\frac{\mathsf{int}_n(x)}{2^f}$.

Operators. For a predicate b, $1\{b\} \in \{0, 1\}$ returns 1 if b is true and 0 otherwise. For $n < \ell$, $x \in \mathbb{U}_N$, $\mathsf{extend}_{n,\ell}(x)$ returns x appended with $(\ell - n)$ 0's on the left. For $x \in \mathbb{U}_N$, $\mathsf{MSB}_n(x) \in \{0, 1\}$ denotes the most-significant bit of x.

Secret Sharing. For $x \in \mathbb{U}_N$, *secret sharing* samples random *shares* $x_0, x_1 \in \mathbb{U}_N$ such that $x = x_0 + x_1 \mod N$ holds, and is denoted by $\mathsf{share}(x)$. When x_0 is held by P_0 and x_1 is held by P_1, we denote the process of exchanging the shares and adding them to reconstruct the underlying value by $\mathsf{reconstruct}(x_b)$ for $b \in \{0, 1\}$.

2.2. Threat Model

This work considers standard 2PC in the preprocessing model [8], [9], [14], [22], [37] that has also received

significant attention in the context of secure machine learning [32], [38], [41], [63], [64], [76]. That is, there are two parties P_0 and P_1 with inputs x_0 and x_1 and they wish to compute a public function $y = f(x_0, x_1)$ without revealing anything more than the function output y to each other. In a preprocessing phase that is independent of the inputs to the function x_0 and x_1, correlated randomness is generated and made available to P_0 and P_1. This randomness can be generated through multiple ways - a trusted dealer [11], [14], [32], [38], [41], [63], [64], [76], generic 2PC protocols [28], [78], or through more efficient specialized 2PC protocols [24]. In this work, we consider the first approach. All our protocols satisfy the standard notion of simulation-based security [17], [28], [47] with security provided against a semi-honest static probabilistic polynomial time (PPT) adversary corrupting either P_0 or P_1.

2.3. Function Secret Sharing

A Function Secret Sharing (FSS) [12], [13] scheme is a pair of algorithms $(\mathsf{Gen}, \mathsf{Eval})$. Gen splits a function g into two function shares (g_0, g_1) and Eval takes as input $b \in \{0,1\}$, function share g_b and input x and returns $g_b(x)$. The correctness property of an FSS scheme requires that $g_0(x) + g_1(x) = g(x)$ for all x. The security property requires that each function share g_b hides the function g.

Definition 1 (FSS: Syntax [12], [13])**.** *A (2-party) FSS scheme is a pair of algorithms* $(\mathsf{Gen}, \mathsf{Eval})$ *such that:*

- $\mathsf{Gen}(1^\lambda, \hat{g})$ *is a PPT key generation algorithm that given* 1^λ *and* $\hat{g} \in \{0,1\}^*$ *(description of a function* g*) outputs a pair of keys* (k_0, k_1)*. We assume that* $\hat{g}$ *explicitly contains descriptions of input and output groups* $\mathbb{G}^{\mathsf{in}}, \mathbb{G}^{\mathsf{out}}$*.*
- $\mathsf{Eval}(b, k_b, x)$ *is a polynomial-time evaluation algorithm that given* $b \in \{0,1\}$ *(party index),* k_b *(key defining* $g_b : \mathbb{G}^{\mathsf{in}} \to \mathbb{G}^{\mathsf{out}}$*) and* $x \in \mathbb{G}^{\mathsf{in}}$ *(input for* g_b*) outputs* $y_b \in \mathbb{G}^{\mathsf{out}}$ *(the value of* $g_b(x)$*).*

(k_0, k_1) are called FSS keys and the number of bits required to store one FSS key is called *key size*. We formally define correctness and security of an FSS scheme in Appendix F.

2.4. 2PC with preprocessing from FSS

Consider secure computation of a circuit with gates $\{g_i\}_i$ and wires $\{w_i\}_i$. We describe the 2PC protocol with preprocessing using FSS from [14] in two phases.

Offline Phase. For each wire w_i, sample a random mask r_i from the appropriate group. Then, for each of the gate g with input wire w_i and output wire w_j, generate an FSS key for its offset function $g^{[r_i, r_j]}(x) = g(x - r_i) + r_j$ and provide one key to each party. For input and output wires of the circuit belonging to party b, that party also learns the masks associated with those wires.

Online Phase. For each input wire w_i with value x_i owned by a party b, party b calculates $\hat{x}_i = x_i + r_i$ and sends it to party $1 - b$. Now, the parties evaluate the circuit gates in topological order. To evaluate a gate g with input and output wire w_i and w_j respectively, both parties evaluate the corresponding FSS key on $\hat{x}_i$ to get secret shares of $\hat{x}_j = g^{[r_i, r_j]}(\hat{x}_i) = g(\hat{x}_i - r_i) + r_j = g(x_i) + r_j$. The parties then reconstruct these shares to get masked value $\hat{x}_j$. For the output wires, the party owning the wire subtracts the corresponding mask to get the final output value.

Protocol Structure and Security for FSS protocols. We use $\hat{(\cdot)}$ to denote masked values. Consider a function F and input x such that $y = F(x)$. Protocol for F, denoted by Π^F, has two phases Gen^F and Eval^F. Gen^F is executed in the preprocessing phase on input and output masks r^{in} and $\mathsf{r}^{\mathsf{out}}$, respectively, to produce the preprocessing material or *keys* for F made available to P_0 and P_1. The number of bits required to store the key for Π^F is called the *key size* and is denoted by $\mathsf{keysize}(\Pi^F)$. Next, Eval^F is the protocol run by P_0 and P_1 in the online phase on masked input $\hat{x} = x + \mathsf{r}^{\mathsf{in}}$ and their respective keys. At the end of Eval^F, P_0 and P_1 learn secret-shares of masked output value $\hat{y} = y + \mathsf{r}^{\mathsf{out}}$. All protocols presented in this paper have the above structure.

Security for $\Pi^F = (\mathsf{Gen}^F, \mathsf{Eval}^F)$ is defined through the following two interactions. 1) A *real interaction* in which Gen^F is executed in the preprocessing phase (with input and output masks r^{in} and $\mathsf{r}^{\mathsf{out}}$) and P_0 and P_1 execute Eval^F in the online phase with keys obtained in the preprocessing phase. This interaction happens in the presence of an adversary $\mathcal{A}$ and the environment $\mathcal{Z}$. 2) An *ideal interaction* in which P_0 and P_1 send their inputs to a functionality that computes the functionality faithfully (i.e., unmasks $\hat{x}$ to get x, computes $y = F(x)$, computes $\hat{y} = y + \mathsf{r}^{\mathsf{out}}$ and provides shares of $\hat{y}$ to P_0 and P_1). We say that protocol Π^F securely realizes function F if for every adversary $\mathcal{A}$ in the real interaction, there is an adversary $\mathcal{S}$ (called the simulator) in the ideal interaction such that no environment $\mathcal{Z}$ can distinguish between the two interactions.

2.5. Distributed Point Function (DPF)

The point function $f^{\bullet}_{\alpha,\beta} : \mathbb{U}_N \to \mathbb{G}^{\mathsf{out}}$ takes as input $x \in \mathbb{U}_N$ and outputs $\beta \in \mathbb{G}^{\mathsf{out}}$ if $x = \alpha$ and 0 otherwise. The corresponding FSS-scheme for point function $(\mathsf{Gen}^{\bullet}_n, \mathsf{Eval}^{\bullet}_n)$ is called *Distributed Point Function* [12], [13]. Notationally, we write $(k^{\bullet}_0, k^{\bullet}_1) \leftarrow \mathsf{Gen}^{\bullet}_n(1^\lambda, \alpha, \beta, \mathbb{G}^{\mathsf{out}})$ and $y_b = \mathsf{Eval}^{\bullet}_n(b, k^{\bullet}_b, x)$, for $x \in \mathbb{U}_N$. For all our protocols, it suffices to have $\mathbb{G}^{\mathsf{out}} = \{0,1\}$ and $\beta = 1$, and this allows us to leverage the construction of DPF with *early termination optimization* (that is applicable for small payloads) [13].

Theorem 1 (Cost of DPF from [13])**.** *Given PRG* $G : \{0,1\}^\lambda \to \{0,1\}^{2\lambda+2}$ *and let* $\nu = \log_2(\lambda+1)$*. When* $n > \nu$*, there exists a DPF for* $f^{\bullet}_{\alpha,1} : \mathbb{U}_N \to \{0,1\}$ *with key size* $(n-\nu)\cdot(\lambda+2) + 2\lambda$*. Number of PRG invocations in* $\mathsf{Gen}^{\bullet}_n$ *is* $2(n-\nu)$ *and in* $\mathsf{Eval}^{\bullet}_n$ *is* $n - \nu$*. When* $n \leqslant \nu$*, keysize of* 2^n *and 0 PRG invocations in* $\mathsf{Gen}^{\bullet}_n$ *and* $\mathsf{Eval}^{\bullet}_n$ *is required.*

Similar to prior FSS works [13], [64], [71], we set $\lambda = 127$ and implement the required length doubling PRG using 2 calls to AES-128 in counter mode. As previously observed [13], [64], a single AES call suffices for $\mathsf{Eval}^{\bullet}_n$ as

only half of the output is used. From here on, we refer to it as an *half-PRG call*.

2.6. Comparisons using DPF Keys

Comparison function $f^{<}_{\alpha,\beta} : \mathbb{U}_N \rightarrow \mathbb{G}^{\mathsf{out}}$ takes as input $x \in \mathbb{U}_N$ and returns $\beta \in \mathbb{G}^{\mathsf{out}}$ if $x < \alpha$ and 0 otherwise. Previous works [11], [32] used a specialized FSS-scheme called *Distributed Comparison Function* (DCF) to realize this functionality. Recent work of [64] showed that when $\mathbb{G}^{\mathsf{out}} = \{0, 1\}$, $\beta = 1$, FSS scheme for comparison function can be constructed using the DPF construction from [13].

Theorem 2 (FSS for comparison using DPF [64])**.** *There exists an algorithm* $\mathsf{Eval}^{<}_n$ *such that* $\forall x, \alpha \in \mathbb{U}_N$*:*

$$(k^{\bullet}_0, k^{\bullet}_1) \leftarrow \mathsf{Gen}^{\bullet}_n(1^{\lambda}, \alpha, 1, \{0, 1\})$$
$$\implies \mathsf{Eval}^{<}_n(0, x, k^{\bullet}_0) + \mathsf{Eval}^{<}_n(1, x, k^{\bullet}_1) = f^{<}_{\alpha,1}(x)$$

and $\mathsf{Eval}^{<}_n$ *invokes DPF half-PRG* $\max(n-\nu, 0)$ *times. Thus,* $(\mathsf{Gen}^{\bullet}_n, \mathsf{Eval}^{<}_n)$ *is an FSS-scheme for comparison function.*

Compared to DCF construction from [11] that requires a length quadrupling PRG, above construction results in $> 2\times$ lower compute cost.

3. Overview of Transformers

3.1. Architecture Overview

Transformers is a neural network architecture used commonly in natural language tasks. At a high level, a transformer architecture consists of an encoder and a decoder [70]. The encoder generates a sequence of hidden states from the given input sequence. The decoder takes the hidden states produced by the encoder and generates the output sequence. Real-world models stack multiple encoder and decoder blocks, as shown in Figure 1, to obtain high accuracy results. Further, transformers can be used in both encoder-decoder (e.g., BERT) and decoder-only mode (e.g., GPT). We discuss the key components of a single transformer block below:

Token embeddings: Transformers represent a natural language input as a sequence of tokens (e.g., each word can be represented as a token) wherein each token is a one-dimensional vector of size d_{model}. The token embedding matrix $W_e \in \mathbb{R}^{d_{model} \times N_V}$, where N_V is the vocabulary size, maps each token to its corresponding embedding vector. Further, each token is also assigned a positional encoding vector of size d_{model} that encodes the token's position in the input sequence [70]. The sum of the token embedding vector and the positional encoding vector is used as input to the transformer block.

Self-attention and multi-head attention (MHA): The self-attention mechanism helps the model attend to different parts in the input sequence. It maps a query and a set of key-value pairs to an output as follows:

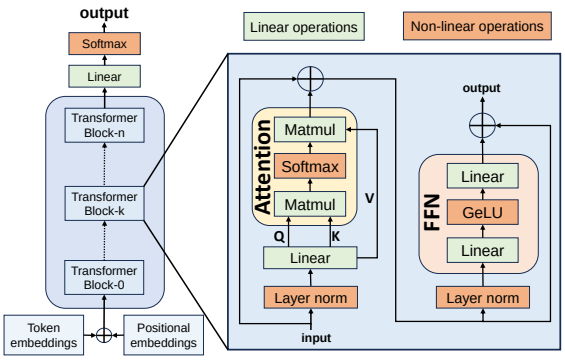

Figure 1: Architecture of a transformer neural network

$$\text{Attention (Q, K, V)} = \text{softmax}(\text{QK}^T/\sqrt{d_{model}})\text{V}$$

where $Q \in \mathbb{R}^{y \times d_{model}}$ is the query matrix and $K, V \in \mathbb{R}^{z \times d_{model}}$ are key and value matrices, respectively (here, y and z represent the length of primary and context sequence.)

The multi-head attention module consists of multiple attention heads that operate in parallel, each over $\frac{d_{model}}{num\_heads}$ in the above formulation (e.g., $num\_heads = 12$ in GPT-2). The outputs of the attention heads are concatenated and linearly transformed to obtain the MHA output.

Softmax: For a vector $\boldsymbol{x} \in \mathbb{R}^k$, define $x_{\max} = \max(x_0, x_1, \ldots, x_{k-1})$. The softmax function on $\boldsymbol{x}$ returns a vector $\boldsymbol{y} \in \mathbb{R}^k$ such that:

$$y[i] = \frac{e^{x[i]}}{\sum_{j=0}^{k-1} e^{x[j]}} = \frac{e^{x[i]-x_{\max}}}{\sum_{j=0}^{k-1} e^{x[j]-x_{\max}}}$$

Since exponentials in the first expression can get arbitrarily large, the second expression is preferred where exponential is only computed on negative values (including 0).

Feed forward network (FFN): FFN consists of two fully connected layers wherein the first layer transforms the input from dimension d_{model} to d_{ff}, and the second layer transforms it back to d_{model} (typically, $d_{ff} = 4 \times d_{model}$). FFN for a matrix $X \in \mathbb{R}^{z \times d_{model}}$ (where z is the sequence length) can be represented as:

$$\mathsf{FFN}(X) = \mathsf{GeLU}(XW_1 + b_1)W_2 + b_2$$

where $W_1 \in \mathbb{R}^{d_{model} \times d_{ff}}$ and $W_2 \in \mathbb{R}^{d_{ff} \times d_{model}}$ are the weight matrices and $b_1 \in \mathbb{R}^{d_{ff}}$, $b_2 \in \mathbb{R}^{d_{model}}$ are the bias vectors for first and second layers within FFN. GeLU is the Gaussian Error Linear Unit activation function [34].

