DarKnight: An Accelerated Framework for Privacy and Integrity Preserving Deep Learning Using Trusted Hardware

Hanieh Hashemi
University of Southern California
Los Angeles, California, US
Hashemis@usc.edu

Yongqin Wang
University of Southern California
Los Angeles, California, US
Yongqin@usc.edu

Murali Annavaram
University of Southern California
Los Angeles, California, US
Annavara@usc.edu

ABSTRACT

Privacy and security-related concerns are growing as machine learning reaches diverse application domains. The data holders want to train or infer with private data while exploiting accelerators, such as GPUs, that are hosted in the cloud. Cloud systems are vulnerable to attackers that compromise the privacy of data and integrity of computations. Tackling such a challenge requires unifying theoretical privacy algorithms with hardware security capabilities. This paper presents DarKnight, a framework for large DNN training while protecting input privacy and computation integrity. DarKnight relies on cooperative execution between trusted execution environments (TEE) and accelerators, where the TEE provides privacy and integrity verification, while accelerators perform the bulk of the linear algebraic computation to optimize the performance. In particular, DarKnight uses a customized data encoding strategy based on matrix masking to create input obfuscation within a TEE. The obfuscated data is then offloaded to GPUs for fast linear algebraic computation. DarKnight's data obfuscation strategy provides provable data privacy and computation integrity in the cloud servers. While prior works tackle inference privacy and cannot be utilized for training, DarKnight's encoding scheme is designed to support both training and inference.

KEYWORDS

data privacy, trusted execution environment, Intel SGX, deep learning, data encoding, neural networks

ACM Reference Format:

Hanieh Hashemi, Yongqin Wang, and Murali Annavaram. 2021. DarKnight: An Accelerated Framework for Privacy and Integrity Preserving Deep Learning Using Trusted Hardware. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '21)*, October 18–22, 2021, Virtual Event, Greece. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3466752.3480112>

1 INTRODUCTION

The need for protecting data privacy in deep learning is growing rapidly in many areas such as health care [5], autonomous vehicles [81], finance [32], communication technologies [20].

This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

MICRO '21, October 18–22, 2021, Virtual Event, Greece
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8557-2/21/10.
<https://doi.org/10.1145/3466752.3480112>

Data holders in many of these application domains, however, are not machine learning experts. They rely on machine learning as a service (MLaaS) platforms such as Microsoft Azure ML [44], Google AI platform [26], Amazon ML [3] to fulfill their ML needs. These services incorporate ML accelerators such as GPUs for high performance and provide easy to use ML runtimes to enable data holders to quickly set up their models and train. While these platforms lower the steep learning curve, they exacerbate the users' concern regarding data privacy. Typically it is not the case that the entire cloud service provider is untrustworthy, but a subset of cloud servers may be compromised by hackers and other adversaries as has been demonstrated in recent attacks [17, 79]. Hence, our work is focused on protecting against a scenario where a subset of machines may be compromised, rather than a wholly untrusted cloud provider where every machine from the cloud provider is compromised and colluding. Our goal is to protect data privacy and computational integrity while still enabling the use of untrusted cloud systems.

There are various algorithmic approaches to protect data privacy, such as Homomorphic Encryption libraries [22, 24, 37, 43], Secure Multi-Party Computing (MPC) [50, 51, 66, 74], Differential Privacy [1, 19, 70] and using Trusted Execution Environments [54, 72]. Each of these methods provides a different privacy guarantee and comes at different cost [46], as we explain in the next section.

This work proposes DarKnight, a framework for accelerating privacy and integrity preserving deep learning while using untrusted accelerators. DarKnight uses a unique collaborative computing model between the hardware-supported trusted execution environments (TEE) and GPU accelerators to tackle both privacy and performance challenges. The training/inference data from a client is only made visible within the TEE which prevents its visibility to an adversary. DarKnight uses a novel data encoding strategy based on matrix masking to obfuscate input data within the TEE before the data is allowed to leave the TEE. The obfuscated data is then offloaded to GPUs to accelerate DNN's linear computations on the encoded data. Computing solely within TEEs can provide data privacy, by blocking access to TEE memory from intruders. Even though recent enhancements to TEEs support large memory workloads [64], TEE-enabled CPUs still do not support massive data parallel computations that accelerators such as GPUs and TPUs support. Therefore, DarKnight distributes compute-intensive linear operations to GPUs. DarKnight's usage of TEEs is limited to protecting the privacy of data through a customized matrix masking, verifying the integrity of computations, and performing non-linear operations (ReLU, Maxpool).

Table 1 shows the relative speedup for various forward and backward propagation operations in the VGG16 DNN model [67]

executing on a GPU (Nvidia GeForce GTX 1080 Ti GPUs) relative to a baseline model executing on Intel SGX (Coffee Lake E-2174G 3.80GHz CPU). Significant speedup can be achieved by offloading linear operations to GPUs. In the rest of the paper, we describe DarkKnight’s approach to collaboratively execute between TEE and GPUs to achieve high performance without compromising privacy. We also describe how DarkKnight can provide computational integrity by verifying computations performed on the GPUs.

Table 1: Speedup in GPU relative to SGX in VGG16 Training on ImageNet. The baseline is implemented fully on Intel SGX

Operations	Linear Ops	Maxpool Time	Relu Time	Total
Forward Pass	126.85	11.86	119.60	119.03
Backward Propagation	149.13	5.47	6.59	124.56

This paper makes the following major contributions:

- We design a data privacy-preserving framework for DNN training and inference that uses a novel matrix masking scheme to encode data within the TEE.
- For high performance DNN training this work offloads linear computations on encoded data to be executed on GPUs. The novelty of the encoding scheme is such that linear computations on encoded data can be decoded within the TEE to accurately extract the necessary result.
- We extend our encoding strategy to guarantee privacy even in the presence of colluding GPUs.
- We provide a low overhead mechanism for verifying computation integrity.
- We implemented DarkKnight using an Intel SGX-enabled CPU to perform matrix masking and non-linear DNN operations while using an Nvidia GPUs to accelerate linear operations. We observe an average of 6.5 x performance improvement for different DNN models.

To the best of our knowledge, this is the first work that uses TEE-GPU collaboration for **training** large DNNs on private data while providing computational integrity and **provable** data privacy.

The rest of the paper is organized as follows. In Section 2, we discuss related work and background. Section 3 describes the system setting and DarkKnight overview. We elaborate the encoding scheme in Section 4. In Section 5 privacy theorem is provided. Implementation and experimental results are presented in Section 6 and Section 7. In Section 8, we draw the conclusion.

2 BACKGROUND AND RELATED WORK

2.1 Intel SGX

TEEs such as ARMTrustZone [2], Intel SGX [11], and Sanctum [12] provide a hardware-assisted secure execution environment where data privacy and computational integrity of the user’s application is guaranteed by the hardware. Some of the cloud providers including IBM and Microsoft Azure have already equipped their cloud systems with Intel SGX [71]. In this work, we also utilize Intel SGX for data privacy purposes. Intel SGX provides three important features: remote attestation, local attestation, and sealing that plays an important role in our scheme. TEEs generally provide a

limited amount of secure memory that is tamper-proof. SGX provides 128 MB as the enclave memory in its current implementation. Although recent enhancements to SGX have relaxed this memory limit, the performance of SGX still suffers when enclaves use very large memory due to the Merkle-tree based encryption and versioning requirements [11, 64]. Furthermore, TEEs are CPU-based. Hence, there are not as well suited for massively parallel matrix operations that are common in DNNs. Although there are different methods proposed to extend the TEEs beyond CPUs, none of them are available in the marketplace yet [73, 77]. While some types of side-channel attacks have been performed on SGX [8, 27, 56, 75], many of these attacks are being fixed actively [76]. As such SGX side channel leakage is outside the scope of this work.

2.2 Related Work

There are a variety of approaches for protecting input and model privacy and computation integrity during DNN training and inference. These methods provide different privacy guarantees [46]. In this section, we briefly explain the most common methods. *Fully Homomorphic encryption (FHE)* techniques encrypt input data and then perform operations directly on encrypted data. They usually provide a high theoretical privacy guarantee on data leakage, albeit with a significant performance penalty, and hence are rarely used in training DNNs. *Secure multi-party computing (MPC)* is another approach, where multiple servers may use custom data exchange protocols to protect input data. They mostly use secret sharing schemes and have super-linear overhead for communication as the number of sharers and colluding entities grows. An entirely orthogonal approach is to use *differential privacy (DP)*, which protects individual users’ information through probabilistic guarantees by inserting noise signals to different parts of the computation. The tradeoff between utility and privacy is a challenge in this line of work. Recently, TEEs attracted attention in ML domain for their privacy and integrity properties [4, 30, 48, 49, 55, 57, 59]. Among TEE-based approaches, [72] introduced Slalom an *inference* framework that uses TEE-GPU collaboration to protect data privacy and integrity. However, as stated in their work their model was not designed for training DNNs. We elaborate on these reasons in Section 7.2. In another line of research, custom hardware accelerators are designed for security [33, 73]. These accelerator-based approaches use DNN-specific optimizations to reduce the protected memory footprint to improve efficiency. *Instahide* [34] combines multiple images from a private dataset, and merges them with a public image set, and uses a sign flip function on pixels as random noise parameters. This method processes the encoded data without any decoding. However, privacy guarantees are purely empirical. That is why in [10] authors designed an attack to break this empirical guarantee. DarkKnight provides a strong cryptography guarantee for privacy protection instead of relying just on empirical quantification which may be compromised. In Table 2, we compare some of these approaches based on their privacy and integrity guarantees, and their applications.