Activation: An activation function applies a non-linear transformation element-wise to the given input vector and its output determines which of the neurons should be activated in the next layer. Popular examples of activation functions include ReLU, GeLU, tanh etc. Language models commonly use GeLU which returns a vector $\boldsymbol{y} \in \mathbb{R}^k$ for $\boldsymbol{x} \in \mathbb{R}^k$ s.t.:

$$y_i = \text{GeLU}(x_i) = \frac{x_i}{2}(1 + \text{erf}\ (\frac{x_i}{\sqrt{2}}))$$

where erf is an integral of a Gaussian [34].

Layer normalization: Layer norm is used to normalize the distribution of activations at each layer of the neural network. For a vector of real values $\boldsymbol{x} \in \mathbb{R}^k$, let $m = \sum x_i/k$

and $v = (\sum(x_i - m)^2)/k$ denote it's mean and variance, respectively. For $z_i = x_i - m$ and model parameters $\gamma, \beta \in \mathbb{R}$, the layer normalization returns a vector $\boldsymbol{y} \in \mathbb{R}^k$ s.t.:

$$y_i = \gamma \cdot \frac{x_i - m}{\sqrt{v}} + \beta = \gamma \cdot \frac{z_i}{\sqrt{\sum z_i^2/k}} + \beta \qquad (1)$$

3.2. Secure Inference of Transformers

Based on the above description and the literature on cryptographic protocols, the layers in a transformer can be classified into two categories - linear and non-linear.

Linear layers. These consist of the matrix multiplications occurring in multihead attention (MHA) and feed forward (FFN) layers. Similar to all prior works on secure inference, we work with fixed-point arithmetic. Here, multiplying two fixed-point values with precision f over integers results in a fixed-point value with implicit precision $2f$. Hence, multiplications must be followed by a truncation operation to bring the precision back to f. For the matrix multiplications over integers, we use the existing protocol [14], [32], [38] that relies on Beaver-triple like correlations generated in preprocessing phase. For truncations, as one of our contributions, we provide a significantly more efficient protocol compared to the prior work [11], [32] (see Section 4.2).

Non-linear layers. These consist of GeLU, Softmax and LayerNorm. In Section 5, we provide novel precise protocols for these non-linearities over fixed-point arithmetic that not only preserve the accuracy of transformers but also lead to efficient secure inference on transformers (Section 7).

Putting things together. For each of the layers of the transformers, we provide a secure protocol where the evaluating parties start with masked input, i.e., for an input x and random mask r^{in}, parties hold $\hat{x} = x + \mathsf{r}^{\mathsf{in}}$ and after the protocol learn masked output, i.e., $\hat{y} = y + \mathsf{r}^{\mathsf{out}}$ for output y and mask $\mathsf{r}^{\mathsf{out}}$. Given this invariant, we are able to trivially put together the secure protocols for each layer to obtain a secure protocol for inference and prove security by invoking the sequential composition theorem [17], [47].

3.3. GPU-accelerated Secure Inference

Graphics Processing Units (GPUs) support thousands of concurrent threads and provide much higher memory bandwidth compared to CPUs [2]. Therefore, GPUs are a natural fit for accelerating transformers in plaintext: (1) several linear layers (e.g., in FFN) in a transformer network involve large matrix multiplications that can be accelerated using GPUs, often by up to two orders of magnitude compared to CPUs. (2) the non-linear layers are memory intensive and hence benefit from the high memory bandwidth of GPUs. Under secure inference, the linear layers can be accelerated similar to plaintext. However, the non-linear layers require several rounds of network communication between the client and the model provider, and transfer of large pre-generated keys from CPU to GPU over the PCIe links. Therefore, communication and key transfer overheads dominate the overall time under secure inference.

We reduce the size of communication and data transfer, at the expense of some extra computation, as follows: (1) we reduce network communication with an efficient packing scheme for non-standard bitwidths. This adds extra computation for packing and unpacking values which we implement efficiently on the GPU itself. (2) we reduce the number of DPF keys needed for GeLU from two to one at the cost of one extra evaluation of the same key per element. These optimizations reduce network communication by $1.2-1.5\times$ and key transfer by $1.8\times$ over a naïve port of our CPU protocols to the GPU. Note that without these optimizations, a GPU's compute units would often remain idle. Hence, the additional computation is essentially free for SIGMA.

4. Crypto Building Blocks

Similar to ORCA [38], we design efficient protocols with multi-round online phase. Our goal is to achieve low key size, online compute and online communication while ensuring small constant round complexity. At the end of Eval^F, the evaluators learn secret-shares of masked output value $\hat{y} = y + \mathsf{r}^{\mathsf{out}}$. Now, Eval^F can be followed by a $\mathsf{reconstruct}$ to obtain the masked output value $\hat{y}$ and we denote this modified protocol by $\hat{\Pi}^F$. As the input and output masks are unknown to the evaluators, the cleartext values remain hidden from the evaluators.

We first provide a summary of protocols for multiplication, selection, and lookup tables from prior works. Then, we describe our novel FSS-based protocols for truncation and comparison. All of these are used as sub-protocols by our novel protocols for complex non-linearities (Section 5).

4.1. Protocols from Previous Works

Multiplication. For secure multiplication of two n-bit integers, [14] provides a beaver-triple based (non-interactive) FSS-protocol Π_n^{Mul} with keysize of $3n$ bits.

Select. The functionality $\mathsf{select}_n : \{0,1\} \times \mathbb{U}_N \to \mathbb{U}_N$ takes as input a selector bit s and a payload x such that $\mathsf{select}_n(s, x) = x$ if $s = 1$ and 0 otherwise. Orca [38] provides a non-interactive protocol Π_n^{select} that realizes select_n securely with keysize $4n$.

SelectLin. Let $\mathsf{selectlin}_{n,\gamma} : \{0,1\}^2 \times \mathbb{U}_N \to \mathbb{U}_N$ be a functionality parameterized by a length 4 vector of pairs of elements, $\gamma = \{(\alpha_{\mathbf{0}}, \beta_{\mathbf{0}}), (\alpha_{\mathbf{1}}, \beta_{\mathbf{1}}), (\alpha_{\mathbf{2}}, \beta_{\mathbf{2}}), (\alpha_{\mathbf{3}}, \beta_{\mathbf{3}})\}$ with $\alpha_i, \beta_i \in \mathbb{U}_N, \forall i \in [4]$. It takes as input two selector bits s_0, s_1, and a payload x, and outputs $\mathsf{selectlin}_{n,\gamma}(s_0, s_1, x) = \alpha_{2s_0+s_1} x + \beta_{2s_0+s_1}$. This functionality can be easily realized using one-time truth tables as described in [21] and results in a non-interactive protocol $\Pi_n^{\mathsf{selectlin}_\gamma}$ with keysize $8n$.

Look-up Table. The functionality $\mathsf{LUT}_{n,\ell,\boldsymbol{T}} : \mathbb{U}_N \to \mathbb{U}_L$ is parameterized by input bitwidth n, output bitwidth ℓ and a public table $\boldsymbol{T} \in \mathbb{U}_L^N$. It takes as input $x \in \mathbb{U}_N$ and returns

$T[x] \in \mathbb{U}_L$. Pika [71] provides a protocol $\Pi^{\mathsf{LUT}}_{n,\ell,\boldsymbol{T}}$ such that $\mathsf{keysize}(\Pi^{\mathsf{LUT}}_{n,\ell,\boldsymbol{T}}) = \mathsf{keysize}(\mathsf{DPF}_{n,1}) + n + 2\ell$. Online phase invokes the DPF PRG $2^{n-\nu} - 1$ times, where $\nu = \log_2(\lambda + 1)$, and communicates 2ℓ bits in 1 round.

4.2. Our Truncation Protocol

As discussed in Section 3.2, linear layers or matrix multiplication needs to be followed by an element-wise truncation to bring down precision. Our protocols for complex non-linearities also use multiple truncations. The literature considers (cheap) local truncations [41], [52], [72], [73], [76] and (expensive) faithful truncations [11], [32], [59]. While local truncations are almost free to implement, a very recent work Li et al. [46] shows that these do not satisfy standard simulation-based security and are insecure. In light of this, in this work, all our protocols for secure inference only use faithful truncations or arithmetic right shifts (ARS). Here, we provide new protocols for truncation that are much more efficient than prior FSS-based protocol from [11], [32].

ARS with guaranteed gap. We first consider the case when the input is known have a *gap* w.r.t. the bitwidth used. In particular, we require that $v \in \mathbb{U}_N$ is such that $v \in [0, 2^{n-2}) \cup [2^n - 2^{n-2}, 2^n)$. Looking ahead, within our protocols for non-linearities, this assumption holds many a times from domain knowledge.

We first use the following relation from [25] to reduce ARS to logical right shift (LRS), i.e., a reduction of shift of signed values to unsigned values. In particular, for n-bit values and shift amount f, when $v \in [0, 2^{n-2}) \cup [2^n - 2^{n-2}, 2^n)$, for $x = v + 2^{n-2}$,

$$\mathsf{ARS}_{n,f}(v) = \mathsf{LRS}_{n,f}(x) - 2^{n-f-2}$$

where $\mathsf{LRS}_{n,f}(x) = \left\lfloor \frac{x}{2^f} \right\rfloor$. Note that constraint on v implies that $x = v + 2^{n-2}$ seen as an unsigned value lies in $[0, 2^{n-1})$ which would be crucial for the optimization that we do.

Now, given the above relation, to construct a protocol for $\mathsf{ARS}_{n,f}(v)$, we construct a protocol for $\mathsf{LRS}_{n,f}(x)$ using the following lemma (also used in [11], [59]).

Lemma 1. *For $x_0 = \hat{x} \mod 2^f$ and $r_0 = \mathsf{r}^{in} \mod 2^f$,*

$$\mathsf{LRS}^{[\mathsf{r}^{in},\mathsf{r}^{out}]}_{n,f}(\hat{x}) = \mathsf{LRS}_{n,f}(\hat{x}) - \mathsf{LRS}_{n,f}(\mathsf{r}^{in}) + 2^{n-f} \cdot 1\{\hat{x} < \mathsf{r}^{in}\} - 1\{x_0 < r_0\} + \mathsf{r}^{out} \quad (2)$$

When $x \in [0, 2^{n-1})$, following observation[3] (proof in Appendix H) provides an efficient way to compute $1\{\hat{x} < \mathsf{r}^{in}\}$.

Lemma 2. *For $\hat{x} = x + \mathsf{r}^{in} \mod N$, if $x < 2^{n-1}$,*

$$1\{\hat{x} < \mathsf{r}^{in}\} = 1\{\mathsf{MSB}_n(\hat{x}) = 0\} \wedge 1\{\mathsf{MSB}_n(\mathsf{r}^{in}) = 1\}$$

We provide a formal description of our protocol for LRS for inputs with a gap in Figure 2 (security proof in Appendix I.1). Here, the term $1\{x_0 < r_0\}$ is computed using DPF-based comparison with 1-bit output to allow for smaller FSS key and lower online compute. Once the evaluators learn the masked value of this bit ($\hat{w}$), they do a local extension ($\hat{z}$). They use $\hat{z}$ and arithmetic shares of the mask ($r^{(w)}$) provided by the dealer to obtain arithmetic shares of $u = 1\{x_0 < r_0\}$. It is trivial to extend this to ARS (with the same cost) and we summarize the cost in Theorem 3.

Logical Right-Shift with Gap $\Pi^{\mathsf{GapLRS}}_{n,f}$

$\mathsf{Gen}^{\mathsf{GapLRS}}_{n,f}(\mathsf{r}^{in}, \mathsf{r}^{out})$:

1: $(k^\bullet_0, k^\bullet_1) \leftarrow \mathsf{Gen}^\bullet_f(1^\lambda, \mathsf{r}^{in} \mod 2^f, 1, \{0,1\})$
2: $c \xleftarrow{\$} \{0,1\}, r^{(w)} = \mathsf{extend}_{1,n}(c)$
3: $m = 2^{n-f} \cdot \mathsf{extend}_{1,n}(\mathsf{MSB}_n(\mathsf{r}^{in}))$
4: $r = \mathsf{r}^{out} - \mathsf{LRS}_{n,f}(\mathsf{r}^{in})$
5: $\mathsf{share}\ (r^{(w)}, m, r)$
6: For $b \in \{0,1\}$, $k_b = k^\bullet_b || r^{(w)}_b || m_b || r_b$

$\mathsf{Eval}^{\mathsf{GapLRS}}_{n,f}(b, k_b, \hat{x})$:

1: Parse k_b as $k^\bullet_b || r^{(w)}_b || m_b || r_b$
2: $\hat{w}_b = \mathsf{Eval}^{\leq}_f(b, k^\bullet_b, \hat{x} \mod 2^f) + r^{(w)}_b \mod 2$
3: $\hat{w} = \mathsf{reconstruct}(\hat{w}_b)$, $\hat{z} = \mathsf{extend}_{1,n}(\hat{w})$
4: $u_b = b\hat{z} + r^{(w)}_b - 2\hat{z}r^{(w)}_b$
5: $t_b = m_b \cdot \mathsf{extend}_{1,n}(1 - \mathsf{MSB}_n(\hat{x}))$
6: **return** $b \cdot \mathsf{LRS}_{n,f}(\hat{x}) + r_b + t_b - u_b$

Figure 2: Protocol for Logical Right-Shift with Gap

Theorem 3. *There exists a protocol $\Pi^{\mathsf{GapARS}}_{n,f}$ that realizes $\mathsf{ARS}_{n,f}$ securely for cleartext inputs in $[0, 2^{n-2}) \cup [2^n - 2^{n-2}, 2^n)$ such that $\mathsf{keysize}(\Pi^{\mathsf{GapARS}}_{n,f}) = \mathsf{keysize}(\mathsf{DPF}_{f,1}) + 3n$. The online phase requires 1 evaluation of $\mathsf{DPF}_{f,1}$ and communication of 2 bits in 1 round.*

Truncate-Reduce. $\mathsf{TR}_{n,f} : \mathbb{U}_N \to \mathbb{U}_{2^{n-f}}$ is defined as dropping the lower f bits of the n-bit input and returning the output as an $(n-f)$-bit number. It can be expressed as:

$$\mathsf{TR}_{n,f}(x) = \mathsf{LRS}_{n,f}(x) \mod 2^{n-f}$$

Note that Equation 2 for LRS does not rely on gap in inputs. Now, as the term $2^{n-f} \cdot 1\{\hat{x} < \mathsf{r}^{in}\}$ cancels out due to mod operation, we can realize truncate-reduce securely using a single comparison for $1\{x_0 < r_0\}$. We omit details and summarize cost below:

Theorem 4. *There exists a protocol $\Pi^{\mathsf{TR}}_{n,f}$ that realizes $\mathsf{TR}_{n,f}$ securely such that $\mathsf{keysize}(\Pi^{\mathsf{TR}}_{n,f}) = \mathsf{keysize}(\mathsf{DPF}_{f,1}) + 2(n-f)$. The online phase requires 1 evaluation of $\mathsf{DPF}_{f,1}$ and communicates 2 bits in 1 round.*

ARS without known gap. Let $\mathsf{SignExt}_{\ell,n} : \mathbb{U}_L \to \mathbb{U}_N$ be defined as sign extending a value in ℓ-bits to equivalent value in n-bits. When input to ARS is not known to have a gap, we express[4] $\mathsf{ARS}_{n,f}$ as $\mathsf{TR}_{n,f}$ followed by $\mathsf{SignExt}_{n-f,n}$. We

3. Similar observation was also used by [25] for their probabilistic LRS protocol that ignores the LSB correction term $1\{x_0 < r_0\}$ and referred to as MSB-to-Wrap optimization in SIRNN [58] and used in various protocols.

4. Similar approach was used in Orca [38] for stochastic truncations.

use our protocol for (faithful) truncate-reduce and replace DCF in the protocol for sign-extension from Orca [38] with DPF-based comparison. We summarize overall costs below:

Theorem 5. *There exists a protocol* $\Pi^{\mathsf{ARS}}_{n,f}$ *that realizes* $\mathsf{ARS}_{n,f}$ *securely such that* $\mathsf{keysize}(\Pi^{\mathsf{ARS}}_{n,f}) = \mathsf{keysize}(\Pi^{\mathsf{TR}}_{n,f}) + \mathsf{keysize}(\mathsf{DPF}_{n-f,1}) + 2n + 1$. *Online phase requires* 1 *evaluation each of* $\mathsf{DPF}_{f,1}$ *and* $\mathsf{DPF}_{n-f,1}$ *and communicates* $2(n-f)+4$ *bits in 3 rounds.*

Cost Comparison. In contrast, [11] gave a protocol for $\mathsf{ARS}_{n,f}$ (also used in [32]) that requires a key size of approximately $n(\lambda+2n)+f(\lambda+n)$ bits and online phase makes $2(n+f-1)$ AES calls. Concretely, for $n=64, f=12$, $\Pi^{\mathsf{GapARS}}_{n,f}$ has $17.5\times$ smaller key size and $30\times$ lower online compute, and $\Pi^{\mathsf{ARS}}_{n,f}$ has $2.5\times$ smaller key size and $3\times$ lower online compute.

4.3. Our DReLU and Comparison Protocols

For an n-bit value $x \in \mathbb{U}_N$ in 2's complement notation, derivative of ReLU or DReLU is defined as

$$\mathsf{DReLU}_n(x) = 1\{x < 2^{n-1}\} = 1 \oplus \mathsf{MSB}_n(x)$$

and the offset function of DReLU_n can be written as

$$\begin{aligned}\mathsf{DReLU}^{[\mathsf{r}^{\mathsf{in}},\mathsf{r}^{\mathsf{out}}]}_n(\hat{x}) &= \mathsf{DReLU}_n(\hat{x} - \mathsf{r}^{\mathsf{in}} \mod N) \oplus \mathsf{r}^{\mathsf{out}} \\ &= \mathsf{MSB}_n(\hat{x} - \mathsf{r}^{\mathsf{in}} \mod N) \oplus 1 \oplus \mathsf{r}^{\mathsf{out}}\end{aligned}$$

Prior FSS works [11], [14], [32], [38] provide a non-interactive protocol for DReLU that uses a DCF key for comparison and evaluates it twice during online phase. In contrast, we provide a non-interactive protocol that does a single evaluation of a DPF key for comparison. In all, we get $> 4\times$ reduction in online compute.

Our protocol builds on the logic used in CrypTFlow2 [59] for MSB computation over secret shares (in $\log n$ rounds). For $x \in \mathbb{U}_N$ such that $x = x_0 + x_1 \mod N$, $y_0 = x_0 \mod 2^{n-1}$ and $y_1 = x_1 \mod 2^{n-1}$,

$$\mathsf{MSB}_n(x) = \mathsf{MSB}_n(x_0) \oplus \mathsf{MSB}_n(x_1) \oplus 1\{y_0 + y_1 \geq 2^{n-1}\}$$

Using this above, we get

$$\begin{aligned}\mathsf{DReLU}^{[\mathsf{r}^{\mathsf{in}},\mathsf{r}^{\mathsf{out}}]}_n(\hat{x}) = \mathsf{MSB}_n(\hat{x}) \oplus \mathsf{MSB}_n(2^n - \mathsf{r}^{\mathsf{in}}) \\ \oplus 1\{2^{n-1} - y_0 - 1 < y_1\} \oplus 1 \oplus \mathsf{r}^{\mathsf{out}}\end{aligned}$$

where $y_0 = \hat{x} \mod 2^{n-1}$ and $y_1 = (2^n - \mathsf{r}^{\mathsf{in}}) \mod 2^{n-1}$.

Based on above equation, we provide a protocol for DReLU_n in Figure 3 (security proof in Appendix I.1) where we compute $1\{2^{n-1} - y_0 - 1 < y_1\}$ using a single evaluation of DPF-based comparison.

Theorem 6. Π^{DReLU}_n *(non-interactively) securely realizes* DReLU_n *with* $\mathsf{keysize}(\Pi^{\mathsf{DReLU}}_n) = \mathsf{keysize}(\mathsf{DPF}_{n-1,1}) + 1$. *Online phase requires* 1 *evaluation of* $\mathsf{DPF}_{n-1,1}$.

Comparison. To compare two values x, y, i.e., to compute $x \geqslant y$, similar to all prior works, we re-write it as $x - y \geqslant 0$ and realize it using a call to Π^{DReLU}_n.

DReLU Π^{DReLU}_n

$\mathsf{Gen}^{\mathsf{DReLU}}_n(\mathsf{r}^{\mathsf{in}}, \mathsf{r}^{\mathsf{out}})$:

1: $x_1 = 2^n - \mathsf{r}^{\mathsf{in}}$
2: $y_1 = x_1 \mod 2^{n-1}$
3: $(k^\bullet_0, k^\bullet_1) \leftarrow \mathsf{Gen}^\bullet_{n-1}(1^\lambda, y_1, 1, \{0,1\})$
4: $r = \mathsf{r}^{\mathsf{out}} \oplus \mathsf{MSB}_n(x_1) \oplus 1$
5: share r
6: For $b \in \{0,1\}$, $k_b = r_b || k^\bullet_b$

$\mathsf{Eval}^{\mathsf{DReLU}}_n(b, k_b, \hat{x})$:

1: Parse k_b as $r_b || k^\bullet_b$
2: $\hat{y} = \hat{x} \mod 2^{n-1}$
3: $t_b \leftarrow \mathsf{Eval}^{<}_{n-1}(b, k^\bullet_b, 2^{n-1} - \hat{y} - 1)$
4: **return** $b \cdot \mathsf{MSB}_n(\hat{x}) \oplus r_b \oplus t_b$

Figure 3: Protocol for DReLU.

5. Our protocols for complex non-linearities

Here, we describe our protocols for various complex non-linearities - GeLU (Section 5.1), softmax (Section 5.2), and layer normalization (Section 5.3). Finally, in Section 5.4, we discuss a few transformers-specific optimizations that allow us to compute these non-linearities over smaller tensors or smaller bitwidths in certain scenarios. Computing these non-linear functions requires efficient computation of various unary functions - GeLU, exponential, inverse, and reciprocal square root. Pika's approach to compute any arbitrary elementary function is to just look up the correct output from a table [71]. However, for an n-bit input, it requires a lookup table (LUT) of size 2^n, and computing it securely requires roughly 2^{n-7} PRG calls. In contrast, Grotto [64] uses custom splines and DPFs to realize a subset of functions required in transformers (see Section 7 for a thorough comparison).

In SIGMA, we devise function-dependent strategies to significantly reduce the size of LUTs used, while ensuring that our protocols provide good numerical approximations and hence, preserve the accuracy of transformers when run securely using our protocols. For $f = 12$ used by all our benchmarks, our protocols use LUTs of size 2^8 for GeLU and exponential, an LUT of size between 2^{13} and 2^{16} for inverse, and an LUT of size 2^{13} for reciprocal square root, independent of bitwidth n. Note that almost all our benchmarks require a bitwidth of around 50 and our techniques result in significantly smaller LUTs than Pika that are very efficient to compute securely. Moreover, our recipe for approximating reciprocal square root is general and applicable to any elementary function.

For each of the non-linearities, we describe our secure protocol as a sequence of calls to protocols described in Section 4 and security trivially holds in the simulation paradigm using sequential composition [17], [47]. While for ease of exposition, we describe our ideas for $f = 12$ that is used by all our transformer benchmarks, they can easily be

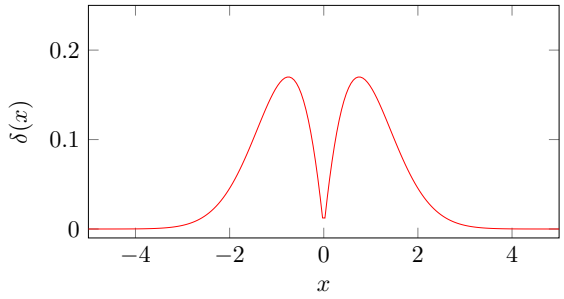

Figure 4: Plot for $\delta(x) = \mathsf{ReLU}(x) - \mathsf{GeLU}(x)$.

generalized to higher precision values by using appropriately larger LUTs.

5.1. GeLU

For a real number x, $\mathsf{GeLU}(x) = 0.5x(1 + \mathsf{erf}(x/\sqrt{2}))$ where erf is the error function [34]. Prior works, e.g., Crypten [41], Grotto [64], provide protocols for GeLU in the same threat model as ours. However, these are an order of magnitude less performant than SIGMA (Section 7.1).

Our main insight is that $\mathsf{GeLU}(x)$ is same as $\mathsf{ReLU}(x) := \mathsf{max}(x, 0)$ almost everywhere except in a small interval around 0. Let $\delta(x) = \mathsf{ReLU}(x) - \mathsf{GeLU}(x)$ (plot shown in Figure 4). Given that $\mathsf{ReLU}(x)$ can be efficiently realized using a call to DReLU and select, it suffices to efficiently compute $\delta(x)$ for x near 0. Finally, we output $\mathsf{GeLU}(x)$ as $\mathsf{ReLU}(x) - \delta(x)$. We calculate $\delta(x)$ using an LUT. However, for efficiency, we need to restrict the input domain of the LUT, while ensuring that the results are precise enough.

First, we observe that $\delta(x)$ becomes negligible outside the range $(-4, 4)$ and for precision $f = 12$, $\delta(-4) = \delta(4) = 0$. Hence, we first restrict the inputs to $(-4, 4)$ or equivalently $[-2^{f+2} + 1, 2^{f+2} - 1]$ using a *clip* operation. Formally, for n-bit values and clipping nodes A, B, we define $\mathsf{Clip}_{n,A,B}(x)$ as (i) A for $x < A$ (ii) x for $x \in [A, B]$, and (iii) B for $x > B$.

Next, we observe that $\delta(x)$ is an even function between $(-4, 4)$. Hence, it suffices to compute the LUT using the absolute value of the clipped input, that lies in $[0, 2^{f+2} - 1]$ and requires $f + 2$ bits to represent. We further reduce the size of input domain to LUT by scaling down to 6-bits of precision, retaining 8-bits of information that are used as input to the LUT to compute $\delta(x)$.

We provide a formal description of our approximation of $\mathsf{GeLU}(x)$ in Figure 5. Here, $A = -2^{f+2} + 1$ and $B = 2^{f+2} - 1$. Also, $\boldsymbol{T} \in \mathbb{U}_N^{256}$ is the table such that for all $i \in \mathbb{U}_{256}$, $T[i] = \lfloor \delta(\frac{i}{2^6}) \cdot 2^f \rfloor$. For $f = 12$, our approximation achieves[5] an ULP error of 31 which suffices to maintain PyTorch accuracy for all benchmarks as shown in Section 7.

Next, we describe how we translate the above cleartext function to secure protocols. We do a re-ordering of operations in the above description to achieve secure operations

5. We compute error by exhaustive testing on all inputs between $(-4, 4)$ as the error is 0 outside this domain.

$\mathsf{GeLU}_{n,f}(x)$:
1: $p = \mathsf{ReLU}_n(x)$
2: $c = \mathsf{Clip}_{n,A,B}(x)$
3: $a = \mathsf{Abs}_n(c)$
4: $t = \mathsf{TR}_{n,f-6}(a) \mod 256$
5: **return** $p - \mathsf{LUT}_{8,n,\boldsymbol{T}}(t)$

Figure 5: Our approximation for $\mathsf{GeLU}_{n,f}(x)$.

GeLU (CPU) $\Pi^{\mathsf{GeLUCPU}}_{n,m,f}(\hat{x})$
1: $\hat{y} \leftarrow \hat{\Pi}^{\mathsf{TR}}_{m,f-6}(\hat{x} \mod 2^m)$
2: $\hat{d} \leftarrow \hat{\Pi}^{\mathsf{DReLU}}_{m-f+6}(\hat{y})$
3: $\hat{p} \leftarrow \hat{\Pi}^{\mathsf{select}}_{m-f+6}(\hat{d}, \hat{y})$
4: $\hat{a} \leftarrow 2 \cdot \hat{p} - \hat{y}$
5: $\hat{i} \leftarrow \hat{\Pi}^{\mathsf{DReLU}}_{m-f+6}(\hat{a} - 256) \oplus 1$
6: $\hat{c} \leftarrow \hat{\Pi}^{\mathsf{select}}_{8}(\hat{i}, \hat{a} - 255 \mod 256) + 255$
7: **return** $\Pi^{\mathsf{select}}_{n}(\hat{d}, \hat{x}) - \Pi^{\mathsf{LUT}}_{8,n,\boldsymbol{T}}(\hat{c})$

Figure 6: CPU-optimized protocol for $\mathsf{GeLU}_{n,m,f}$

on lower bitwidths, resulting in lower keysize, online compute, and communication. Moreover, since the performance bottlenecks are different on CPU and GPU, we provide two different versions of the GeLU protocol. Looking ahead, for GPUs, we trade-off lower keysize and communication with higher compute compared to CPU.