3 DARKNIGHT

System Structure: Our system model for learning is shown in Figure 1. The client is the data owner who wishes to train a model

Table 2: Comparison of applications and security guarantees of various prior techniques on neural networks' security (○ means method doesn't support that feature and ● means it supports the feature)

Method	Training	Inference	DP	MPC	HE	TEE	Data Privacy	Model Privacy(Client)	Model Privacy(Server)	Integrity	GPU Acceleration	Large DNNs
SecureNN [74]	●	●	○	●	○	○	●	●	●	○	●	○
Chiron [35]	●	●	○	○	○	○	●	●	●	●	○	○
MSP [36]	●	●	○	○	○	○	●	●	●	●	○	○
Gazelle [37]	○	●	○	○	●	○	●	○	○	○	●	●
MiniONN [43]	○	●	○	●	●	○	●	●	○	○	●	●
CryptoNets [24]	○	●	○	●	●	○	●	●	○	○	●	●
Slalom [72]	○	●	○	○	○	●	●	●	○	○	●	●
Origami [54]	○	●	○	○	○	○	●	○	○	○	●	●
Shredder [45]	○	●	○	○	○	○	●	○	○	○	●	●
Occlumency [40]	○	●	○	○	○	○	●	●	●	○	○	●
Delphi [47]	○	●	○	●	●	○	●	●	○	○	●	●
DarKnight	●	●	○	●	○	●	●	●	○	●	●	●

Figure 1: Overview of DarKnight components, its communication and, computation

using the cloud service. The cloud server is equipped with TEE which is responsible for protecting data privacy through encoding and performing non-linear operations. The TEE hardware guarantees code verification and authentication, where a data holder can verify the code and data usage within the enclave code. The system model uses K' GPU accelerators ($\text{GPU}_1, \dots, \text{GPU}_{K'}, K' > 1$) that participate in linear computations on data that is encoded in the TEE. In our implementation, we use Intel SGX as a TEE. Communication channels between the client, server, and GPUs are encrypted. For example, all the client data is first encrypted before being sent to the TEE. A pairwise secure channel between TEE and each GPU can be established using a secret key exchange protocol at the beginning of the session. While GPUs perform most of the training computations, input privacy on GPUs is guaranteed by our proposed encoding scheme (described next), which obfuscates the original input.

Threat Model: While adversaries can perform various attacks to exfiltrate DNN model parameters [61], DarKnight focuses on attacks that expose the datasets used in training or inference and attacks that modify computational results on untrusted hardware. We assume that data stored in TEE is protected from an adversary. Side-channel attacks are out of the scope of this work. The threat model on the cloud is a dynamic malicious adversary. This means at any given time, accelerator GPUs may try to glean private information from the data shared with them by the TEE. To provide perfect privacy it is the responsibility of the TEE to provide the encoded data that is drawn from a uniformly random distribution which is decoupled from the actual raw input data (More details in Section 5). Moreover, a subset of GPUs may try to extract information by collaborating with each other. We refer to them as *colluding GPUs*. Since the GPUs can be malicious, they may also inject faults

in the computation to sabotage training or inference. Thus, computational integrity will be explicitly verified by DarKnight in the presence of such an adversarial capability assumption.

Note that we assume only a subset of all the available GPUs collude. Such an assumption is not unique to this work as most MPC approaches assume only a subset of parties collude. Furthermore, this assumption is also true in scenarios where a client may request a subset of GPUs across different geographic regions, as is possible with current cloud services [60]. Clients may also request for nodes across multiple cloud providers to prevent collusion across the entire network. Nonetheless, it is useful to note that DarKnight does not provide provable privacy guarantees against a scenario when all the GPUs participating in distributed training collude.

In a system with K' accelerator GPUs, DarKnight simultaneously provides:

Data Privacy: DarKnight provides perfect privacy by making encoded data and raw inputs completely independent. Perfect privacy achieves when the data seen by any GPU does not give any information about the original input client data. More formally, perfect privacy is when the mutual information between encoded data and raw data is equal to zero. Namely, $I(\tilde{X}_k : X_k) = 0$, where I is the mutual information [13, 79].

Integrity: DarKnight is $(K'-1)$ -secure, namely it can *detect* any malicious computation even if $K'-1$ GPUs send erroneous results to TEE.

Collusion tolerance: DarKnight provide perfect privacy and integrity verification when M GPUs collude, where M is a function of K' (total GPUs in the system) and the number of inputs that can be encoded, as described at the end of this section 4.

Model Privacy: DarKnight does not reveal anything about the model to the client. However, model privacy on the server-side is out of the scope of this work. One common defense is using central differential privacy to keep the model private. Central differential privacy can be used on top of DarKnight [18].

3.1 DarKnight Flow

The initial machine learning model (W) that a user wants to train is loaded into the cloud server and is made accessible to the untrusted GPUs as well. DarKnight then uses the following steps:

- (1) A batch of training/inference input data set is encrypted by the client and sent to the TEE enclave on the server.
- (2) TEE initiates encoding on decrypted inputs.

- (3) During the forward/backward pass of training, each layer requires linear and nonlinear operations. The linear operations are compute-intensive and will be offloaded to GPUs. DarKnight's encoding mechanism is used to *seal* the data before sending the data to GPU accelerators. To seal the data, DarKnight uses the notion of a *virtual batch*, where K inputs and a random noise are linearly combined to form $K + M$ coded inputs, as described in the next section; M is the collusion tolerance as described above. The size of the virtual batch is limited by the size of the TEE memory that is necessary to encode K images, typically 4-8 images at a time. This virtual batch size may be different from the traditional batch size used in machine learning.
- (4) The encoded data is offloaded to GPUs for linear operation. Each GPU receives only one encoded data inputs.
- (5) GPUs perform linear operations on different encoded data sets and return the results to TEE.
- (6) The TEE decodes the received computational outputs using DarKnight's decoding strategy and then performs any non-linear operations within the TEE.
- (7) This process is repeated both for forward pass and backward propagation of each layer.

4 DARKNIGHT: PRIVACY IN TRAINING

In this section, we provide an overview of the forward and backward propagation of deep neural network training and the novel encoding and decoding process that we designed for privacy and integrity verification.

We first start with the encoding/decoding in forward pass which is the first phase in training and then explain how backward propagation works. Please note that forward pass and inference are similar in terms of encoding and decoding functions and hence the forward propagation strategy can be applied directly to inference only systems.

In the rest of this section, we assume that we have a DNN with L layers which is being trained with a virtual batch of K inputs, the model parameters $\mathbf{W}_l$ at layer l are updated using the well known SGD process.

Also for simplicity, we first show how this mechanism works for a system in which GPUs are not colluding, and next we expand the encoding to support a system with M colluding GPUs in section 4.5.

4.1 Forward Pass

At a layer l of the forward pass, we need to compute $\mathbf{y}_l = \langle \mathbf{W}_l, \mathbf{x}_l \rangle$, where $\mathbf{W}_l$ and $\mathbf{x}_l$ represent the model parameters and inputs in layer l , and $\langle \cdot, \cdot \rangle$ corresponds to the bilinear operation at that layer (e.g. matrix product, convolution, etc.). After the linear operation finishes, an activation function ($g(\cdot)$) creates the next layer input $\mathbf{x}_{l+1} = g(\mathbf{y}_l)$. Within this context, DarKnight first receives a set of K inputs $\mathbf{x}_0^{(1)}, \dots, \mathbf{x}_0^{(K)}$ for a batch training from a client. Our goal is to perform linear calculations of $\mathbf{y}_0^{(1)} = \langle \mathbf{W}_0, \mathbf{x}_0^{(1)} \rangle, \dots, \mathbf{y}_0^{(K)} = \langle \mathbf{W}_0, \mathbf{x}_0^{(K)} \rangle$ on the GPUs without exposing the raw inputs to the GPUs. Note that the subscript 0 in all these variables refers to the first layer. At this point, we drop the subscript for a more clear notation. Also, we use the notation $\mathbf{x}$ for the inputs that need to be

Figure 2: An example of DarKnight Encoding Scheme for $K=4$

protected and $\bar{\mathbf{x}}$ for the encoded inputs. DarKnight must protect $\mathbf{x}_1^{(i)}$ for each layer of the DNN when the layer's operations are outsourced to GPUs.

Key Insight: The main idea behind DarKnight's privacy protection scheme is the fact that the most computationally intensive operator (such as convolutions) is *bilinear*. Thus, instead of asking a GPU to calculate $\langle \mathbf{W}, \mathbf{x}^{(i)} \rangle$, which exposes the inputs, DarKnight uses matrix masking to linearly combine the inputs and add a random noise to them. Due to the bilinear property, any linear operation on K masked inputs can be recovered if there are K different linear computations performed.

Matrix Masking: Introduced by [14, 15], matrix masking scheme can be used for a variety of reasons such as noise addition, sampling, etc. Matrix masking uses the general form $\mathbf{B}\mathbf{X}\mathbf{A} + \mathbf{C}$ for protecting Matrix $\mathbf{X}$. In the above formula $\mathbf{B}$, $\mathbf{A}$, and $\mathbf{C}$ are called record transformation masks, attribute transformation masks, and displacing masks, respectively. Any of these additive and multiplicative masks can be used for encoding data based on the data privacy goal. Prior approaches have used different combinations of masks to protect matrix data [38, 69, 79]. DarKnight's encoding scheme is a form of matrix masking that is optimized for DNN's linear operations.