5.1.1. CPU Protocol. We make the following optimizations.

Optimization 1. Since $A = -B$, it holds that $\mathsf{Abs}_n(\mathsf{Clip}_{n,A,B}(x)) = \mathsf{Clip}_{n,0,B}(\mathsf{Abs}_n(x))$. Hence, we switch the steps (2) and (3) in Figure 5 to $a = \mathsf{Abs}_n(x)$; $c = \mathsf{Clip}_{n,0,B}(a)$. This switch has 2 benefits. First, the absolute value can be calculated for free given ReLU as $\mathsf{Abs}_n(x) = 2 \cdot \mathsf{ReLU}_n(x) - x$. Second, since the input to Clip is now guaranteed to be a positive number, it can be realized by 1 comparison (with B) instead of 2 before (one each with A and B).

Optimization 2. Since the lower $f - 6$ bits are going to be discarded anyways, and do not affect the outcome of comparisons in ReLU or Clip, it is safe to perform this operation as the very first step. This reduces the bitwidth of comparisons in ReLU and Clip by $f - 6$.

Optimization 3. This applies when domain knowledge helps in restricting the inputs of GeLU to a sub-domain of $\mathbb{U}_N$. For instance, in all transformers, GeLU is always preceded by a linear layer that invokes a truncation by f after a matrix multiplication. Due to this, the effective input bitwidth of the GeLU input is $m = n - f$. Combining this with the above, the comparisons can happen over $m - (f - 6)$ bits.

Based on the above optimizations, we present our CPU-optimized protocol $\Pi^{\mathsf{GeLUCPU}}_{n,m,f}$ for $\mathsf{GeLU}_{n,m,f}$ in Figure 6,

where input/output bitwidths are n, effective input bitwidth is m, and precision is f.

Cost Analysis. $\Pi^{\mathsf{GeLUCPU}}_{n,m,f}$ requires a key size equal to the key size of 2 $\Pi^{\mathsf{DReLU}}_{m-f+6}$, 1 $\Pi^{\mathsf{LUT}}_{8,n,\boldsymbol{T}}$, 1 $\Pi^{\mathsf{TR}}_{m,f-6}$ and 3 calls to Π^{select} of bitwidths n, $m-f+6$ and 8. Online phase compute consists of a single evaluation of each of these and communication of $4(m-f)+2n+46$ bits in 6 rounds.

5.1.2. GPU Protocol. We note that the performance bottlenecks on CPU and GPU are quite different. CPU implementations are bottlenecked by compute (i.e., number of AES calls). However, once AES calls are accelerated well on GPU, performance bottlenecks become key transfer from CPU RAM to GPU memory and communication between the two parties. Thus, when creating a secure version of Figure 5 for the GPU, we focus on reducing key size and communication while tolerating a higher compute. We later argue that this trade-off results in lower runtime compared to a naïve port of the CPU protocol.

Our starting point is the protocol outlined in Figure 6. To allow computing on smaller bitwidths, we keep optimizations 2 and 3 intact. Thus, we start by computing $y = \mathsf{TR}_{m,f-6}(x \mod 2^m)$. Crucially, we let go of optimization 1, and combine ReLU and Clip differently. First, we compute DReLU bit $d = \mathsf{DReLU}(y)$. We additionally compute an interval containment bit $i = 1\{-255 \leqslant y \leqslant 255\} = \mathsf{DReLU}(y-256) - \mathsf{DReLU}(y+255)$. In doing so, we compute one more DReLU than the CPU, i.e., a total of 3. However, crucially, since all the DReLU evaluations are on y shifted by a constant, they can all use the same key. Hence, unlike GeLUCPU, this requires a *single* DPF key.

Given i and d, we compute $\mathsf{Abs}(\mathsf{Clip}(y))$ as

$$\mathsf{Abs}(\mathsf{Clip}(y)) = \begin{cases} 255 & i=0, d=0 \\ 255 & i=0, d=1 \\ -y & i=1, d=0 \\ y & i=1, d=1 \end{cases}$$

As an optimization, similar to CPU, before a selection, we first reduce y to 8 (relevant) bits and compute $\mathsf{Abs}(\mathsf{Clip}(z))$, where $z = y \mod 256$. Note that since i is computed on y, it only allows the value of z to propagate when $-255 \leqslant y \leqslant 255$. Since d already contains the sign of y, the last 8 bits of y (captured by z), suffice to correctly compute $\mathsf{Abs}(\mathsf{Clip}(y))$. For this selection based on i and d, we invoke $\hat{\Pi}^{\mathsf{selectlin}_\gamma}_8(i,d,z)$ with $\gamma = \{(\mathbf{0},\mathbf{255}),(\mathbf{0},\mathbf{255}),(-\mathbf{1},\mathbf{0}),(\mathbf{1},\mathbf{0})\}$. This gives us $c = \mathsf{Abs}(\mathsf{Clip}(z))$.

We provide the formal GPU protocol in Figure 7. We also note that unlike the CPU version, this does not require reconstructing $\mathsf{ReLU}(x)$ over $m-f-6$ bits (Step 3 in Figure 6). This is because we extract the interval containment bit needed for Clip from x and not from $\mathsf{Abs}(x)$. This allows us to save on communication as well as one round, resulting in efficient GPU implementation.

GeLU (GPU) $\Pi^{\mathsf{GeLUGPU}}_{n,m,f}(\hat{x})$

1: $\hat{y} \leftarrow \hat{\Pi}^{\mathsf{TR}}_{m,f-6}(\hat{x} \mod 2^m)$
2: $\hat{d}_b \leftarrow \Pi^{\mathsf{DReLU}}_{m-f+6}(\hat{y})$
3: $\hat{i}_b \leftarrow \Pi^{\mathsf{DReLU}}_{m-f+6}(\hat{y}+255) \oplus \Pi^{\mathsf{DReLU}}_{m-f+6}(\hat{y}-256)$
4: $(\hat{i},\hat{d}) = \mathsf{reconstruct}(\hat{i}_b,\hat{d}_b)$
5: $\hat{z} = \hat{y} \mod 256$
6: $\hat{c} \leftarrow \hat{\Pi}^{\mathsf{selectlin}_\gamma}_8(\hat{i},\hat{d},\hat{z})$
7: **return** $\Pi^{\mathsf{select}}_n(\hat{d},\hat{x}) - \Pi^{\mathsf{LUT}}_{8,n,\boldsymbol{T}}(\hat{c})$

Figure 7: GPU-optimized protocol for $\mathsf{GeLU}_{n,m,f}$. The calls to Π^{DReLU} in steps 2-3 can use same key.

Cost Analysis. $\Pi^{\mathsf{GeLUGPU}}_{n,m,f}$ requires a key size equal to the key size of 1 DPF_{m-f+5} key (for the 3 calls to $\hat{\Pi}^{\mathsf{DReLU}}_{m-f+6}$), 1 $\Pi^{\mathsf{LUT}}_{8,n,\boldsymbol{T}}$ call, 1 $\hat{\Pi}^{\mathsf{TR}}_{m,f-6}$ call and 2 calls to Π^{select} for bitwidths 8 and n. The online phase communicates $2(m-f)+2n+34$ bits in 4 rounds.

Compared to CPU protocol for $n=64, m=52, f=12$, the GPU protocol has $1.8\times$ smaller keysize, $1.3\times$ less communication, and $1.5\times$ larger number of half-PRG calls. Empirically, on a microbenchmark of 1 million GeLUs, our protocol takes about 70ms, of which 34ms is key transfer, 16ms is compute (of which about 88% is DReLU) and 20ms is communication. This is about $1.4\times$ faster than a naïve port of the CPU protocol.

5.2. Softmax

For a vector $\boldsymbol{x} \in \mathbb{R}^k$ and $x_{\max} = \max(x_0, x_1, \ldots, x_{k-1})$, softmax on $\boldsymbol{x}$ returns a vector $\boldsymbol{y} \in \mathbb{R}^k$ such that:

$$y[i] = \frac{e^{x[i]-x_{\max}}}{\sum_{j=0}^{k-1} e^{x[j]-x_{\max}}}$$

Overview. We need protocols for max, exponentiation of negative values and inverse. $x_{\max}$ can be computed using $k-1$ invocations of our protocols for comparison of 2 elements (Section 4.3) and select in $2\lceil \log_2(k) \rceil$ rounds. Now, we can subtract $x_{\max}$ from every element $x[i]$ to obtain $x[i]-x_{\max}$ and invoke the exponentiation protocol on this value to obtain $z[i]$. We can then compute $z = \sum_{j=0}^{k-1} z[i]$, invoke our protocol for inverse on z to obtain z^{-1}, and compute $y[i] = z^{-1} \cdot z[i]$ followed by truncation. We use $\Pi^{\mathsf{GapARS}}_{n,f}$ for the final truncation as $y[i] \in [0,1]$ with precision $2f$ (due to being a probability distribution) resulting in the required *gap*.

Below, we describe our novel protocols for exponential and inverse, both of which use the domain knowledge of softmax for efficiency.

5.2.1. Negative Exponential. Define $\mathsf{nExp}(x) = e^{-x}$ for $x \in \mathbb{R}^+$. We first observe that nExp is a monotonically decreasing function and for $f=12$, $x \geqslant 16$, fixed-point representation of $\mathsf{nExp}(x)$, i.e., $\lfloor e^{-x} \cdot 2^{12} \rfloor = 0$. Hence, we

Negative Exponential $\Pi^{\mathsf{nExp}}_{n,m,f}(\hat{x})$

1: $\hat{d} \leftarrow \hat{\Pi}^{\mathsf{DReLU}}_{m}((\hat{x} - 2^{16}) \mod 2^m) \oplus 1$
2: $\hat{c} \leftarrow \hat{\Pi}^{\mathsf{select}}_{16}(\hat{d}, \hat{x} - (2^{16} - 1) \mod 2^{16}) + (2^{16} - 1)$
3: $\hat{c}_1 \leftarrow \hat{\Pi}^{\mathsf{TR}}_{16,8}(\hat{c})$; $\hat{c}_0 \leftarrow \hat{c} \mod 256$
4: $\hat{t}_1 \leftarrow \hat{\Pi}^{\mathsf{LUT}}_{8,n,\boldsymbol{T_1}}(\hat{c}_1)$; $\hat{t}_0 \leftarrow \hat{\Pi}^{\mathsf{LUT}}_{8,n,\boldsymbol{T_0}}(\hat{c}_0)$
5: $\hat{t} \leftarrow \hat{\Pi}^{\mathsf{Mul}}_{n}(\hat{t}_0, \hat{t}_1)$
6: **return** $\Pi^{\mathsf{GapARS}}_{n,f}(\hat{t})$

Figure 8: Protocol for $\mathsf{nExp}_{n,m,f}$

first clip the inputs to the interval $[0, 16) \subset \mathbb{R}^+$ followed by using an LUT to compute nExp for this interval. When $x \in [0, 16)$, we need 16-bits to represent fixed-point values with precision $f = 12$. Now, directly using lookup for exponentiation would require an LUT of size 2^{16}, which is expensive.

Next, we use the technique from Seedot [29] for nExp (also used in [58]) that allows reducing one 16-bit LUT to two 8-bit LUTs. Let $c = c_1||c_0$ be the 16-bit clipped value with $f = 12$, where c_1 is upper 8-bits and c_0 is lower 8-bits. These can be calculated as $c_1 = \mathsf{TR}_{16,8}(c)$ and $c_0 = c \mod 256$. Seedot showed that

$$\left\lfloor \mathsf{nExp}\left(\frac{c}{2^{12}}\right) \cdot 2^f \right\rfloor \approx \mathsf{ARS}_{n,f}\left(\boldsymbol{T_1}[c_1] \cdot \boldsymbol{T_0}[c_0]\right)$$

where $\boldsymbol{T_1}, \boldsymbol{T_0}$ are 8-bit LUTs with n-bit values such that $T_1[i] = \lfloor \mathsf{nExp}(i/2^4) \cdot 2^f \rfloor$ and $T_0[i] = \lfloor \mathsf{nExp}(i/2^{12}) \cdot 2^f \rfloor$ for all $i \in \mathbb{U}_{2^8}$. Here, $\Pi^{\mathsf{GapARS}}_{n,f}$ suffices to perform $\mathsf{ARS}_{n,f}$ as its input is always less than 2^{2f}, leading to a gap. Compared to using 16-bit LUT, the above approach reduces online compute by $100\times$ (1022 half-PRG calls to 10 half-PRG calls including TR and ARS).

We provide a formal description of our protocol $\Pi^{\mathsf{nExp}}_{n,m,f}$ in Figure 8. Here, similar to GeLU, we introduce an additional parameter m that captures effective bitwidth and helps reduce cost when possible from domain knowledge.

5.2.2. Inverse. We calculate inverse using an LUT of carefully chosen size. It is easy to see that for a softmax of size k, the input to inverse $z \in [1, k]$. That is, it has a non-zero integer part which is also upper bounded. Hence, without losing any information, we reduce the bitwidth of input from n to $p = f + \lceil \log_2(k+1) \rceil$ retaining precision f. Next, we create an approximate input with lower precision by chopping off few lower bits[6]. In our specific case, we reduce precision to 6, creating an input with bitwdth $q = 6 + \lceil \log_2(k+1) \rceil$. Finally, we use a q-bit LUT to read the output of inverse. The protocol for inverse $\Pi^{\mathsf{Inv}}_{n,f}$ returns $\Pi^{\mathsf{LUT}}_{q,n,\boldsymbol{T}}\left(\hat{\Pi}^{\mathsf{TR}}_{p,f-6}(\hat{x} \mod 2^p)\right)$, where $\boldsymbol{T} \in \mathbb{U}^{2^q}_N$ is a table such that $T[i] = \left\lfloor 2^{f+6}/i \right\rfloor$ for all $i \in \mathbb{U}_{2^q}$.

6. While doing this in general can lose all information from the input and result in garbage result for inverse, it is still safe to do in our setting because the initial input has a meaningful lower bound.

5.3. Layer Normalization

Equation 1, Section 3.1 provides the mathematical expression for layer normalization. We note that all sub-expressions in the equation can be implemented using our existing protocols barring reciprocal square root. Below we provide an overview of our protocol for reciprocal square root and defer the details of the overall protocol and an additional optimization to Appendix G.

5.3.1. Reciprocal Square Root. While we aim to approximate the reciprocal square root using an LUT, securely computing an n-bit LUT for a large n (e.g., 50) is not efficient. So far, we have exploited two main ideas to reduce the size of LUTs significantly. Either, the function is non-zero only in a small domain (e.g., $\mathsf{GeLU}(x) - \mathsf{ReLU}(x)$, $\mathsf{nExp}(x)$) or we use domain knowledge to restrict the input domain (e.g., inverse in softmax). However, both these ideas are inapplicable here. Although reciprocal square root is a monotonically decreasing function, it only approximates to 0 for very large values. Moreover, we do not have any useful lower or upper bound on the input. Hence, our idea is to shift to a representation that allows representing a large dynamic range with a small number of bits. This is exactly what floating-point representations allow. We use domain knowledge to design a custom 13-bit floating-point representation to encode the input and use it to index an LUT. We provide a formal description in Appendix G.2.

5.4. Global Optimizations

Effective Bitwidth. In case of transformers, GeLU is always preceded by a linear layer which invokes a truncation after matrix multiplication. This means, for n bit inputs to the linear layer, the output of truncation by f lies in range $[-2^{n-f-1}, 2^{n-f-1})$. Hence, the effective bitwidth of the input to GeLU is only $m = n - f$. This allows us to perform comparisons on a smaller bitwidth of m instead of n.

Similarly, softmax is also preceded by a linear layer. As the first step of softmax is to find the $\max$ element, all the comparisons in $\max$ calculation can happen over an effective bitwidth of $m = n - f$. Then, the max element is subtracted from all the elements in the input vector before being passed to the protocol for nExp. As both input vector elements and max element have effective bitwidth of $n - f$, the input to nExp has effective bitwidth of $m = n - f + 1$.

Attention Mask. In transformer models, for input with sequence length k, the input to softmax is always a batch of k vectors of size k. In many GPT models, including those that we evaluate on, the upper triangular elements of the softmax input are *masked*, i.e., their nExp is set to 0 in the softmax computations. Hence, we can avoid calling the $\max$ and nExp protocols for the masked elements and reduce their number of calls to half.

6. Implementation

We have implemented two versions of SIGMA, one which is optimized for CPUs and the other for GPUs.

GPU. The GPU-accelerated part has around 9K lines of C++/CUDA code. For the GPU version, our starting point is Orca [38], which is currently the state-of-the-art in GPU-accelerated FSS. Similar to [38], [76], we use CUTLASS [1] to implement linear layers. We borrow Orca's ideas on AES acceleration, memory layout and payload packing to build an efficient GPU-accelerated DPF kernel. Securely realizing $\mathsf{LUT}_{n,\ell,\boldsymbol{T}}$ (Section 4.1) requires computing $\mathsf{Eval}^{\bullet}_n(b, k^{\bullet}_b, x),\ \forall x \in \mathbb{U}_N$ [71]. For this, we follow the depth-first approach of [43], while using Orca's AES kernel.

Building on our optimized kernels for DPFs and LUTs, we provide efficient GPU implementations of our protocols for GeLU, Softmax, and LayerNorm. We carefully use templating as in Orca [38] and Piranha [76] to ensure that compute happens on lower bitwidths wherever possible. In GeLU, for example, we use the fact that $\mathsf{selectlin}$ (Step 6 in Figure 7) runs on $z \in \mathbb{U}_{256}$ to run the protocol with the `uint8_t` data-type on the GPU. This helps us reduce key size, which, in turn, reduces the time to transfer keys from CPU to GPU memory.

We also *fuse* kernels wherever possible to avoid repeated accesses to GPU's global memory. We notice that our protocols often evaluate a DPF on a unary function of the input variable, or return a linear function of the DPF input and output variables. Consider the protocol for DReLU outlined in Figure 3. A naïve implementation of this protocol would first run a GPU kernel to compute $\hat{z} = 2^{n-1} - \hat{y} - 1$, invoke the DPF kernel on $\hat{z}$ to compute t_b, and then run a third kernel to compute the linear combination $b \cdot \mathsf{MSB}_n(\hat{x}) \oplus r_b \oplus t_b$ in Step 3. Together, these kernels require four loads and three stores to GPU global memory. On the other hand, a fused kernel for DReLU only requires 1 load for the input and 1 store for the output which is more efficient.

Once the compute has been accelerated, key transfer and communication dominate most of the runtime. For example, communication and key transfer consume 35% and 44% of the total runtime. To lower communication, we observe that our protocols operate on non-powers-of-2 bitwidths. Hence, there is often a gap between the size of a ring element and the corresponding C++ data-type e.g., `uint64_t`. In some cases, this gap can be quite large, e.g., secure inference for BERT-large communicates ring elements with bitwidths 50 in linear layers, 44 in GeLU, and 39 in Softmax. Therefore, we *pack* elements before transmitting them over the network to achieve significant communication savings over a naïve implementation that transmits standard data-types[7]. For example, we reduce communication by 35% for BERT-large.

We provide kernels for packing and unpacking elements of arbitrary bitwidths on the GPU as a part of SIGMA. For packing, we make each GPU thread responsible for writing a segment of 8 bytes of data. It uses the size of the ring elements it needs to communicate to fetch the elements that belong to its segment. It also performs any shifts necessary to accommodate 'partial' elements in its segment (e.g. to pack only the first 8 bits of an element). This allows us to ensure that the packing is tight.

Since packing and unpacking require additional computation, we are implicitly trading lower communication for more computation. We find that GPUs can effectively handle this additional computation due to their high degree of parallelism. However, the cost of packing and unpacking values on CPUs overshadows the benefit of lower communication. Therefore, we do not use this optimization in our CPU implementation.

CPU. The CPU code is written with 7500 lines of C++ and uses OMP for multithreading, Eigen [30] for matrix multiplications, and `cryptoTools` [61] for PRG implementations that use native x86 AES instructions.

SyTorch Frontend. We also develop SYTORCH, a C++-based frontend, for specifying the architecture of machine learning models to be used for secure inference. It allows users to express models in a PyTorch-like high-level description and run them with various *backends*, e.g., fixed-point cleartext or SIGMA's protocols for CPUs/GPUs. We provide a sample SYTORCH code snippet in Figure 11 (Appendix C). Given a SyTorch model and an input, the outputs from all backends are bitwise equivalent.

The SyTorch program is compiled to a control flow graph (CFG), which is automatically transformed, e.g., relevant truncations are inserted and effective bitwidths are set (Section 5.4). The final optimized graph is then interpreted. For each operation occurring in the graph, the corresponding protocol is executed.

Model	# GeLU	# Softmax	# Rsqrt	# blocks	h	d_{model}
BERT-tiny	131072	512	512	2	2	128
BERT-base	4718592	18432	3072	12	12	768
BERT-large	12582912	49152	6144	24	16	1024
GPT-2	4718592	18432	3072	12	12	768
GPT-Neo	25165824	49152	6144	24	16	2048

Table 1: Number of scalar GeLU, 128-length Softmax, scalar reciprocal square roots, blocks, attention heads h and embedding length d_{model} for transformers.

	AES or half-PRG		Comm. (Bytes)		Key Size (KB)	
	Grotto	SIGMA	Grotto	SIGMA	Grotto	SIGMA
Gelu	753	78	320	58	1.97	1.43
Inverse	1092	254	320	36	1.97	0.17
Rsqrt	4215	1840	320	106	1.97	1.93

Table 2: SIGMA has lower computation, communication, and key size than Grotto [64].

7. Evaluation

We provide empirical results to justify the following claims. SIGMA's protocols for complex non-linearities are up to $10\times$ more efficient than (FSS-based) Grotto [64] (Table 2) and up to $38\times$ more efficient than CrypTen [41] (Table 3). For end-to-end evaluation of transformers, CrypTen [41] is our primary baseline. CrypTen is the state-of-the-art that supports the operations present in transformers, works in 2PC with preprocessing model, and provides

7. While Orca packs 1 or 2-bit values, we support packing for *all* non-powers-of-2 bitwidths in SIGMA, providing benefit in all our protocols. While reporting improvements, we use the baseline that packs 1 or 2-bit values but uses standard data-types for rest.

Model	GeLU						Softmax						LayerNorm					
	Time (s)			Communication (GB)			Time (s)			Communication (GB)			Time (s)			Communication (GB)		
	CT	S-CPU	S-GPU	CT	S-CPU	S-GPU	CT	S-CPU	S-GPU	CT	S-CPU	S-GPU	CT	S-CPU	S-GPU	CT	S-CPU	S-GPU
BERT-tiny	0.27	0.06	0.007	0.10	0.01	0.003	0.71	0.09	0.02	0.09	0.01	0.005	0.60	0.03	0.03	0.003	0.004	0.002
BERT-base	4.59	3.76	0.25	3.45	0.25	0.16	7.53	4.42	0.44	3.27	0.37	0.26	4.31	0.67	0.25	0.11	0.15	0.11
BERT-large	11.50	9.84	0.66	9.19	0.66	0.42	17.35	11.94	1.13	8.72	1.00	0.69	8.75	1.78	0.55	0.29	0.40	0.30
GPT-2	4.47	3.76	0.25	3.45	0.25	0.16	6.89	2.76	0.27	3.27	0.19	0.13	3.94	0.69	0.25	0.11	0.15	0.11
GPT-Neo	20.35	20.35	1.33	18.38	1.69	0.86	16.33	7.55	0.66	8.72	0.50	0.36	8.91	3.39	0.80	0.57	0.80	0.60

Table 3: SIGMA outperforms CrypTen (GPU) on GeLU, Softmax and LayerNorm. CT denotes CrypTen, and S-CPU and S-GPU stand for SIGMA running on CPU and GPU respectively.

GPU-accelerated implementations. But unlike SIGMA that provides standard 2PC security, CrypTen provides imperfect security because it uses fast but insecure local truncations [46]. End-to-end secure inference of transformers with SIGMA is $11.5 - 19.4\times$ faster than CrypTen and requires $8.4 - 11.6\times$ lower communication (Table 4). We observe that SIGMA running on CPUs is already faster than CrypTen running on GPUs. Furthermore, SIGMA on GPUs is up to an order of magnitude faster than SIGMA running on CPU (Figure 9). Finally, we show that SIGMA scales efficiently with the number of model parameters (Figure 10) by evaluating on GPT models with up to 13 billion parameters, while CrypTen crashes on the larger models.

Models and datasets: We evaluate BERT-tiny, BERT-base, and BERT-large models [69] on the SST2, QNLI, and MRPC classification tasks from GLUE benchmark [74]. These models have 4.4 million, 110 million, and 330 million parameters respectively. The prior work of Iron [33] also considers these models and datasets. For a billion parameter model, we evaluate the GPT-Neo-1.3B model at huggingface [4] (225 thousand downloads in last month) on the challenging Lambada dataset [54], which has next-word-prediction tasks. We evaluate GPT-2 with 124 million parameters from huggingface (downloaded 15 million times within last month) on Lambada as well. These models use GeLU and Softmax in abundance (Table 1). We also report the number of reciprocal square roots arising because of layer normalizations. Prior works have observed that these non-linearities are the performance bottlenecks in secure inference of transformers [33], [45]. Following Iron [33], we evaluate all models on inputs of sequence length[8] 128. We set the precision $f = 12$, and the bitwidths to be large enough so that SIGMA's accuracy matches that of 32-bit floating-point PyTorch (Appendix A). In particular, for BERT-tiny a bitwidth of 37 suffices, whereas the other models require a bitwidth of 50 or 51.

Hardware platform: We evaluate on two machines connected via LAN with 9.4 Gbps bandwidth and 0.05 ms ping time. Each machine has 1 TB RAM, an A6000 GPU with 46GB GPU memory, and an AMD Epyc 7742 processor. Evaluation of SIGMA running on CPUs uses 4 threads.

7.1. Non-linearities

We show our performance improvements for GeLU, Softmax, and layer normalization over the baselines.

8. We evaluate other sequence lengths in Appendix E.

7.1.1. Comparison with Grotto. Grotto [64] is a recent work that provides FSS-based protocols for GeLU, inverse (that arises in softmax), and reciprocal square root (that arises in layer normalization). Table 2 shows that, for each of these functions, SIGMA beats Grotto in all aspects: computation, communication, and key size. Since the source code of Grotto is unavailable, we cannot evaluate it on our setup. However, the communication and the key size are independent of the setup. The compute cost of FSS-based protocols like Grotto and SIGMA is heavily dominated by PRG calls, and we use these as a proxy for the computation overheads.

7.1.2. Comparison with Orca. Orca [38] is the state-of-the-art in GPU-accelerated FSS and it proposes the recipe of using 2PC floating-point protocols [57] for complex non-linearities like Softmax. The communication overheads of this approach are severe and would require 7 GB (for BERT-tiny) to 1.1 TB (for GPT-Neo) of communication for evaluating GeLU and Softmax layers. In contrast, SIGMA's communication is between 20 MB and 4 GB (Table 4).

7.1.3. Comparison with CrypTen. We compare SIGMA (both CPU and GPU implementations) and CrypTen by measuring their latency and communication in evaluating GeLU, Softmax and LayerNorm (Table 3). For GeLU and Softmax, SIGMA's communication is an order of magnitude lower than CrypTen. Due to this, SIGMA's protocols running on CPUs outperform CrypTen on GPUs on all transformers. For LayerNorm, CrypTen's communication is low because of its use of local truncation. However, our protocols for reciprocal square root is more efficient and our runtimes for LayerNorm on CPUs are $2.6-20\times$ better. Furthermore, with GPU acceleration, SIGMA outperforms CrypTen by at least $10\times$ for all three non-linearities on all transformers. Finally, the lower communication of SIGMA running on GPUs (vs. CPUs) is due to communication packing (Section 6).

7.2. Transformers

We evaluate SIGMA on end-to-end[9] transformer inference to show that it beats CrypTen in latency and communication, GPU acceleration is helpful for SIGMA, and SIGMA scales well to larger models.

9. The preprocessing costs are not included for both CrypTen and SIGMA; we describe SIGMA's preprocessing cost in Appendix B.

Model	Time (s)			Communication (GB)	
	CrypTen	SIGMA	Speedup	CrypTen	SIGMA
BERT-tiny	1.71	0.09	19.4×	0.20	0.02
BERT-base	21.55	1.84	11.7×	8.34	0.99
BERT-large	54.53	4.73	11.5×	23.36	2.63
GPT-2	20.45	1.61	12.7×	8.34	0.82
GPT-Neo	108.30	7.43	14.6×	46.89	4.02

Table 4: SIGMA vs CrypTen on end-to-end inference.

7.2.1. Comparison with CrypTen. Table 4 shows the performance of transformer models with CrypTen and SIGMA, both running on GPUs. There are two factors at play here. 1) SIGMA is using secure but more expensive truncations which take more time and communication than CrypTen's local truncations, and 2) SIGMA's protocols for non-linearities have massive improvements over CrypTen (Section 7.1.3). Overall, for end-to-end transformer inference, SIGMA outperforms CrypTen by $11.5 - 19.4\times$ in latency and $8.4 - 11.6\times$ in communication.