DarKnight Encoding: The SGX based enclave first receives a set of inputs (such as a set of images) from a data holder. Then the DarKnight scheme creates $K + 1$ encoding within the SGX from K data inputs $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)})$, as follows,

$$\bar{\mathbf{x}}^{(i)} = \alpha_{1,i} \mathbf{x}^{(1)} + \dots + \alpha_{K,i} \mathbf{x}^{(K)} + \alpha_{(K+1),i} \mathbf{r} \quad (1)$$

Where $i = 1, \dots, (K + 1)$. The scalars $\alpha_{i,j}$, and the noise vector $\mathbf{r}$ are randomly generated; and the size of $\mathbf{r}$ matches that of $\mathbf{x}$. The scalars $\alpha_{i,j}$'s are represented by matrix $\mathbf{A} \in \mathbb{R}^{(K+1),(K+1)}$, which are dynamically generated for each virtual batch and securely stored inside SGX for decoding. As we prove in a later section, by revealing the values $\bar{\mathbf{x}}^{(i)}$ s to GPUs, we do not disclose any information about the inputs $\mathbf{x}^{(i)}$ s. An example of encoding is shown in Figure 2.

At the next step, each encoded data $\bar{\mathbf{x}}^{(i)}$ is sent to a GPU_i which performs the following computation:

$$\bar{\mathbf{y}}^{(i)} = \langle \mathbf{W}, \bar{\mathbf{x}}^{(i)} \rangle, \quad i = 1, \dots, (K + 1).$$

Please note that each GPU receives *at most one* encoded data.

DarKnight Decoding: The $K + 1$ outputs $\bar{\mathbf{y}}^{(i)}$ returned from the GPUs must be decoded within the SGX to extract the original results $\mathbf{y}^{(i)}$. These values can be extracted as follows:

$$\bar{\mathbf{Y}} = \langle \mathbf{W}, [\bar{\mathbf{x}}^{(1)}, \dots, \bar{\mathbf{x}}^{(K+1)}] \rangle = \underbrace{\langle \mathbf{W}, [\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}, \mathbf{r}] \rangle}_{\mathbf{Y}} \cdot \mathbf{A} \Rightarrow \mathbf{Y} = \bar{\mathbf{Y}} \cdot \mathbf{A}^{-1} \quad (2)$$

4.2 Backward Propagation

The goal of training is to find the weight updates for each layer. For computing weight updates using SGD, we need to compute:

$$\mathbf{W}_l^{\text{new}} = \mathbf{W}_l^{\text{old}} - \eta \times \nabla \mathbf{W}_l, \quad \nabla \mathbf{W}_l = \frac{1}{K} \sum_{i=1}^K \langle \delta_l^{(i)}, \mathbf{x}_l^{(i)} \rangle \quad (3)$$

Here $\mathbf{x}_l^{(i)}$ is the i^{th} input of layer l . η is the learning rate, and $\delta_l^{(i)}$ is the gradient of the loss for the i^{th} point in the training batch, with respect to the output of layer l .

There are two sets of computation intensive linear operations in each layer for computing backward propagation.

(1) The first linear operation is to compute the derivative of loss with respect to weights which is computed as follows:

$$\nabla \mathbf{W}_l = \frac{\partial \mathcal{L}}{\partial \mathbf{W}_l} = \frac{1}{K} \sum_{i=1}^K \langle \delta_l^{(i)}, \mathbf{x}_l^{(i)} \rangle$$

In this equation $\mathbf{x}_l^{(i)}$ needs to be protected for data privacy. Hence, when we want to offload these computations to GPUs, encoded data should be used instead of raw data.

(2) In the above equation $\delta_l^{(i)}$ represents the gradient of layer l for input data point i , which is computed as a second linear operation that computes the derivative of loss with respect to the output of the layer. $\delta_l^{(i)}$ can be computed as follows, where g_l is the activation function of layer l .

$$\delta_l^{(i)} = \frac{\partial \ell(\mathbf{y}_l^{(i)}, \mathbf{y}_{exp}^{(i)})}{\partial \mathbf{y}_l} = \langle \delta_{l+1}^{(i)}, g_l'(\mathbf{y}_l^{(i)}) \rangle$$

As shown this computation does not contain any sensitive $\mathbf{x}^{(i)}$ information. Therefore, it can be offloaded to the GPUs without any data encoding.

In the rest of this section we explain how data encoding works for computing $\nabla \mathbf{W}$.

Key Insight: As explained in the previous section, the encoding process for forward pass exploited the invariant property of model parameter for any given input such that $\langle \mathbf{W}, [\bar{\mathbf{x}}^{(1)}, \dots, \bar{\mathbf{x}}^{(K+1)}] \rangle = \langle \mathbf{W}, [\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}, \mathbf{r}] \rangle \cdot \mathbf{A}$, meaning that a single $\mathbf{W}$ was shared between all the inputs of that layers. However, during the backward propagation process, we have different $\delta_l^{(i)}$ for each input $\mathbf{x}_l^{(i)}$. Thus, DarKnight designed a novel decoding scheme to extract the $\langle \delta_l^{(i)}, \mathbf{x}_l^{(i)} \rangle$ from obfuscated inputs $\langle \delta_l^{(i)}, \bar{\mathbf{x}}_l^{(i)} \rangle$.

The decoding scheme is based on a key insight. While backward propagation operates on a batch of inputs, it is not necessary to compute the $\langle \delta_l^{(i)}, \mathbf{x}_l^{(i)} \rangle$ for each input $\mathbf{x}^{(i)}$. Instead, the training process only needs to compute *cumulative parameter updates* for the entire batch of inputs. Hence, what is necessary to compute is the entire $\nabla \mathbf{W}_l$ which is an average over all updates corresponding to inputs in the batch.

DarKnight Encoding: DarKnight exploits this insight to protect privacy without significantly increasing the encoding and decoding complexity of the blinding process. As shown in Equation (3), there are K inputs on which gradients are computed. DarKnight calculates the overall weight update in the backward propagation

by summing up the following $K + 1$ equations each of which are computed on a different GPU,

$$\nabla \mathbf{W} = \frac{1}{K} \sum_{j=1}^{K+1} \gamma_j \text{Eq}_j, \quad \text{Eq}_j = \left\langle \sum_{i=1}^K \beta_{j,i} \delta^{(i)}, \bar{\mathbf{x}}^{(j)} \right\rangle \quad (4)$$

In the above equations, the encoded input $\bar{\mathbf{x}}^{(j)}$ to a layer is the same that was already calculated during the forward pass using Equation (1). As a result, the TEE can simply reuse the forward pass encoding without having to re-compute. As shown $\delta^{(i)}$ s are multiplied with the $\beta_{j,i}$ in the GPUs after which the GPUs compute the bi-linear operation to compute Eq_j .

In the above computation DarKnight selects $\alpha_{i,j}$'s, $\beta_{j,i}$'s, and γ_i 's such that they follow a very specific mathematical property as shown in Equation 5. If the values are selected to follow this property, the overall parameter updates $\nabla \mathbf{W}$ can be decoded very efficiently by scaling each GPU_i 's result with a corresponding γ_i 's and summing up all the $K + 1$ outputs.

$$\mathbf{B}^T \cdot \Gamma \cdot \mathbf{A}^T = \begin{bmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & 1 & 0 \end{bmatrix}_{K \times (K+1)} \quad (5)$$

Assuming batch size is equal to K , the $\beta_{j,i}$ parameters used for scaling δ values is gathered in the $K + 1$ by K matrix, $\mathbf{B}$. $\alpha_{i,j}$'s are gathered in the $K + 1$ by $K + 1$ matrix $\mathbf{A}$, and γ_i 's form the diagonal of a $K + 1$ by $K + 1$ matrix Γ , that gives us the proper parameters for efficient decoding. Note that the TEE keeps matrix Γ and $\mathbf{A}$ as secret within the enclave memory, while providing $\mathbf{B}$ to the GPUs so each GPU can compute Eq_j .

DarKnight Decoding: Given the constraint imposed on $\alpha_{j,i}$'s, $\beta_{j,i}$'s and γ_i 's the decoding process is trivially simple to extract $\nabla \mathbf{W}$. It is easy to see that if the scalars $\alpha_{i,j}$'s, $\beta_{i,j}$'s and γ_i 's satisfy the relation (5), we will have

$$\frac{1}{K} \sum_{j=1}^{K+1} \gamma_j \text{Eq}_j = \frac{1}{K} \sum_{i=1}^K \langle \delta_l^{(i)}, \mathbf{x}_l^{(i)} \rangle = \nabla \mathbf{W}_l \quad (6)$$

In other words, the decoding process only involves calculating a linear combination of the values in Equation (4). In particular, the TEE enclave receives Eq_j from each GPU_i and multiplies that value with the secretly held γ_i . It then computes the average over all the Eq_j values received to compute the overall weight update for the batch.

In Equation (4), the computation of $\beta_{j,i} \delta^{(i)}$ has no privacy implication. We use this encoding only to simplify the secure aggregation mechanism to compute the weight update $\nabla \mathbf{W}$. Thus we do not need to protect matrix $\mathbf{B}$ in the enclave.

Note that even though the above equations are computed over K inputs in a virtual batch, it is possible for the SGX enclave to securely store multiple $\nabla \mathbf{W}_l$ associated with multiple virtual batches that comprise the training batch. Once all the inputs in the training batch are processed the SGX enclave can do a single aggregation to generate a batch-wide weight update. Details are provided in Section 6.