7.2.2. GPU acceleration. Figure 9 shows the speedups of SIGMA running on CPUs and GPUs over CrypTen as the baseline. We observe that for end-to-end transformer inference, SIGMA running on CPUs is always faster than CrypTen running on GPUs. SIGMA's protocols for GPUs are an order of magnitude faster compared to their CPU counterparts for all models except BERT-tiny, which is too small to leverage GPUs effectively.

7.3. Scaling to larger models

To evaluate how well SIGMA would scale to larger models, we increase the number of parameters by increasing the number of blocks, the embedding length, and the number of attention heads (Appendix D). We evaluate SIGMA on GPT models with 1.3, 2.7, 7, and 13 billion parameters in Figure 10. SIGMA scales efficiently to larger models and SIGMA running on GPUs is able to perform inference of a 13 billion parameter GPT model within 30 seconds. In contrast, CrypTen overflows GPU memory on the 7 billion and 13 billion parameter models and crashes.

8. Related Work

Secure inference (with MPC or with other techniques like TEEs [68] or FHE [27]) has a vast literature and we don't attempt to survey it. Here, we focus on works related to transformers, GPU acceleration of MPC, and FSS.

After the success of large models like GPT3/GPT3.5 with 175 billion parameters, there are ongoing efforts to reduce the cost and latency of inference by using smaller models [6], [66], [67]. For example, phi-1 outperforms GPT-3.5 models while using only 1.3 billion parameters [31]. Another approach to reduce the latency of secure inference involves replacing complex non-linearities that are expensive in MPC with simple non-linearities. The simple approximations significantly impact accuracy but, at least for BERT class models, this accuracy loss can be recovered by further

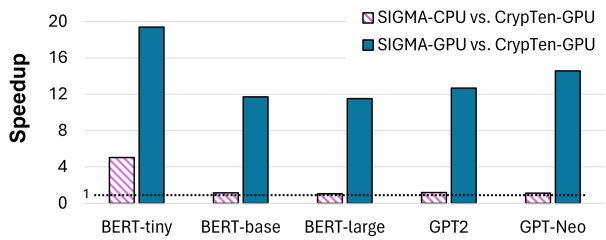

Figure 9: SIGMA CPU and GPU speedups over CrypTen.

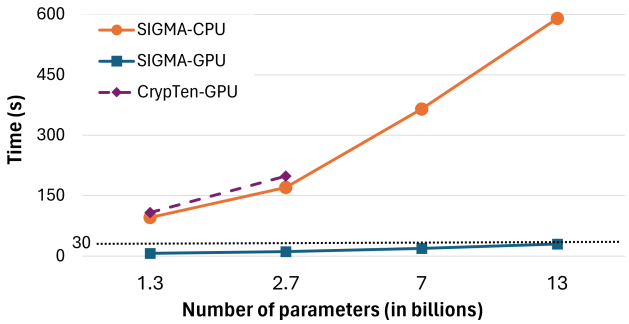

Figure 10: SIGMA scaling to larger models. CrypTen runs out of memory for 7B and 13B models.

retraining of the simplified models [49]. THE-X [19], MPCFormer [45] and Privformer [7] use simple non-linearities. In contrast, Iron [33], CrypTen [41], and SIGMA use precise approximations of complex non-linearities and there is no accuracy loss. Recent pipelining optimizations have improved the performance of CrypTen by up to 13% [75] and such optimizations can benefit SIGMA as well.

There are several works that focus on accelerating secure inference with GPUs, but to support CNNs and not transformers. CryptGPU accelerates 3-party secure inference with GPUs [65]. Piranha is a general framework that supports various number of parties [76]. Delphi performs a network architecture search to navigate performance-accuracy tradeoffs. GForce uses custom training approaches to improve inference efficiency [53]. Beyond inference, Visor [56] focuses on video analytics and general protocols like Yao's garbled circuits have also been accelerated with GPUs [36].

Several recent works consider 2PC in the preprocessing model based on FSS techniques. The work of [14] initiated this study and showed how to construct 2PC protocols for any computation comprising of gates for which FSS constructions exist for the corresponding offset gate. The work of [11] provides various FSS protocols for functions occurring in fixed-point arithmetic, while [32], [63], [71], [77] provides specialized FSS protocols for ML operations. The works of [63] and [38] accelerate FSS protocols on GPUs while [5] and [64] consider FSS protocols for various elementary functions such as sigmoid, GeLU, and so on.

9. Conclusion

We build SIGMA, the first system for FSS-based secure inference of transformers. To this end, we build novel proto-

cols for GeLU, Softmax, and layer normalization. For end-to-end transformer inference, SIGMA satisfies standard 2PC security, matches PyTorch accuracy, while being an order of magnitude faster than the baselines. Although we have focused on GeLU activations, as prior work evaluates on models that use them, the same techniques generalize to construct efficient protocols for other activations such as sigmoid, SiLU, etc. Similar to all prior works on secure inference of transformers, SIGMA focuses on semi-honest security and we leave security against malicious adversaries [18], [20], [25], [26], [40], [44] for future work.

References

[1] "CUTLASS," https://github.com/NVIDIA/cutlass.

[2] "Nvidia a100 tensor core gpu architecture," https://images.nvidia.com/aem-dam/en-zz/Solutions/data-center/nvidia-ampere-architecture-whitepaper.pdf.

[3] "GPT-2," https://huggingface.co/gpt2, 2023.

[4] "GPT Neo," https://huggingface.co/docs/transformers/model_doc/gpt_neo, 2023.

[5] A. Agarwal, S. Peceny, M. Raykova, P. Schoppmann, and K. Seth, "Communication efficient secure logistic regression," *IACR Cryptol. ePrint Arch.*, p. 866, 2022. [Online]. Available: https://eprint.iacr.org/2022/866

[6] L. A. Agrawal, A. Kanade, N. Goyal, S. K. Lahiri, and S. K. Rajamani, "Guiding language models of code with global context using monitors," 2023.

[7] Y. Akimoto, K. Fukuchi, Y. Akimoto, and J. Sakuma, "Privformer: Privacy-preserving transformer with mpc," in *EuroS&P*, 2023.

[8] D. Beaver, "Efficient multiparty protocols using circuit randomization," in *CRYPTO*, '91.

[9] R. Bendlin, I. Damgård, C. Orlandi, and S. Zakarias, "Semi-homomorphic encryption and multiparty computation," in *EUROCRYPT*, 2011.

[10] F. Boemer, R. Cammarota, D. Demmler, T. Schneider, and H. Yalame, "MP2ML: a mixed-protocol machine learning framework for private inference," in *ARES*, 2020.

[11] E. Boyle, N. Chandran, N. Gilboa, D. Gupta, Y. Ishai, N. Kumar, and M. Rathee, "Function secret sharing for mixed-mode and fixed-point secure computation," in *EUROCRYPT*, 2020.

[12] E. Boyle, N. Gilboa, and Y. Ishai, "Function secret sharing," in *EUROCRYPT*, 2015.

[13] ——, "Function secret sharing: Improvements and extensions," in *CCS*, 2016.

[14] ——, "Secure computation with preprocessing via function secret sharing," in *TCC*, 2019.

[15] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," 2020.

[16] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M. T. Ribeiro, and Y. Zhang, "Sparks of artificial general intelligence: Early experiments with gpt-4," 2023.

[17] R. Canetti, "Security and Composition of Multiparty Cryptographic Protocols," *J. Cryptology*, 2000.

[18] N. Chandran, D. Gupta, S. L. B. Obbattu, and A. Shah, "SIMC: ML Inference Secure Against Malicious Clients at Semi-Honest Cost," in *USENIX Security Symposium*, 2022.

[19] T. Chen, H. Bao, S. Huang, L. Dong, B. Jiao, D. Jiang, H. Zhou, J. Li, and F. Wei, "THE-X: privacy-preserving transformer inference with homomorphic encryption," in *ACL*, 2022.

[20] A. P. K. Dalskov, D. Escudero, and M. Keller, "Secure evaluation of quantized neural networks," *PoPETs*, 2020.

[21] I. Damgård, J. B. Nielsen, M. Nielsen, and S. Ranellucci, "The tinytable protocol for 2-party secure computation, or: Gate-scrambling revisited," in *CRYPTO*, 2017.

[22] I. Damgård and S. Zakarias, "Constant-overhead secure computation of boolean circuits using preprocessing," in *TCC*, 2013.

[23] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," 2019.

[24] J. Doerner and A. Shelat, "Scaling ORAM for secure computation," in *CCS*, 2017.

[25] D. Escudero, S. Ghosh, M. Keller, R. Rachuri, and P. Scholl, "Improved primitives for MPC over mixed arithmetic-binary circuits," in *CRYPTO*, 2020.

[26] T. K. Frederiksen, T. P. Jakobsen, and J. B. Nielsen, "Faster maliciously secure two-party computation using the gpu," in *SCN*, 2014.

[27] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. E. Lauter, M. Naehrig, and J. Wernsing, "Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy," in *ICML*, 2016.

[28] O. Goldreich, S. Micali, and A. Wigderson, "How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority," in *STOC*, 1987.

[29] S. Gopinath, N. Ghanathe, V. Seshadri, and R. Sharma, "Compiling kb-sized machine learning models to tiny iot devices," in *PLDI*, 2019.

[30] G. Guennebaud and B. Jacob, "Eigen v3," http://eigen.tuxfamily.org, 2010.

[31] S. Gunasekar, Y. Zhang, J. Aneja, C. C. T. Mendes, A. D. Giorno, S. Gopi, M. Javaheripi, P. Kauffmann, G. de Rosa, O. Saarikivi, A. Salim, S. Shah, H. S. Behl, X. Wang, S. Bubeck, R. Eldan, A. T. Kalai, Y. T. Lee, and Y. Li, "Textbooks are all you need," 2023.

[32] K. Gupta, D. Kumaraswamy, N. Chandran, and D. Gupta, "Llama: A low latency math library for secure inference," in *PETS*, 2022.

[33] M. Hao, H. Li, H. Chen, P. Xing, G. Xu, and T. Zhang, "Iron: Private inference on transformers," in *NeurIPS*, 2022.

[34] D. Hendrycks and K. Gimpel, "Bridging nonlinearities and stochastic regularizers with gaussian error linear units," *CoRR*, vol. abs/1606.08415, 2016. [Online]. Available: http://arxiv.org/abs/1606.08415

[35] Z. Huang, W. jie Lu, C. Hong, and J. Ding, "Cheetah: Lean and fast secure two-party deep neural network inference," in *USENIX Security Symposium*, 2022.

[36] N. Husted, S. Myers, A. Shelat, and P. Grubbs, "Gpu and cpu parallelization of honest-but-curious secure two-party computation," in *ACSAC*, 2013.

[37] Y. Ishai, E. Kushilevitz, S. Meldgaard, C. Orlandi, and A. Paskin-Cherniavsky, "On the power of correlated randomness in secure computation," in *TCC*, 2013.

[38] N. Jawalkar, K. Gupta, A. Basu, N. Chandran, D. Gupta, and R. Sharma, "Orca: Fss-based secure training with gpus," Cryptology ePrint Archive, Paper 2023/206, 2023.

[39] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, "Gazelle: A low latency framework for secure neural network inference," in *USENIX Security Symposium*, 2018.

[40] M. Keller, "MP-SPDZ: A versatile framework for multi-party computation," in *CCS*, 2020.

[41] B. Knott, S. Venkataraman, A. Hannun, S. Sengupta, M. Ibrahim, and L. van der Maaten, "CrypTen: Secure multi-party computation meets machine learning," in *NeurIPS*, 2021.

[42] N. Kumar, M. Rathee, N. Chandran, D. Gupta, A. Rastogi, and R. Sharma, "CrypTFlow: Secure tensorflow inference," in *IEEE S&P*, 2020.

[43] M. Lam, J. Johnson, W. Xiong, K. Maeng, U. Gupta, M. Rhu, H.-H. S. Lee, V. J. Reddi, G.-Y. Wei, D. Brooks, and E. Suh, "GPU-based Private Information Retrieval for On-Device Machine Learning Inference," *CoRR*, vol. abs/2301.10904, 2023.

[44] R. Lehmkuhl, P. Mishra, A. Srinivasan, and R. A. Popa, "Muse: Secure inference resilient to malicious clients," in *USENIX Security Symposium*, 2021.

[45] D. Li, H. Wang, R. Shao, H. Guo, E. Xing, and H. Zhang, "MPCFormer: Fast, performant and private Transformer inference with MPC," in *ICLR*, 2023.

[46] Y. Li, Y. Duan, Z. Huang, C. Hong, C. Zhang, and Y. Song, "Efficient 3PC for Binary Circuits with Application to Maliciously-Secure DNN Inference," in *USENIX Security Symposium*, 2023.

[47] Y. Lindell, "How to simulate it – a tutorial on the simulation proof technique," *Tutorials on the Foundations of Cryptography*, 2017.

[48] J. Liu, M. Juuti, Y. Lu, and N. Asokan, "Oblivious Neural Network Predictions via MiniONN Transformations," in *CCS*, 2017.

[49] N. W. Ming, Z. Wang, C. Liu, R. S. M. Goh, and T. Luo, "MA-BERT: Towards matrix arithmetic-only BERT inference by eliminating complex non-linear functions," in *ICLR*, 2023.

[50] P. Mishra, R. Lehmkuhl, A. Srinivasan, W. Zheng, and R. A. Popa, "Delphi: A cryptographic inference service for neural networks," in *USENIX Security Symposium*, 2020.

[51] P. Mohassel and P. Rindal, "ABY^3: A Mixed Protocol Framework for Machine Learning," in *CCS*, 2018.

[52] P. Mohassel and Y. Zhang, "SecureML: A System for Scalable Privacy-Preserving Machine Learning," in *IEEE S&P*, 2017.

[53] L. K. L. Ng and S. S. M. Chow, "Gforce: Gpu-friendly oblivious and rapid neural network inference," in *USENIX Security Symposium*, 2021.

[54] D. Paperno, G. Kruszewski, A. Lazaridou, N. Q. Pham, R. Bernardi, S. Pezzelle, M. Baroni, G. Boleda, and R. Fernandez, "The LAMBADA dataset: Word prediction requiring a broad discourse context," in *ACL*, 2016.

[55] A. Patra, T. Schneider, A. Suresh, and H. Yalame, "ABY2.0: Improved Mixed-Protocol secure Two-Party computation," in *USENIX Security Symposium*, 2021.

[56] R. Poddar, G. Ananthanarayanan, S. Setty, S. Volos, and R. A. Popa, "Visor: Privacy-preserving video analytics as a cloud service," in *USENIX Security Symposium*, 2020.