DarKnight Training Complexity: The size of the α , δ , and γ matrices is just proportional to the square of the virtual batch size

that is being processed at one time. Therefore, generating them for every batch has a negligible performance overhead. Even with 8-64 batch size, (commonly used in VGG training [9, 29, 53]) these scaling values are substantially smaller than the model parameters $\mathbf{W}$. Hence, the order complexity of encoding/decoding operations is much less than the linear operations $\langle \mathbf{W}, \mathbf{x} \rangle$ in a DNN with millions of parameters. In addition to that, the process of decoding K inputs with one random noise requires $K + 1$ computations. During decoding, we extract $\mathbf{W} \cdot \mathbf{r}$, but that value is just dropped. Thus, DarKnight trades only $\frac{1}{K}$ additional computations to provide provably perfect privacy.

4.3 Decoding Correctness Proof

In this section we prove that the above described decoding approach correctly produces the aggregate weight update within the TEE when using Equation 6. Recall that DarKnight generates $\tilde{\mathbf{X}} = \mathbf{X}\mathbf{A}$ within the TEE. The TEE also generates the random matrix $\mathbf{B}$, $\mathbf{\Gamma}$ to satisfy Equation 5. The $\tilde{\mathbf{X}}$ along with the $\mathbf{B}$ are sent to the GPUs. The GPUs generate $\tilde{\delta} = \delta\mathbf{B}^T$ and then compute $\tilde{\mathbf{Y}} = \tilde{\delta}^T \tilde{\mathbf{X}}$.

$$\begin{aligned} \mathbf{X} &= [\mathbf{x}^1, \dots, \mathbf{x}^K, \mathbf{r}] \in \mathbb{R}^{N \times K+1}, \\ \tilde{\mathbf{X}} &= [\tilde{\mathbf{x}}^1, \dots, \tilde{\mathbf{x}}^{K+1}] \in \mathbb{R}^{N \times K+1}, \\ \delta &= [\delta^1, \dots, \delta^K] \in \mathbb{R}^{N \times K}, \\ \tilde{\delta} &= [\tilde{\delta}^1, \dots, \tilde{\delta}^K] \in \mathbb{R}^{N \times (K+1)}, \\ \mathbf{\Gamma} &= \text{Diag}(\gamma^1, \dots, \gamma^{K+1}) \in \mathbb{R}^{(K+1) \times (K+1)}, \\ \text{and } \mathbf{A} &\in \mathbb{R}^{K+1 \times K+1}, \mathbf{B} \in \mathbb{R}^{K+1 \times K}. \end{aligned} \quad (7)$$

The goal is to compute $\sum_{i=1}^K \delta^{(i)} \mathbf{x}^{(i)}$ by using Equation (6). We make use of the following algebraic equations,

$$\begin{aligned} \text{Tr}[\mathbf{X}] &= \sum_{i=1}^K X_{ii} \quad \text{definition of Matrix Trace} \\ \text{Tr}[\mathbf{X}\mathbf{Y}] &= \text{Tr}[\mathbf{Y}\mathbf{X}] \\ (\mathbf{A}\mathbf{B})^T &= \mathbf{B}^T \mathbf{A}^T \\ (\mathbf{A}^T)^T &= \mathbf{A} \end{aligned} \quad (8)$$

Now, as long as the equation 5 holds We will have

$$\begin{aligned} \sum_{j=1}^{K+1} \gamma_j \text{Eq}_j &= \text{Tr}(\mathbf{\Gamma} \tilde{\delta}^T \tilde{\mathbf{X}}) = \text{Tr}(\mathbf{\Gamma} \mathbf{B} \delta^T \mathbf{X} \mathbf{A}) = \\ \text{Tr}(\mathbf{A} \mathbf{\Gamma} \mathbf{B} \delta^T \mathbf{X}) &= \text{Tr}((\mathbf{B}^T \mathbf{\Gamma} \mathbf{A}^T)^T \cdot (\delta^T \mathbf{X})) = \sum_{i=1}^K \langle \delta_l^{(i)}, \mathbf{x}_l^{(i)} \rangle, \end{aligned} \quad (9)$$

which concludes the proof.

4.4 Computational Integrity

Apart from protecting data privacy, DarKnight's encoding method can be extended to detect computational integrity violations by untrusted GPUs. In this case, the linear computations performed by GPUs must also be verified. Recall that DarKnight creates $K + 1$ encoded inputs $\tilde{\mathbf{x}}^{(1)}, \dots, \tilde{\mathbf{x}}^{(K+1)}$ for K original inputs. To provide integrity, DarKnight creates one additional linear combination of

inputs (say $\tilde{\mathbf{x}}^{(K+2)}$), using the same approach as in Equation (1). This leads to having $K + 2$ linear equations for recovering $K + 1$ unknowns. This additional equation allows us to verify the accuracy of each result $\mathbf{y}^{(i)}$ by computing it redundantly at least twice using at least two sets of equations. If the results computed are not consistent across the two, one of the GPU cores may not function properly or their data is modified by an attacker. Once an integrity violation is detected, TEE may perform additional corrective action, such as executing on another GPU worker or perform additional redundant computations. But these actions are outside the scope of our current work.

4.5 Colluding GPUs

In this section, we investigate the scenario in which multiple GPUs can collaborate to extract information from the encoded data. With K' GPUs and virtual batch size of K , we can tolerate $M < K' - K$ colluding GPUs without compromising privacy. We show how we can securely outsource calculating $\langle \mathbf{W}, \mathbf{x}^{(i)} \rangle$, $i = 1, \dots, K$, to the GPUs. We first create $S = M + K$ encoded data vectors, $\tilde{\mathbf{x}}^i$, $i = 1, \dots, S$, using M noise vectors $\mathbf{r}^1, \dots, \mathbf{r}^M$, as follows.

$$\begin{aligned} \tilde{\mathbf{X}} &= \mathbf{X}\mathbf{A}_1 + \mathbf{R}\mathbf{A}_2, \quad \text{where}, \\ \tilde{\mathbf{X}} &= [\tilde{\mathbf{x}}^1, \dots, \tilde{\mathbf{x}}^S] \in \mathbb{R}^{N \times S}, \\ \mathbf{X} &= [\mathbf{x}^1, \dots, \mathbf{x}^K] \in \mathbb{R}^{N \times K}, \\ \mathbf{R} &= [\mathbf{r}^1, \dots, \mathbf{r}^M] \in \mathbb{R}^{N \times M}, \\ \text{and } \mathbf{A}_1 &\in \mathbb{R}^{K \times S}, \mathbf{A}_2 \in \mathbb{R}^{M \times S}. \end{aligned} \quad (10)$$

Here, the matrices $\mathbf{A}_1$ and $\mathbf{A}_2$ are the encoding coefficients similar to the initial scheme we used for DarKnight. Same as before, we can calculate the weight updates using the following equations:

$$\nabla \mathbf{W} = \sum_{j=1}^S \gamma_j \text{Eq}_j, \quad \text{Eq}_j = \left\langle \sum_{i=1}^K \beta_{j,i} \delta^{(i)}, \tilde{\mathbf{x}}^{(j)} \right\rangle \quad (11)$$

We now define

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}, \quad \mathbf{B} = [\beta_{j,i}], \quad \mathbf{\Gamma} = \text{Diag}(\gamma_1, \dots, \gamma_S) \quad (12)$$

Now, it is easy to show that we can recover $\nabla \mathbf{W}$ using Equation (11) if:

$$\mathbf{B}^T \cdot \mathbf{\Gamma} \cdot \mathbf{A} = \begin{bmatrix} 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 1 & 0 & \dots & 0 \end{bmatrix}_{K \times S} \quad (13)$$

As explained with this data encoding scheme, for a virtual batch of size K , and in a system with M colluding GPUs, we need $K + M$ GPUs for data privacy reasons. Also, we need one extra GPU for integrity checks which performs redundant computation to verify the results.

In summary: In a system with K' GPUs and virtual batch size K , DarKnight can provide data privacy and computational integrity by tolerating up to M colluding malicious GPUs, where $K + M + 1 \leq K'$.

5 PRIVACY GUARANTEE

In this section, we explain how we provide data privacy to DarKnight. As we explained in the previous section, for each virtual batch, each GPU receives only one encoded data for each layer. To provide data privacy, encoded data should be completely independent of the raw data. If we make every encoded data, $\bar{x}_i \quad \forall i \in [1, \dots, S]$, appear uniformly random in the Finite Field $\mathbb{F}_p$ (p is a large prime) that is representing it, the adversary cannot infer any information about raw data. This is due to lemma 1, since the values $\bar{x}_i$ s are simply a linear combination of the variables x_i 's plus uniform random variables over the operating field. More concretely, $I(\bar{X}_M : X_K) = 0$ and knowing $\bar{x}_i$ leaks no information about x_i . If every encoded data each GPU receives appears uniformly random, it will be impossible for a GPU to distinguish between different raw input data by just observing the encoded data (it's the formal definition of computational indistinguishability in [25]). In the other words to achieve this goal, we make sure every encoded data exposed to GPUs looks uniformly random over Finite Field $\mathbb{F}_p$ for a large prime p . Since each GPU observes only one encoded data from a virtual batch, this privacy guarantee is one of the strongest privacy guarantees in the field of cryptography which is known as the one-time pad and used in many of the prior works [16, 72, 79]. For this, first, DarKnight quantizes all the inputs and the model weights and maps them to the field $\mathbb{F}_p$. In this work, we choose $l = 8$ and $p = 2^{25} - 39$ for ResNet, VGG and MobileNet models which is the largest prime with 25 bits.