[57] D. Rathee, A. Bhattacharya, R. Sharma, D. Gupta, N. Chandran, and A. Rastogi, "SecFloat: Accurate Floating-Point meets Secure 2-Party Computation," in *IEEE S&P*, 2022.

[58] D. Rathee, M. Rathee, R. K. K. Goli, D. Gupta, R. Sharma, N. Chandran, and A. Rastogi, "SIRNN: A math library for secure inference of RNNs," in *IEEE S&P*, 2021.

[59] D. Rathee, M. Rathee, N. Kumar, N. Chandran, D. Gupta, A. Rastogi, and R. Sharma, "CrypTFlow2: Practical 2-Party Secure Inference," in *CCS*, 2020.

[60] M. S. Riazi, C. Weinert, O. Tkachenko, E. M. Songhori, T. Schneider, and F. Koushanfar, "Chameleon: A hybrid secure computation framework for machine learning applications," in *ASIACCS*, 2018.

[61] P. Rindal, "cryptoTools," https://github.com/ladnir/cryptoTools.

[62] B. D. Rouhani, M. S. Riazi, and F. Koushanfar, "DeepSecure: Scalable Provably-Secure Deep Learning," in *DAC*, 2018.

[63] T. Ryffel, D. Pointcheval, and F. Bach, "ARIANN: Low-interaction privacy-preserving deep learning via function secret sharing," in *PETS*, 2022.

[64] K. Storrier, A. Vadapalli, A. Lyons, and R. Henry, "Grotto: Screaming fast $(2+1)$-pc for $\mathbb{Z}_{2^n}$ via (2, 2)-dpfs," Cryptology ePrint Archive, Paper 2023/108, 2023.

[65] S. Tan, B. Knott, Y. Tian, and D. J. Wu, "Cryptgpu: Fast privacy-preserving machine learning on the GPU," in *IEEE S&P*, 2021.

[66] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto, "Stanford alpaca: An instruction-following llama model," https://github.com/tatsu-lab/stanford_alpaca, 2023.

[67] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample, "Llama: Open and efficient foundation language models," 2023.

[68] F. Tramèr and D. Boneh, "Slalom: Fast, verifiable and private execution of neural networks in trusted hardware," in *ICLR*, 2019.

[69] I. Turc, M. Chang, K. Lee, and K. Toutanova, "Well-read students learn better: The impact of student initialization on knowledge distillation," *CoRR*, vol. abs/1908.08962, 2019.

[70] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017.

[71] S. Wagh, "Pika: Secure Computation using Function Secret Sharing over Rings," *PoPETs*, 2022.

[72] S. Wagh, D. Gupta, and N. Chandran, "SecureNN: 3-party secure computation for neural network training," *PoPETs*, 2019.

[73] S. Wagh, S. Tople, F. Benhamouda, E. Kushilevitz, P. Mittal, and T. Rabin, "Falcon: Honest-majority maliciously secure framework for private deep learning," *PoPETs*, 2021.

[74] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "GLUE: A multi-task benchmark and analysis platform for natural language understanding," in *ICLR*, 2019.

[75] Y. Wang, R. Rajat, and M. Annavaram, "Mpc-pipe: an efficient pipeline scheme for secure multi-party machine learning inference," *CoRR*, vol. abs/2209.13643, 2022.

[76] J.-L. Watson, S. Wagh, and R. A. Popa, "Piranha: A GPU Platform for Secure Computation," in *USENIX Security Symposium*, 2022.

[77] P. Yang, Z. L. Jiang, S. Gao, J. Zhuang, H. Wang, J. Fang, S. Yiu, and Y. Wu, "Fssnn: Communication-efficient secure neural network training via function secret sharing," Cryptology ePrint Archive, Paper 2023/073, 2023.

[78] A. Yao, "How to Generate and Exchange Secrets (Extended Abstract)," in *FOCS*, 1986.

Appendix A. Accuracy Results

Model	Dataset	Train Size	Val Size	PyTorch Acc	SIGMA Acc	BW
BERT-tiny	SST2	67k	872	81.19%	81.42%	37
	MRPC	3.7k	408	72.54%	72.79%	37
	QNLI	105K	5463	81.64%	81.73%	37
BERT-base	SST2	67k	872	90.59%	90.25%	50
	MRPC	3.7k	408	84.31%	83.82%	50
	QNLI	105K	5463	88.72%	89.03%	50
BERT-large	SST2	67k	872	88.99%	88.99%	50
	MRPC	3.7k	408	78.67%	78.92%	50
	QNLI	105K	5463	92.23%	92.31%	50
GPT2	Lambada	-	5153	32.46%	33.28%	50
GPT-Neo	Lambada	-	5153	57.46%	57.81%	51

Table 5: For different models and datasets, we show the size of the training set (BERT models need finetuning), the size of validation set on which accuracy is measured, the accuracy of PyTorch floating-point, SIGMA's accuracy, and the bitwidth BW used by SIGMA to get this accuracy.

Appendix B. Preprocessing cost

We use a dealer to generate FSS keys and transfer them to the machines performing secure inference. Since the dealer has been accelerated with GPUs, the time to generate the keys is small (even smaller than the secure inference time) and the bulk of the preprocessing time goes in transferring the keys from the dealer machines (Table 6). Note that CPU key size is roughly $1.25\times$ larger than the GPU key size for the models in Table 6, due to differences such as the protocols for GeLU (Section 5.1).

Model	Key size (GB)	Generation time (s)	Transfer time (s)	Online time (s)
BERT-tiny	0.32	0.06	0.27	0.09
BERT-base	16.69	1.43	14.20	1.84
BERT-large	45.06	3.75	38.35	4.73
GPT2	14.17	1.26	12.06	1.61
GPT-Neo	75.57	6.25	64.32	7.43

Table 6: For different models, we show the size of FSS keys, the time taken by the dealer to generate them, the time to transfer them on the network, and online time of SIGMA.

Appendix C. Sample SyTorch Code

```
TransformerBlock(u64 n_heads, u64 n_embd)
{
    attn = new MultiHeadAttention<T>(n_heads,
                                     n_embd);
    ffn = new FFN<T>(n_embd, 4*n_embd);
    ln0 = new LayerNorm<T>(n_embd);
    ln1 = new LayerNorm<T>(n_embd);
}

Tensor<T> &_forward(Tensor<T> &input)
{
    auto &ln0_out = ln0->forward(input);
    auto &attn_out = attn->forward(ln0_out);
    auto &attn_ip = add(attn_out, input);
    auto &ln1_out = ln1->forward(attn_ip);
    auto &ffn_out = ffn->forward(ln1_out);
    auto &ffn_out_add = add(ffn_out, attn_ip);
    return ffn_out_add;
}
```

Figure 11: SYTORCH code for a GPT-2 Transformer block.

Appendix D. Large Model details

# Parameters	# blocks	h	d_{model}
1.3B	24	16	2048
2.7B	32	20	2560
7B	32	32	4096
13B	40	40	5120

Table 7: Number of parameters, blocks, attention heads h and embedding length d_{model} for the larger transformers.

Appendix E. Sequence Length

We evaluate SIGMA on input token sequences of lengths between 64 and 1024 in Table 8. For reference, the lengths for inputs in the Lambada dataset are below 180. The speedups of SIGMA over CrypTen don't vary much with sequence length. As sequence length increases, the number of GeLUs increases linearly but the compute of softmax increases super-linearly. A sequence length of k requires evaluating k softmax operations with inputs of length k.

Sequence length	Time (s)		Comm (GB)	
	CrypTen	SIGMA	CrypTen	SIGMA
64	14.22	0.96	3.92	0.37
128	20.45	1.61	8.34	0.82
256	36.68	3.26	21.11	1.98
512	85.75	8.01	63.73	5.29
1024	269.06	23.17	228.97	15.92

Table 8: Secure inference of GPT2 with SIGMA and CrypTen with varying sequence length.

Appendix F.
FSS Correctness and Security

Definition 2 (FSS: Correctness and Security [12], [13])**.** *Let $\mathcal{G} = \{g\}$ be a function family, $P_{\mathcal{G}} = \{\hat{g}\}$ be the set of descriptions of functions in $\mathcal{G}$, and* Leak *be a function specifying the allowable leakage about $\hat{g}$. When* Leak *is omitted, it is understood to output only $\mathbb{G}^{\mathsf{in}}$ and $\mathbb{G}^{\mathsf{out}}$. We say that* (Gen, Eval) *as in Definition 1 is an FSS scheme for $\mathcal{G}$ (with respect to leakage* Leak*) if it satisfies the following.*

- **Correctness:** *For all $\hat{g} \in P_{\mathcal{G}}$ describing $g : \mathbb{G}^{\mathsf{in}} \to \mathbb{G}^{\mathsf{out}}$, and every $x \in \mathbb{G}^{\mathsf{in}}$, if $(k_0, k_1) \leftarrow \mathsf{Gen}(1^\lambda, \hat{g})$ then* $\Pr\left[\mathsf{Eval}(0, k_0, x) + \mathsf{Eval}(1, k_1, x) = g(x)\right] = 1$.
- **Security:** *For each $b \in \{0, 1\}$ there is a PPT algorithm Sim_b (simulator), such that for every sequence $(\hat{g}_\lambda)_{\lambda \in \mathbb{N}}$ of polynomial-size function descriptions from $\mathcal{G}$ and polynomial-size input sequence x_λ for g_λ, the outputs of the following* Real *and* Ideal *experiments are computationally indistinguishable:*
  - Real_λ*:* $(k_0, k_1) \leftarrow \mathsf{Gen}(1^\lambda, \hat{g}_\lambda)$*; Output k_b.*
  - Ideal_λ*: Output* $\mathsf{Sim}_b(1^\lambda, \mathsf{Leak}(\hat{g}_\lambda))$.

Appendix G.
Layer Normalization

The functionality of layer normalization, as defined in Section 5.3, calls reciprocal square root with variance of the input vector as an input. Our protocol for reciprocal square root (Appendix G.2) makes use of the protocol for interval lookup (Appendix G.1). Finally, we provide the overall optimized protocol for layer normalization in Appendix G.3.

G.1. Interval Lookup

Let $\boldsymbol{p}, \boldsymbol{q} \in \mathbb{U}_N^k$ be arrays defining k disjoint intervals $[p[i], q[i]] \forall i \in [k]$, constrained with $p[i+1] = q[i] \forall i \in [k-1]$, $p[0] = 0$ and $q[k-1] = 2^n - 1$. Let $\boldsymbol{v} \in \mathbb{U}_L^k$ be a payload array. We define the functionality $\mathsf{IntervalLookup}_{n, \mathbb{U}_L, \boldsymbol{p}, \boldsymbol{q}, \boldsymbol{v}} : \mathbb{U}_N \to \mathbb{U}_L$ which returns $v[i]$ when $x \in [p[i], q[i]]$ for some $i \in [k]$. Since this functionality is equivalent to a 0-degree spline, we use the protocol for splines from Grotto [64] to implement this. Even though the protocol invokes DPF evaluation k times, they significantly reduce the number of half PRG calls compared to nk using the memoization technique, which caches the intermediate seeds in DPF tree to be reused in subsequent evaluations. We omit details and directly summarize the costs of the protocol:

Theorem 7. *Let $\ell = \lceil \log_2(|\mathbb{G}|) \rceil$ and $\boldsymbol{p}, \boldsymbol{q} \in \mathbb{U}_N^k, \boldsymbol{v} \in \mathbb{G}^k$ be arrays of size k. There exists a protocol $\Pi^{\mathsf{IntervalLookup}}_{n, \mathbb{G}, \boldsymbol{p}, \boldsymbol{q}, \boldsymbol{v}}$ which securely realizes $\mathsf{IntervalLookup}_{n, \mathbb{G}, \boldsymbol{p}, \boldsymbol{q}, \boldsymbol{v}}$ such that $\mathsf{keysize}(\Pi^{\mathsf{IntervalLookup}}_{n, \mathbb{G}, \boldsymbol{p}, \boldsymbol{q}, \boldsymbol{v}}) = \mathsf{keysize}(\mathsf{DPF}_{n,1}) + 3\ell$. In the online phase, the protocol requires k memoized evaluations of $\mathsf{DPF}_{n,1}$ and communication of 4ℓ bits in 1 round.*

G.2. Reciprocal Square Root

For bitwidth n, input precision f^{in} and output precision f^{out}, we define the function $\mathsf{RecSqrt}_{n, f^{\mathsf{in}}, f^{\mathsf{out}}}$ to be the approximation of the reciprocal square root of a fixed-point number $x \in \mathbb{U}_N$ with scale f^{in}. It returns a fixed-point number $y \in \mathbb{U}_N$ with scale f^{out}, i.e., $\mathsf{uint}_n(y) \approx \sqrt{2^{f^{\mathsf{in}}}/x} \cdot 2^{f^{\mathsf{out}}}$.

As discussed in Section 5.3.1, since the inputs of reciprocal square root occurring in layer normalization are unconstrained, to get a small LUT, we first convert the input to a custom floating point representation. This allows us to represent a large dynamic range using only a small number of bits. A similar protocol for converting fixed-point numbers to IEEE 32-bit floating-point numbers was provided by Orca [38].

A floating-point representation has a sign bit, exponent bits, and mantissa bits. Taking inspiration from the `bfloat16` datatype which is being extensively used in ML, we also use a 7-bit mantissa. As we are only interested in non-zero positive n-bit integers with $n \leq 64$, a 6-bit exponent suffices and we don't need a sign bit. This 13-bit index is used to look-up the fixed-point output.