LEMMA 1. Assume that $x \in \mathbb{F}_p$ is a scalar in the field $\mathbb{F}_p$, and z is a random variable, uniformly chosen over the field $\mathbb{F}_p$, and let $y = x + z$. Then

$$I(y; x) = 0. \quad (14)$$

In the other words, knowing y leaks no information about x .

PROOF. Since the random variable z is uniform over the field, the random variables x and y will be independent. Therefore, their mutual information will be zero. $\square$

we can apply Lemma 1 to our work, by replacing the scalar x with the linear combination of x_i 's, and replacing the random variable z with the linear combination of uniform random variables R_i 's (note that linear combination of uniform random variables in the field would be a uniform random variable).

Quantization: There are various quantization schemes proposed in literature [28, 68], which strive to preserve accuracy while training with quantized data. Algorithm 1 presents our quantization code. At the first step of the quantization, we convert all the values from the floating-point domain to the fixed-point domain. To do so, after selecting the number of fractional bits, l , we multiply all the values by 2^l and then call the Round function to convert the values to integer representation as shown in lines 12 to 18. At the next step, we use function Field to transfer the computations to Finite Field. To do so, for negative numbers, we first add p to them to transfer every value to $\mathbb{F}_p$. After that, we compute the remainder of the computation to p . Please note that both X^l 's and W^l 's are multiplied by 2^l before linear operations. Hence, bias should be multiplied by 2^{2l} so that after the linear operation we will be able to convert everything back to floating-point representation correctly.

After the linear computation on GPUs, TEE receives $\bar{Y}$. TEE then subtracts p from all the elements larger than $\frac{p}{2}$ to restore negative numbers. At the next step, the decoding function is called to cancel the noise from linear operation results (Y_q). Finally, $\text{Round}(Y_q \times 2^{-l})$ is called and at the end by multiplying 2^{-l} to the output, we have the result of the desired linear operation. Now, the original floating-point format is recovered for computing Non-Linear operations in TEE.

Please note that quantizing the weight and bias values does not have a privacy implication. These values may be stored and quantized in the GPUs so that the linear operation operates correctly in the Finite Field. For ResNet and MobileNet DNNs, this static quantization shows a good accuracy because these DNNs have normalization layers that keep the values within a range. However, for VGG models a slightly different quantization is used to dynamically normalize the values of inputs and weights if they pass the limits as suggested in [55, 78]. We normalize the values by dividing them to the maximum absolute entry of the vector.

Random Numbers: Random scalars (A, B, Γ) and random vectors (R) are generated in the Finite Field $\mathbb{F}_p$ and encoding/decoding computations are defined over Field $\mathbb{F}_p$.

Algorithm 1 The pseudo-code for quantization

```

1: procedure LINEAR OPERATIONS( $W, X, b$ )
2:    $W_q = \text{Field}(\text{Round}(W \times 2^l))$ 
3:    $b_q = \text{Field}(\text{Round}(b \times 2^{2l}))$ 
4:    $X_q = \text{Field}(\text{Round}(X \times 2^l))$ 
5:   Generate  $R$  and all the random scalars  $A, B, \Gamma$  in the field  $F_p$ .
6:    $\bar{X} = \text{Encode}(X, r, A)$  in the Field using equation (10).
7:   GPUs compute  $\bar{Y} = \langle W_q, \bar{X}, b \rangle$ 
8:    $Y_q = \text{Field}(\bar{Y}A^{-1})$ 
9:    $Y = \text{Round}(Y_q \times 2^{-l}) \times 2^{-l}$ 
10:  return  $Y$ 
11: procedure ROUND( $X$ )
12:  for  $\forall X_i \in X$  do
13:    if  $(X_i - \lfloor X_i \rfloor < 0.5)$  then
14:       $X_i^r \leftarrow \lfloor X_i \rfloor$ 
15:    else
16:       $X_i^r \leftarrow \lfloor X_i \rfloor + 1$ 
17:  return  $X^r$ 
18: procedure FIELD( $X$ )
19:  for  $\forall X_i \in X$  do
20:    if  $(X_i^f < 0)$  then
21:       $X_i^f \leftarrow X_i^f + p$ 
22:     $X_i^f \leftarrow X_i \bmod p$ 
23:  return  $X^r$ 

```

Colluding GPUs: As we explained in Section 4.5, we are using K' GPUs among which at most M collude. Therefore, we need at least M noise vectors according to DarKnight's scheme as shown in Equation 10. Now, the matrix RA_2 is N by S dimensional. Assume that the matrix A_2 is full rank and the noises in the matrix R are independent and uniform over the field F_p . Now for any subset $I \subseteq [K']$ with $|I| \leq M$ of GPUs that collide, they will share the

equations $XA_{1,I} + RA_{2,I}$ (where A_I is a sub-matrix of A whose columns are chosen from the set I). Since A_2 is full-rank, any subset of its columns are also full rank. Therefore, there is no linear combination that can vanish the noise in the M shared equations, and any linear combination of uniform random variables over a field is also uniform. Hence, no matter how the M equations are combined linearly, the result of the combination will seem uniform over the field F_p to GPUs' perspective.

6 DARKNIGHT IMPLEMENTATION

In this section we provide the details on how to efficiently implement DarKnight on off-the-shelf hardware consisting of Intel SGX CPU attached with Nvidia GPUs.

Encoded Data Storage During Forward Pass: The encoded intermediate feature map ($\bar{x}^{(i)}$) that is received by GPUs during the forward propagation is also used during backward propagation for computing the weight update, as shown in Equation 4. Rather than requesting TEE to resend the encoded inputs during back propagation, our current implementation of DarKnight stores these encoded inputs within the GPU memory.

Large Batch Aggregation: One of the benefits of the decoding process employed by DarKnight's back propagation is that an aggregate weight update across all the inputs in a virtual batch is computed by the TEE. Hence, individual weights associated with each input are not revealed (and not even computed by the TEE). The aggregate weight update over a virtual batch is essentially a customized version of secure aggregation for weight updates [7]. As explained in section 4, model parameters ($\mathbf{W}$) are stored outside the enclave and are visible to GPUs. Once the aggregate weight update $\nabla \mathbf{W}$ is computed in the TEE, that data is sent to the GPUs so they can update the model. Some prior works [21, 82] have shown that $\nabla \mathbf{W}$ may leak some information about the intermediate features which may eventually leads to input leakage as a side channel information.

While side channels are outside the scope of this work, we believe that in practice the side channel information leakage can be curtailed. In particular, this prior work [82] also observed that one solution to reduce leakage drastically is to increase the batch size over which the aggregate weight updates are computed. Thus, using a large batch size, one can eliminate nearly all the side channel leakage. DarKnight can be easily adapted to expose the $\nabla \mathbf{W}$ of a large batch to the GPU memory. Earlier, we introduced the notion of a *Virtual Batch*, which is essentially the largest number of images that we can be processed at the same time and fits inside SGX. The size of the virtual batch is purely a limitation of SGX memory.

In our implementation we enable large batch weight aggregation as follows. The TEE still computes $\nabla \mathbf{W}$ at the granularity of a virtual batch. However, instead of updating the weights immediately, our implementation aggregates the $\nabla \mathbf{W}^v$ of multiple virtual batches $v > 1$ within the TEE without disclosing the weight updates to the GPUs. The TEE internally accumulates $\nabla \mathbf{W}^v$ associated with all the virtual batches from a large batch.

One challenge we faced with this implementation is that to store multiple $\nabla \mathbf{W}^v$ for all the virtual batches inside SGX exceeds the memory limitation. To resolve this problem, after each virtual batch computation, we encrypt the pages storing $\nabla \mathbf{W}^v$ and then write

it back to the untrusted memory. Note that this encryption and eviction is outside the critical path. Once all the virtual batches are processed the TEE then incrementally reloads each of the $\nabla \mathbf{W}^v$, decrypts them and creates an aggregate update.

We further optimized the reloading and aggregation of $\nabla \mathbf{W}^v$ using a sharding technique. Note that the aggregation is highly parallel task. We break down the $\nabla \mathbf{W}^v$ into multiple shards before storing them to the untrusted memory. For example, each shard may be a set of DNN layers. Then during the reloading process we perform shard-wise aggregation and then send those updates to the GPUs to incrementally update the model parameters. This pipelined approach to shard-wise aggregation essentially eliminates all the performance penalties associated with large batch aggregation.

Algorithm 2 The pseudo-code for parameters update

```

1: procedure BACKWARD( $\mathbf{W}$ )
2:    $V \leftarrow \text{VirtualBatch.size}()$ 
3:    $N \leftarrow \text{LargeBatch.size}()$ 
4:    $M \leftarrow \frac{N}{V}$ 
5:   Initialize memory Pointer
6:   for  $v = 1, 2, \dots, M$  do                                 $\triangleright$  for each virtual batch
7:     for  $l = 1, 2, \dots, L$  do                                 $\triangleright$  for each layer
8:       Compute  $\nabla \mathbf{W}_l^v$ 
9:        $\nabla \mathbf{W}'^v = \text{Encrypt}(\nabla \mathbf{W}^v)$ 
10:      Pointer.append (Evict( $\nabla \mathbf{W}'^v$ ))
11:    $\nabla \mathbf{W} \leftarrow \text{UpdateAggregation}(\text{Pointer})$ 
12:    $\mathbf{W}^{\text{new}} = \mathbf{W}^{\text{new}} - \eta \times \nabla \mathbf{W}$ 
13:   return  $\mathbf{W}^{\text{new}}$ 
14: procedure UPDATEAGGREGATION(Pointer)
15:    $V \leftarrow \text{VirtualBatch.size}()$ 
16:    $N \leftarrow \text{LargeBatch.size}()$ 
17:    $M \leftarrow \frac{N}{V}$ 
18:   for  $v = 1, 2, \dots, M$  do
19:      $\nabla \mathbf{W}^v = \text{Decrypt}(\text{Pointer} + v \times \mathbf{W}.size())$ 
20:      $\nabla \mathbf{W} = \nabla \mathbf{W} + \nabla \mathbf{W}^v$ 
21:   return  $\nabla \mathbf{W}$ 

```

Pseudo Code of Aggregation: Algorithm 2 shows the steps of backward propagation. In line 4, the actual batch size is divided to the virtual batch size to get the number of virtual batches we have to process before updating the weight. Line 5 initializes the pointer to the untrusted memory where we store $\nabla \mathbf{W}$ s before the aggregation. Line 6 repeats the for loop for each virtual batch. Line 7 to 8 shows how $\nabla \mathbf{W}$ is computed for each layer of DNN. Line 9 Encrypt $\nabla \mathbf{W}^v$ that is containing all the weight updates of the network for that virtual batch. Line 10 calls evict function containing $\nabla \mathbf{W}'^v$. As a result, this page is moved to the untrusted memory with all the precautions. Line 11 calls Aggregation function. This function has a reference to all the pages containing part of the $\nabla \mathbf{W}$ and construct the whole $\nabla \mathbf{W}$. Finally in Line 12 the $\nabla \mathbf{W}$ is sent to GPUs for the weight updates.

In Figure 3, we show how different sizes of virtual batch size can speedup the aggregation time. Having a larger virtual batch size helps with the less number of encryption/eviction and decryption. However, increasing a size of virtual batch at a certain point, will

Figure 3: Effect of different virtual batch sizes (number of GPUs(K)-1 when $M = 0$) on the aggregation (decoding) speedup relative to virtual batch size $K = 1$

increase the latency because the whole data cannot fit inside the enclave memory. As depicted in this figure, if we want to update the weight for batches of size 128, for all our three networks, virtual batch size of 4 ($K = 4$) shows the best performance. Note that this value of K will likely increase in future SGX implementation as the protected memory segment size grows. But in our implementation on current Intel SGX systems $K=4$ shows the best performance.

7 EXPERIMENTAL SETUP AND RESULTS

We implemented DarKnight on a server consisting of an Intel Coffee Lake E-2174G 3.80GHz processor with SGX support, and Nvidia GeForce GTX 1080 Ti GPUs. The server has 64 GB RAM and supports Intel Soft Guard Extensions (SGX). DarKnight’s training scheme and the related unique coding requirements are implemented as an SGX enclave thread where both the decoding and encoding are performed. For SGX implementations, we used Intel Deep Neural Network Library (DNNL) for designing the DNN layers including the Conv layer, ReLU, MaxPooling, and Eigen library for Dense layer. We used Keras 2.1.5, Tensflow 1.8.0, and Python 3.6.8.

To evaluate the method, we used three different DNN models with different structures: VGG16 [67] with 138 million parameters, ResNet50 [31] with 23 million parameters and, MobileNetV2 [63] with 3.4 million parameters. We chose MobileNetV2 because it is the worst-case benchmark for our model as it reduces linear operations considerably (using depth-wise separable convolution), thereby reducing the need for GPU acceleration. We used CIFAR-10 [39] that has 50,000 training images evenly distributed between 10 categories, and ImageNet [62] with about 1.2 million images and 1000 categories as our datasets.

7.1 Training Results

For evaluating training performance, two critical aspects are examined: accuracy impact and speedup of the training.

Effect of Random Noise and Quantization on Accuracy: Quantization, using fixed-point arithmetic and adding the uniform random noise to encode the data for privacy reasons may cause accuracy degradation on GPUs. To study the impact, Figure 4 shows

the training accuracy on clean data when there is no privacy scheme in the system and the accuracy of training in the presence of data encoding on VGG16, ResNet50, and MobileNetV2 on CIFAR-10. For training on CIFAR-10, as depicted in Figure 4 after 100 epochs, only a negligible accuracy degradation (less than 0.01 for all the applications) is observed. Very similar behavior is observed across a wide range of models. For MobileNetV2, even a slight increase is observed. This behaviour is observed in prior works as quantization may remove the noise on data [41, 42].

Training Execution Time, Non-Pipelined: Figure 5 shows the speedup of training using DarKnight with $K = 2$ images encoded and offloaded to 3 GPUs relative to the baseline fully implemented on SGX. The encoding operation in TEE and the linear operations in GPUs are serialized in this design, which we refer to as the non-pipelined version. The figure breaks down the execution time into linear and non-linear categories. Non-linear category includes all the operations performed in TEE including the encoding and decoding overheads, while linear costs include operations performed on GPUs plus the communication cost to move the encoded data to the GPUs and bring the computed results back to the TEE. For emulating the communication time we used Infiniband [58, 65] network switch between TEE and each GPU which has a 40 Gbps bandwidth. We assume a non-pipelined implementation where the data encoding and decoding process are performed sequentially.

The results show that DarKnight speeds up the total linear operation time of VGG16 by 23x by exploiting GPUs parallelism. Even the non-linear operations see a 1.89X speedup. This result is due to the observation that the baseline has to encrypt/decrypt some of the large intermediate feature maps that do not fit within SGX memory. Overall for VGG16 the execution time is improved by more than 8X with DarKnight. Both ResNet and MobileNet models have batch normalization layers which are non-linear operations that are computation-intensive and cannot be offload to GPU accelerators. Even in this worst-case scenario, performance gains of 4.2X and 2.2X are achieved.

Table 3 shows a more detailed breakdown of the fraction of time spent in various operations in each setting. In the baseline majority of time is spent in the linear operations. For VGG16 baseline spends 84% of the time on linear operations. Due to batch-normalization overheads and reduced linear operation counts ResNet50 and MobileNetV2 spend around 60% in linear execution time. Using the DarKnight the distribution of execution time is reversed. VGG16 spends nearly 50% of the time on non-linear operations as GPUs accelerate linear operations. It also pays 19% overhead for encoding and decoding. For ResNet50 and MobileNetV2 this overhead is lower because batch normalization dominates the execution time. DarKnight also pays a new communication overhead to move the data between TEE and GPUs. Across all three models, about 20% of the total training time in DarKnight is spent in this communication phase.

Training Execution Time, Pipelined: The non-pipelined implementation in the previous section assumes that the data encoding/decoding in TEE is serialized with the GPU executions. However, the communication overhead can be easily hidden by overlapping communication and computations. Meaning that, while data communication happens between GPUs and TEE, TEE can encode the next data batch of data. Pipelined implementation with

Figure 4: Training accuracy of DarKnight for CIFAR-10 with (a) VGG16 (b) ResNet50 (c) MobileNetV2

Figure 5: ImageNet Training Speedup for None-Pipeline and Pipeline Settings

Table 3: ImageNet Training Time Breakdown for Different Networks (values are percentage relative to the total execution time)

Operation	VGG16		ResNet50		MobileNetV2	
	DarKnight	Baseline	DarKnight	Baseline	DarKnight	Baseline
Linear	0.04	0.84	0.04	0.61	0.06	0.62
NonLinear	0.50	0.16	0.75	0.39	0.63	0.38
Encoding-Decoding	0.19	0	0.01	0	0.08	0
Communication	0.26	0	0.2	0	0.23	0

asynchronous SGD has been designed in prior work [52, 80]. A pipelined implementation encodes one virtual batch and launches it to the GPUs. While GPUs are performing linear operations, the next virtual batch is encoded under the shadow of GPUs execution time. Figure 5 shows the results. By Overlapping the communication and computation, DarKnight speeds up the total linear operation time by 20 – 158x with pipelined design which leads to an overall higher speedup.

Non-Private Model: Table 4 shows the training speedup of using three unprotected GPUs with Pytorch library over SGX-only computation or using DarKnight (3 GPUs). However, unprotected GPUs does not provide any privacy guarantee. Using only SGX to provide privacy leads to two orders of magnitude slowdown. DarKnight is designed to bridge this gap between TEE and GPU speed while protecting the privacy of data. Please note that, other privacy preserving methods, such as fully homomorphic encryption and secure multi-party computing, are orders of magnitude slower than GPU implementation [23, 37].

Parallelism in Baseline: Using multiple SGX threads one can train concurrently on different batches. Figure 7 shows the effect of multi-threading. The latency grows with more threading within

Table 4: Non-Private Training Speedup on 3-GPUs Relative to SGX-only and DarKnight (3 GPUs) for Training on ImageNet

Model	VGG16	ResNet50	MobileNetV2
Speedup over DarKnight	23.93	41.01	27.51
Speedup over SGX	273.26	216.62	80.31

SGX. The main reason for this counter-intuitive behavior is that training large models are very memory intensive. But current implementations of SGX have limited memory encryption unit capacity. When the training data exceeds the memory limit there is substantial performance overhead in moving data between protected memory and untrusted DRAM. The same behaviour is observed in [72]. So the the baseline with that multi-threading cannot provide a better performance for such memory intensive computations.