Let $x \in \mathbb{U}_N$ be the input to RecSqrt. We convert the integer representation of x to float-like representation and input precision f^{in} would be handled in the LUT later. Let $m \in \mathbb{U}_{128}, e \in \mathbb{U}_{64}$ represent the mantissa and exponent of the floating point representation of x. So, it must hold that:

$$\mathsf{uint}_n(x) \approx 2^{\mathsf{uint}_6(e)} \cdot (1 + \frac{\mathsf{uint}_7(m)}{128}) \tag{3}$$

From here on, we suppress $\mathsf{uint}(\cdot)$ whenever it is clear from context. Let $k \in \mathbb{U}_{64}$ be a number such that $2^{k-1} \leq x < 2^k$. As $1 \leq (1 + m/128) < 2$, it holds that $2^e \leq 2^e \cdot (1 + m/128) < 2^{e+1}$ and hence, we can set $e = k - 1$. To calculate m, we plug $e = k - 1$ in Equation 3:

$$x \approx 2^{k-1} \cdot (1 + \frac{m}{128})$$
$$\implies m \approx \frac{x \cdot 128}{2^{k-1}} - 128 = \frac{x \cdot 2^{n-k}}{2^{n-8}} - 128$$

Let $u = 2^{n-k} \in \mathbb{U}_N$. As $x < 2^k$, $x \cdot 2^{n-k} < 2^n$ and can be encoded in n bits. So, we can approximate m as:

$$\begin{aligned} m &\approx \mathsf{TR}_{n,n-8}(x \cdot u) - 128 \mod 128 \\ &= \mathsf{TR}_{n,n-8}(x \cdot u) \mod 128 \end{aligned} \tag{4}$$

To securely calculate $e = k - 1$ and $u = 2^{n-k}$, we can use the protocol for interval lookup (Appendix G.1). Let $\mathbb{G} = \mathbb{U}_{2^{13}} \times \mathbb{U}_{2^n}$. Let $\boldsymbol{p}, \boldsymbol{q} \in \mathbb{U}_N^n, \boldsymbol{v} \in \mathbb{G}^n$ be arrays s.t. $p[0] = 0$, $q[0] = 1$, $v[0] = (0, 2^{n-1})$, and $\forall i \in [1, n-1]$:

$$p[i] = q[i-1] + 1,\ q[i] = 2^{i+1} - 1,\ v[i] = (i, 2^{n-i-1})$$

Then, it trivially holds that:

$$(\mathsf{extend}_{6,13}(e), u) = \mathsf{IntervalLookup}_{n, \mathbb{G}, \boldsymbol{p}, \boldsymbol{q}, \boldsymbol{v}}(x)$$

Finally, we can calculate m using Equation 4 and concatenate e to get the required floating point representation as:

$$p = m || e = \mathsf{extend}_{7,13}(m) \cdot 2^6 + \mathsf{extend}_{6,13}(e)$$

Reciprocal Square Root $\Pi^{\mathsf{RecSqrt}}_{n,f^{\mathsf{in}},f^{\mathsf{out}}}(\hat{x})$

1: $(\hat{e}_b, \hat{u}_b) \leftarrow \Pi^{\mathsf{IntervalLookup}}_{n,\mathbb{G},\boldsymbol{p},\boldsymbol{q},\boldsymbol{v}}(\hat{x})$
2: $\hat{u} = \mathsf{reconstruct}(\hat{u}_b)$
3: $\hat{t} \leftarrow \hat{\Pi}^{\mathsf{Mul}}_n(\hat{x}, \hat{u})$
4: $\hat{m}_b \leftarrow \Pi^{\mathsf{TR}}_{n,n-8}(\hat{t}) \mod 128$
5: $\hat{p}_b \leftarrow \mathsf{extend}_{7,13}(\hat{m}_b) \cdot 2^6 + \hat{e}_b$
6: $\hat{p} = \mathsf{reconstruct}(\hat{p}_b)$
7: **return** $\Pi^{\mathsf{LUT}}_{13,n,\boldsymbol{T}}(\hat{p})$

Figure 12: Protocol for $\mathsf{RecSqrt}_{n,f^{\mathsf{in}},f^{\mathsf{out}}}$

Note that local extension suffices in case of m, as the result is being multiplied by 2^6, due to which wrap error vanishes. Now, we construct the required 13-bit look-up table. Let $\boldsymbol{T} \in \mathbb{U}_N^{2^{13}}$ be a table such that for all $i \in \mathbb{U}_{2^{13}}$, $i = m||e$ where $m \in \mathbb{U}_{128}$ and $e \in \mathbb{U}_{64}$, we have:

$$q = 2^e \cdot (1 + \frac{m}{128}), T[i] \quad = \left\lfloor \sqrt{2^{f^{\mathsf{in}}}/q} \cdot 2^{f^{\mathsf{out}}} \right\rfloor \mod N$$

Based on the above discussion and using the table $\boldsymbol{T}$, we describe the protocol $\Pi^{\mathsf{RecSqrt}}_{n,f^{\mathsf{in}},f^{\mathsf{out}}}$ in Figure 12.

G.3. Overall Protocol for Layer Normalization

Naïve Protocol. A protocol for layer normalization for fixed-point numbers can be implemented as follows. We first locally add the elements of the vector $\boldsymbol{x}$, locally multiply the result with $\lfloor 2^f/k \rfloor$ and truncate to get m. Then, we locally subtract m from each element in $\boldsymbol{x}$ to get $\boldsymbol{z}$. We then use a beaver-like protocol to compute the sum of squares of the elements in $\boldsymbol{z}$ and call it s. Note that s has precision $2f$. Hence, we truncate by f. Next, we locally multiply the result with $\lfloor 2^f/k \rfloor$ and again truncate by f to get the variance v. We then use the protocol $\Pi^{\mathsf{RecSqrt}}_{n,f,f}$ (Section 5.3.1) to calculate the fixed-point number corresponding to the reciprocal square root of v, which we securely multiply with each element of $\boldsymbol{z}$ followed by truncation. Finally, we multiply the result with γ, truncate and locally add β.

Optimization. As s is truncated and divided by k before eventually being passed to $\Pi^{\mathsf{RecSqrt}}_{n,f,f}$, we can avoid the truncation and division by k in the protocol by setting $f^{\mathsf{in}} = 2f + \log_2(k)$ while invoking the reciprocal square root protocol. Note that even though fixed-point precision is an integer, here we can use real valued precision as the protocol $\mathsf{RecSqrt}_{n,f^{\mathsf{in}},f^{\mathsf{out}}}$ doesn't impose any restriction on the input precision f^{in} and it is only handled while computing the entries of the LUT.

Based on the above discussion, we provide the protocol $\Pi^{\mathsf{LayerNorm}}_{n,k,f}$ for layer normalization in Figure 13. To avoid invoking reciprocal square root on 0 we add 1 to s in line 6. Here, we note that as the elements of $\boldsymbol{p}$ have an absolute value less than $\sqrt{k}$ (with precision $2f$), leading to a gap, we can use $\Pi^{\mathsf{GapARS}}_{n,f}$ to perform this truncation cheaply. Similarly, as the model weight γ is a number

Layer Normalization $\Pi^{\mathsf{LayerNorm}}_{n,k,f}(\hat{\boldsymbol{x}}, \hat{\gamma}, \hat{\beta})$

1: $\hat{y} = \lfloor 2^f/k \rfloor \cdot \sum_{i=0}^{k-1} \hat{x}[i]$
2: $\hat{m} \leftarrow \hat{\Pi}^{\mathsf{ARS}}_{n,f}(\hat{y})$
3: $\hat{\boldsymbol{z}} = \hat{\boldsymbol{x}} - \hat{m}$
4: $\hat{s}_b = \sum_{i=0}^{k-1} \Pi^{\mathsf{Mul}}_n(\hat{z}[i], \hat{z}[i])$
5: $\hat{s} = \mathsf{reconstruct}(\hat{s}_b)$
6: $\hat{t} \leftarrow \hat{\Pi}^{\mathsf{RecSqrt}}_{n,2f+\log_2(k),f}(\hat{s}+1)$
7: $\hat{\boldsymbol{p}} \leftarrow \hat{\Pi}^{\mathsf{Mul}}_n(\hat{\boldsymbol{z}}, \hat{t})$
8: $\hat{\boldsymbol{q}} \leftarrow \hat{\Pi}^{\mathsf{GapARS}}_{n,f}(\hat{\boldsymbol{p}})$
9: $\hat{\boldsymbol{u}} \leftarrow \hat{\Pi}^{\mathsf{Mul}}_n(\hat{\boldsymbol{q}}, \hat{\gamma})$
10: $\hat{\boldsymbol{v}}_b \leftarrow \Pi^{\mathsf{GapARS}}_{n,f}(\hat{\boldsymbol{u}}) + b \cdot \hat{\beta}$
11: **return** $\hat{\boldsymbol{v}}_b$

Figure 13: Protocol for $\mathsf{LayerNorm}_{n,k,f}$

with small magnitude and multiplication with elements of $\boldsymbol{q}$ (bounded by $\sqrt{k}$ in precision f) results in elements with a gap, $\Pi^{\mathsf{GapARS}}_{n,f}$ can be used to truncate vector $\boldsymbol{u}$ as well.

Appendix H. Proof of Lemma 2

To calculate $1\{\hat{x} < \mathsf{r}^{\mathsf{in}}\}$, consider four cases:

1) *Case 1:* $\mathsf{MSB}_n(\hat{x}) = 1$ *and* $\mathsf{MSB}_n(\mathsf{r}^{in}) = 0$.
   Since $\hat{x} \geq 2^{n-1} > \mathsf{r}^{\mathsf{in}}$, $1\{\hat{x} < \mathsf{r}^{\mathsf{in}}\} = 0$ follows trivially.
2) *Case 2:* $\mathsf{MSB}_n(\hat{x}) = 0$ *and* $\mathsf{MSB}_n(\mathsf{r}^{in}) = 1$.
   Since $\hat{x} < 2^{n-1} \leq \mathsf{r}^{\mathsf{in}}$, $1\{\hat{x} < \mathsf{r}^{\mathsf{in}}\} = 1$ follows trivially.
3) *Case 3:* $\mathsf{MSB}_n(\hat{x}) = \mathsf{MSB}_n(\mathsf{r}^{in}) = 0$.
   As $x < 2^{n-1}$ and $\mathsf{r}^{\mathsf{in}} < 2^{n-1}$, $x + \mathsf{r}^{\mathsf{in}} < 2^n \implies \hat{x} = x + \mathsf{r}^{\mathsf{in}} \mod 2^n = x + \mathsf{r}^{\mathsf{in}} \geq \mathsf{r}^{\mathsf{in}} \implies 1\{\hat{x} < \mathsf{r}^{\mathsf{in}}\} = 0$.
4) *Case 3:* $\mathsf{MSB}_n(\hat{x}) = \mathsf{MSB}_n(\mathsf{r}^{in}) = 1$.
   As $x < 2^{n-1}$ and $2^{n-1} \leq \mathsf{r}^{\mathsf{in}} < 2^n$, $x + \mathsf{r}^{\mathsf{in}} \in [2^{n-1}, 2^n + 2^{n-1})$. But as $\hat{x} \geq 2^{n-1}$, $x + \mathsf{r}^{\mathsf{in}} < 2^n$. Hence, $\hat{x} = x + \mathsf{r}^{\mathsf{in}} \mod 2^n = x + \mathsf{r}^{\mathsf{in}} \geq \mathsf{r}^{\mathsf{in}} \implies 1\{\hat{x} < \mathsf{r}^{\mathsf{in}}\} = 0$.

Hence, $1\{\hat{x} < \mathsf{r}^{\mathsf{in}}\} = 1\{\mathsf{MSB}_n(\hat{x}) = 0 \text{ and } \mathsf{MSB}_n(\mathsf{r}^{\mathsf{in}}) = 1\} = \mathsf{MSB}_n(\mathsf{r}^{\mathsf{in}}) \cdot (1 - \mathsf{MSB}_n(\hat{x}))$.

Appendix I. Security Proofs

Let $\mathsf{Sim}^{<}_n$ be the simulator for the FSS-scheme of comparison function from Theorem 2. As we use $\mathsf{Gen}^{\bullet}_n$ from [13] directly in this FSS-scheme, Definition 2 implies that the security of the FSS-scheme for comparison trivially follows from the security of DPF construction of [13].

I.1. DReLU

For $b \in \{0, 1\}$, let $\mathsf{Sim}^{\mathsf{DReLU}}_b$ be the simulator for the protocol Π^{DReLU}_n. It is given the input $\hat{x} \in \mathbb{U}_N$ and output $u_b \in \{0, 1\}$. It simulates the view of party b, by simulating the message $r_b||k^{\bullet}_b$ from dealer by following these steps:

1) Set $\hat{y} = \hat{x} \mod 2^{n-1}$
2) Invoke $\mathsf{Sim}_n^<$ to simulate the DPF key$k^\bullet_{b,\mathsf{sim}}$
3) Set $t_{b,\mathsf{sim}} \leftarrow \mathsf{Eval}^<_{n-1}(b, k^\bullet_{b,\mathsf{sim}}, 2^{n-1} - \hat{y} - 1)$
4) Set $r_{b,\mathsf{sim}} = b \cdot \mathsf{MSB}_n(\hat{x}) \oplus u_b \oplus t_{b,\mathsf{sim}}$.
5) Output $r_{b,\mathsf{sim}} || k^\bullet_{b,\mathsf{sim}}$.

I.2. LRS with Gap

For $b \in \{0,1\}$, let $\mathsf{Sim}_b^{\mathsf{GapLRS}}$ be the simulator for the protocol $\Pi^{\mathsf{GapLRS}}_{n,f}$. It is given the input $\hat{x} \in \mathbb{U}_N$ and output $y_b \in \mathbb{U}_N$. It simulates the view of party b, by simulating the message $k^\bullet_b || r^{(w)}_b || m_b || r_b$ from dealer and $\hat{w}_{1-b}$ from the other party, by following these steps:

1) Sample $r^{(w)}_{b,\mathsf{sim}}, \hat{w}_{1-b,\mathsf{sim}} \xleftarrow{\$} \{0,1\}$.
2) Invoke $\mathsf{Sim}_f^<$ to simulate DPF keys $k^\bullet_{b,\mathsf{sim}}$
3) Set $\hat{w}_{b,\mathsf{sim}} = \mathsf{Eval}^<_f(b, k^\bullet_{b,\mathsf{sim}}, \hat{x} \mod 2^f) \oplus r^{(w)}_{b,\mathsf{sim}}$
4) Set $\hat{w}_{\mathsf{sim}} = \hat{w}_{b,\mathsf{sim}} \oplus \hat{w}_{1-b,\mathsf{sim}}$, $\hat{z}_{\mathsf{sim}} = \mathsf{extend}_{1,n}(\hat{w}_{\mathsf{sim}})$
5) Set $u_{b,\mathsf{sim}} = b\hat{z}_{\mathsf{sim}} + r^{(w)}_{b,\mathsf{sim}} - 2\hat{z}_{\mathsf{sim}} r^{(w)}_{b,\mathsf{sim}}$
6) Sample $m_{b,\mathsf{sim}} \xleftarrow{\$} \mathbb{U}_N$.
7) Set $t_{b,\mathsf{sim}} = m_{b,\mathsf{sim}} \cdot \mathsf{extend}_{1,n}(1 - \mathsf{MSB}_n(\hat{x}))$
8) Set $r_{b,\mathsf{sim}} = y_b - b \cdot \mathsf{LRS}_{n,f}(\hat{x}) - t_{b,\mathsf{sim}} + u_{b,\mathsf{sim}}$
9) Output $k^\bullet_{b,\mathsf{sim}} || r^{(w)}_{b,\mathsf{sim}} || m_{b,\mathsf{sim}} || r_{b,\mathsf{sim}}$ and $\hat{w}_{1-b,\mathsf{sim}}$.