7.2 Inference Results

In this section we briefly explain Slalom [72] a prior work that used SGX and GPU for inference and why that strategy is unable to handle training. Then we evaluate the speedup of DarKnight in the inference phase.

Slalom Inference Slalom is an inference framework that uses TEEs to protect data privacy and integrity. Slalom uses the Intel SGX enclave to blind input data x from a client with an additive stream cipher noise r . The blinded data $(x + r)$ is then sent to an untrusted GPU where linear operations are performed. The computed data $W \cdot (x + r)$ is then returned to the enclave which can decode the correct computational output $W \cdot x$ by subtracting the precomputed $W \cdot r$. Here W is the model parameter matrix. But this encoding approach cannot be applied for training, since it precomputes $W \cdot r$. Precomputing the blinding factors is not feasible during training since the model parameters W are updated after processing every batch. Computing $W \cdot r$ inside the SGX after every batch also defeats the purpose of offloading the linear computations to GPU. This fundamental challenge make it impossible to use Slalom encoding for training. As such DarKnight developed a more comprehensive encoding strategy to handle training on private data while providing rigorous privacy guarantee. Nonetheless DarKnight can perform inference on private data. Hence, we compare the inference performance of DarKnight and Slalom. For a fair comparison, we implemented DarKnight using Eigen library which is the design used in Slalom. Eigen is a high-performance C++ based linear algebra library.

Figure 6: a) Inference speedup comparison with different implementations relative to SGX for VGG16, and MobileNetV1 on ImageNet. b) Inference speedup comparison of different operations relative to DarkKnight(1) for different virtual batch-sizes for VGG16 on ImageNet.

Figure 7: Effect of Multi-threading on the execution time of training VGG16 relative to the single thread execution time

Inference Speedup: Fig. 6 (a) compares the speedup of the inference for VGG16 and MobileNetV1 across five different configurations. The baseline bar (SGX) performs all the calculations within SGX. The red bar uses Slalom blinding while trusting GPU that results are always correct, DarKnight(4) is our model while using a virtual batch size of 4. Slalom+Integrity bar shows the performance when Slalom’s data integrity verification (using Freivalds algorithm [6]) is deployed to verify GPU computations. DarKnight(3)+Integrity uses DarKnight with a virtual batch size of 3 and an additional equation to redundantly compute all the results twice for integrity verification.

For VGG16, DarKnight(4) provides 15X speedup, compared to the SGX only baseline, and 30% improvement over Slalom. Slalom’s implementation encrypts $W \cdot r$ and stores them outside of SGX memory. At each layer, they retrieve the necessary unblinding factors into SGX, then decrypt them before using them. This approach is required since their approach is required to compute the unblinding factors. DarKnight does not need such a pre-computation. Instead it performs $1/K$ additional computations to decode the. Hence, DarKnight performs additional computations on a GPU but reduces the SGX memory pressure. When providing the additional integrity checks, DarKnight(3) provides about 13X speedup over baseline and 1.45X speedup over Slalom.

Effect of Virtual Batch Size: Recall that virtual batch size is the number of images that are linearly combined in the Equation (1). Fig. 6(b) quantifies the effect of batch size on the inference time. In the figure, DarKnight(K) is used to denote a virtual batch size of K . For the same number of input data points with different batch

sizes, we issue inference requests and divided the total inference time across four categories of operations: decoding, encoding, Relu, and Maxpooling operations. We used DarKnight(1) as the baseline. It represents the case where a single image is combined with a uniform random noise r to create two equations using Equation (1). As the virtual batch size increases the total speedup improved as long as the virtual batch size fits within SGX memory limits. As the virtual batch size exceeds 4, the execution time gets worse due to SGX memory overflow in our current experimental setup. But as the SGX memory limitation is relaxed in future we believe large virtual batches can be processed to further improve DarKnight’s performance.

8 CONCLUSION

This work proposes DarKnight, a privacy and integrity preserving framework for DNNs’ training and inference. DarKnight uses a hybrid execution model where TEE provide data obfuscation and untrusted GPUs provide computation acceleration. We design a data encoding in TEE for linear operations of DNN training and we provide a rigorous data privacy guarantee. DarKnight can also provide computation integrity and it is robust even in the presence of a malicious colluding GPUs. We evaluated different models and datasets and observe an average of 6.5X training speedup and 12.5X inference speedup without accuracy degradation over the baseline fully implemented inside TEE.

ACKNOWLEDGMENTS

We would like to express our special gratitude to Hsien-Hsin Sean Lee, Edward Suh, Wenjie Xiong, Chuan Guo and Mark Tygert for sharing their wisdom with us. We are immensely grateful to Krishna Giri Narra and Caroline Tripple for their valuable feedbacks on this project. We sincerely thank all the reviewers for their time and constructive comments. This material is based upon work supported by Defense Advanced Research Projects Agency (DARPA) under Contract Nos. HR001117C0053, HR001120C0088, NSF award number 5345039074, and Facebook AI Research Award Numbers 2215031173 and 2215031183. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 308–318.
- [2] Tiago Alves. 2004. Trustzone: Integrated hardware and software security. *White paper* (2004).
- [3] Amazon. 2020. *Machine Learning on AWS*. <https://aws.amazon.com/machine-learning>
- [4] Aref Asvadishehijani, Murat Kantarcioglu, and Bradley Malin. 2020. GOAT: GPU Outsourcing of Deep Learning Training With Asynchronous Probabilistic Integrity Verification Inside Trusted Execution Environment. *arXiv preprint arXiv:2010.08855* (2020).
- [5] Marcelo Blatt, Alexander Gusev, Yuriy Polyakov, and Shafi Goldwasser. 2020. Secure large-scale genome-wide association studies using homomorphic encryption. *Proceedings of the National Academy of Sciences* 117, 21 (2020), 11608–11613.
- [6] Lenore Blum and Manuel Blum. 1975. Toward a mathematical theory of inductive inference. *Information and control* 28, 2 (1975), 125–155.
- [7] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 1175–1191.
- [8] Ferdinand Brasser, Urs Müller, Alexandra Dmitrienko, Kari Kostiaainen, Srdjan Capkun, and Ahmad-Reza Sadeghi. 2017. Software grand exposure: {SGX} cache attacks are practical. In *11th {USENIX} Workshop on Offensive Technologies ({WOOT} 17)*.
- [9] Alfredo Canziani, Adam Paszke, and Eugenio Culurciello. 2016. An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678* (2016).
- [10] Nicholas Carlini, Samuel Deng, Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody, Shuang Song, Abhradeep Thakurta, and Florian Tramer. 2020. An Attack on InstaHide: Is Private Learning Possible with Instance Encoding? *arXiv preprint arXiv:2011.05315* (2020).
- [11] Victor Costan and Srinivas Devadas. 2016. Intel SGX Explained. *IACR Cryptology ePrint Archive* 2016, 086 (2016), 1–118.
- [12] Victor Costan, Ilija Lebedev, and Srinivas Devadas. 2016. Sanctum: Minimal hardware extensions for strong software isolation. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*. 857–874.
- [13] Thomas M Cover. 1999. *Elements of information theory*. John Wiley & Sons.
- [14] LH Cox. 1994. Matrix masking methods for disclosure limitation in microdata. *Surv. Methodol.* 20 (1994), 165–169.
- [15] Lawrence H Cox. 1980. Suppression methodology and statistical disclosure control. *J. Amer. Statist. Assoc.* 75, 370 (1980), 377–385.
- [16] Fu-Guo Deng and Gui Lu Long. 2004. Secure direct communication with a quantum one-time pad. *Physical Review A* 69, 5 (2004), 052319.
- [17] Yanjie Dong, Julian Cheng, Md Jahangir Hossain, and Victor CM Leung. 2019. Secure distributed on-device learning networks with byzantine adversaries. *IEEE Network* 33, 6 (2019), 180–187.
- [18] Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. 2019. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2468–2479.
- [19] Úlfar Erlingsson, Vasily Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. 1054–1067.
- [20] Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in neural information processing systems*. 2137–2145.
- [21] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. Inverting Gradients—How easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053* (2020).
- [22] Craig Gentry. 2009. Fully homomorphic encryption using ideal lattices. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 169–178.
- [23] Zahra Ghodsi, Tianyu Gu, and Siddharth Garg. 2017. Safetynets: Verifiable execution of deep neural networks on an untrusted cloud. *arXiv preprint arXiv:1706.10268* (2017).
- [24] Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. 2016. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning*. 201–210.
- [25] Oded Goldreich. 2007. *Foundations of cryptography: volume 1, basic tools*. Cambridge university press.
- [26] Google. 2020. *Google AI platform*. <https://cloud.google.com/products/ai>
- [27] Johannes Götzfried, Moritz Eckert, Sebastian Schinzel, and Tilo Müller. 2017. Cache attacks on Intel SGX. In *Proceedings of the 10th European Workshop on Systems Security*. 1–6.
- [28] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. 2015. Deep learning with limited numerical precision. In *International Conference on Machine Learning*. 1737–1746.
- [29] Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149* (2015).
- [30] Hanieh Hashemi, Yongqin Wang, Chuan Guo, and Murali Annaram. 2021. Byzantine-Robust and Privacy-Preserving Framework for FedML. *arXiv preprint arXiv:2105.02295* (2021).
- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [32] JB Heaton, NG Polson, and Jan Hendrik Witte. 2017. Deep learning for finance: deep portfolios. *Applied Stochastic Models in Business and Industry* 33, 1 (2017), 3–12.
- [33] Weizhe Hua, Muhammad Umar, Zhiru Zhang, and G Edward Suh. 2020. GuardNN: Secure DNN Accelerator for Privacy-Preserving Deep Learning. *arXiv preprint arXiv:2008.11632* (2020).
- [34] Yangsibo Huang, Zhao Song, Kai Li, and Sanjeev Arora. 2020. Instahide: Instance-hiding schemes for private distributed learning. In *International Conference on Machine Learning*. PMLR, 4507–4518.
- [35] Tyler Hunt, Congzheng Song, Reza Shokri, Vitaly Shmatikov, and Emmett Witchel. 2018. Chiron: Privacy-preserving machine learning as a service. *arXiv preprint arXiv:1803.05961* (2018).
- [36] Nick Hynes, Raymond Cheng, and Dawn Song. 2018. Efficient deep learning on multi-source private data. *arXiv preprint arXiv:1807.06689* (2018).
- [37] Chirag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. 2018. {GAZELLE}: A low latency framework for secure neural network inference. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 1651–1669.
- [38] Jay J Kim. 1986. A method for limiting disclosure in microdata based on random noise and transformation. In *Proceedings of the section on survey research methods*. American Statistical Association Alexandria, VA, 303–308.
- [39] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. *online: http://www.cs.toronto.edu/kriz/cifar.html* (2009).
- [40] Taegyeon Lee, Zhiqi Lin, Saumay Pushp, Caihua Li, Yunxin Liu, Youngki Lee, Fengyuan Xu, Chenren Xu, Lintao Zhang, and Junehua Song. 2019. Occlumency: Privacy-preserving remote deep-learning inference using SGX. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–17.
- [41] Darryl Lin, Sachin Talathi, and Sreekanth Annapureddy. 2016. Fixed point quantization of deep convolutional networks. In *International conference on machine learning*. PMLR, 2849–2858.
- [42] Ji Lin, Chuang Gan, and Song Han. 2019. Defensive quantization: When efficiency meets robustness. *arXiv preprint arXiv:1904.08444* (2019).
- [43] Jian Liu, Mika Juuti, Yao Lu, and Nadarajah Asokan. 2017. Oblivious neural network predictions via minionn transformations. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 619–631.
- [44] Microsoft. 2020. *Azure Machine Learning*. <https://azure.microsoft.com/en-us/services/machine-learning>
- [45] Fatemehsadat Miresghallah, Mohammadkazem Taram, Prakash Ramrakhiani, Ali Jalali, Dean Tullsen, and Hadi Esmaeilzadeh. 2020. Shredder: Learning noise distributions to protect inference privacy. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 3–18.
- [46] Fatemehsadat Mirshghallah, Mohammadkazem Taram, Praneeth Vepakomma, Abhishek Singh, Ramesh Raskar, and Hadi Esmaeilzadeh. 2020. Privacy in deep learning: A survey. *arXiv preprint arXiv:2004.12254* (2020).
- [47] Pratyush Mishra, Ryan Lehmkuhl, Akshayaram Srinivasan, Wenting Zheng, and Raluca Ada Popa. 2020. DELPHI: A cryptographic inference service for neural networks. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*.
- [48] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. 2021. PPFL: privacy-preserving federated learning with trusted execution environments. *arXiv preprint arXiv:2104.14380* (2021).
- [49] Fan Mo, Ali Shahin Shamsabadi, Kleomenis Katevas, Soteris Demetriou, Ilias Leontiadis, Andrea Cavallaro, and Hamed Haddadi. 2020. Darknetz: towards model privacy at the edge using trusted execution environments. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 161–174.
- [50] Payman Mohassel and Peter Rindal. 2018. ABY3: A mixed protocol framework for machine learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 35–52.
- [51] Payman Mohassel and Yupeng Zhang. 2017. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 19–38.
- [52] Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R Devanur, Gregory R Ganger, Phillip B Gibbons, and Matei Zaharia. 2019. PipeDream: generalized pipeline parallelism for DNN training. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. 1–15.

- [53] Krishna Giri Narra, Zhifeng Lin, Mehrdad Kiamari, Salman Avestimehr, and Murali Annavaram. 2019. Slack squeeze coded computing for adaptive straggler mitigation. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–16.
- [54] Krishna Giri Narra, Zhifeng Lin, Yongqin Wang, Keshav Balasubramaniam, and Murali Annavaram. 2019. Privacy-Preserving Inference in Machine Learning Services Using Trusted Execution Environments. *arXiv preprint arXiv:1912.03485* (2019).
- [55] Lucien KL Ng, Sherman SM Chow, Anna PY Woo, Donald PH Wong, and Yongjun Zhao. 2021. Goten: GPU-Outsourcing Trusted Execution of Neural Network Training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 14876–14883.
- [56] Oleksii Oleksenko, Bohdan Trach, Robert Krahn, Mark Silberstein, and Christof Fetzer. 2018. Varys: Protecting {SGX} enclaves from practical side-channel attacks. In *2018 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)*. 227–240.
- [57] Joongun Park, Naeyeong Kang, Taehoon Kim, Youngjin Kwon, and Jaehyuk Huh. 2020. Nested enclave: supporting fine-grained hierarchical isolation with SGX. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 776–789.
- [58] Gregory F Pfister. 2001. An introduction to the infiniband architecture. *High performance mass storage and parallel I/O* 42, 617–632 (2001), 102.
- [59] Saurav Prakash, Hanieh Hashemi, Yongqin Wang, Murali Annavaram, and Amir Salman Avestimehr. 2020. Mitigating byzantine attacks in federated learning. *arXiv preprint arXiv:2010.07541* (2020).
- [60] Monika Rani, Riju Nayak, and OP Vyas. 2015. An ontology-based adaptive personalized e-learning system, assisted by software agents on cloud storage. *Knowledge-Based Systems* 90 (2015), 33–48.
- [61] M Sadeq Riaz, Bitu Darvish Rouani, and Farinaz Koushanfar. 2019. Deep learning on private data. *IEEE Security & Privacy* 17, 6 (2019), 54–63.
- [62] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115, 3 (2015), 211–252.
- [63] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4510–4520.
- [64] Jeffrey Schwartz. 2021. Intel Makes 3rd Gen Xeon Scalable Processor Rollout Official. <https://www.channelfutures.com/data-centers/intel-makes-3rd-gen-xeon-scalable-processor-rollout-official>
- [65] Tom Shanley. 2003. *InfiniBand network architecture*. Addison-Wesley Professional.
- [66] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. 1310–1321.
- [67] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [68] Jinhyun So, Basak Guler, and A Salman Avestimehr. 2020. Byzantine-Resilient Secure Federated Learning. *arXiv preprint arXiv:2007.11115* (2020).
- [69] Nancy Spruill. 1983. The confidentiality and analytic usefulness of masked business microdata. *Proceedings of the Section on Survey Research Methods*, 1983 (1983), 602–607.
- [70] ADP Team. 2017. Learning with privacy at scale. *Apple Mach. Learn. J* 1, 9 (2017).
- [71] SGX team. 2021. *Intel SGX in clouds*. <https://software.intel.com/content/www/us/en/develop/topics/software-guard-extensions.html>
- [72] Florian Tramèr and Dan Boneh. 2018. Slalom: Fast, Verifiable and Private Execution of Neural Networks in Trusted Hardware. In *International Conference on Learning Representations*.
- [73] Stavros Volos, Kapil Vaswani, and Rodrigo Bruno. 2018. Graviton: Trusted execution environments on gpus. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 681–696.
- [74] Sameer Wagh, Divya Gupta, and Nishanth Chandran. 2019. Secureenn: 3-party secure computation for neural network training. *Proceedings on Privacy Enhancing Technologies* 2019, 3 (2019), 26–49.
- [75] Wenhao Wang, Guoxing Chen, Xiaorui Pan, Yinqian Zhang, XiaoFeng Wang, Vincent Bindschaedler, Haixu Tang, and Carl A Gunter. 2017. Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 2421–2434.
- [76] Yuanzhong Xu, Weidong Cui, and Marcus Peinado. 2015. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 640–656.
- [77] Zhenyu Xu, Thomas Maudlin, Zheyi Yao, Shuyi Pei, Tao Wei, and Qing Yang. 2020. A bus authentication and anti-probing architecture extending hardware trusted computing base off CPU chips and beyond. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 749–761.
- [78] Guandao Yang, Tianyi Zhang, Polina Kirichenko, Junwen Bai, Andrew Gordon Wilson, and Chris De Sa. 2019. SWALP: Stochastic weight averaging in low precision training. In *International Conference on Machine Learning*. PMLR, 7015–7024.
- [79] Qian Yu, Songze Li, Netanel Raviv, Seyed Mohammadreza Mousavi Kalan, Mahdi Soltanolkotabi, and Salman A Avestimehr. 2019. Lagrange coded computing: Optimal design for resiliency, security, and privacy. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 1215–1225.
- [80] Wei Zhang, Suyog Gupta, Xiangru Lian, and Ji Liu. 2015. Staleness-aware asyn-sgd for distributed deep learning. *arXiv preprint arXiv:1511.05950* (2015).
- [81] Jiajun Zhu, David I Ferguson, and Dmitri A Dolgov. 2014. System and method for predicting behaviors of detected objects. US Patent 8,660,734.
- [82] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep leakage from gradients. In *Advances in Neural Information Processing Systems*. 14747–14756.