

2021

## Aspis: A Robust Detection System for Distributed Learning

Konstantinos Konstantinidis  
*Iowa State University*, [kostas@iastate.edu](mailto:kostas@iastate.edu)

Aditya Ramamoorthy  
*Iowa State University*, [adityar@iastate.edu](mailto:adityar@iastate.edu)

Follow this and additional works at: [https://lib.dr.iastate.edu/ece\\_pubs](https://lib.dr.iastate.edu/ece_pubs)



Part of the [Artificial Intelligence and Robotics Commons](#)

The complete bibliographic information for this item can be found at [https://lib.dr.iastate.edu/ece\\_pubs/314](https://lib.dr.iastate.edu/ece_pubs/314). For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

---

This Article is brought to you for free and open access by the Electrical and Computer Engineering at Iowa State University Digital Repository. It has been accepted for inclusion in Electrical and Computer Engineering Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact [digirep@iastate.edu](mailto:digirep@iastate.edu).

---

# Aspis: A Robust Detection System for Distributed Learning

## Abstract

State of the art machine learning models are routinely trained on large scale distributed clusters. Crucially, such systems can be compromised when some of the computing devices exhibit abnormal (Byzantine) behavior and return arbitrary results to the parameter server (PS). This behavior may be attributed to a plethora of reasons including system failures and orchestrated attacks. Existing work suggests robust aggregation and/or computational redundancy to alleviate the effect of distorted gradients. However, most of these schemes are ineffective when an adversary knows the task assignment and can judiciously choose the attacked workers to induce maximal damage. Our proposed method Aspis assigns gradient computations to worker nodes using a subset-based assignment which allows for multiple consistency checks on the behavior of a worker node. Examination of the calculated gradients and post-processing (clique-finding in an appropriately constructed graph) by the central node allows for efficient detection and subsequent exclusion of adversaries from the training process. We prove the Byzantine resilience and detection guarantees of Aspis under weak and strong attacks and extensively evaluate the system on various large-scale training scenarios. The main metric for our experiments is the test accuracy for which we demonstrate significant improvement of about 30% compared to many state-of-the-art approaches on the CIFAR-10 dataset. The corresponding reduction of the fraction of corrupted gradients ranges from 16% to 98%.

## Disciplines

Artificial Intelligence and Robotics

## Comments

This is a pre-print of the article Konstantinidis, Konstantinos, and Aditya Ramamoorthy. "Aspis: A Robust Detection System for Distributed Learning." *arXiv preprint arXiv:2108.02416* (2021). Posted with permission.

## Creative Commons License



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

---

# Aspis: A Robust Detection System for Distributed Learning

---

**Konstantinos Konstantinidis**

Department of Electrical and Computer Engineering  
Iowa State University  
Ames, IA 50011 USA  
kostas@iastate.edu

**Aditya Ramamoorthy**

Department of Electrical and Computer Engineering  
Iowa State University  
Ames, IA 50011 USA  
adityar@iastate.edu

## Abstract

State of the art machine learning models are routinely trained on large scale distributed clusters. Crucially, such systems can be compromised when some of the computing devices exhibit abnormal (*Byzantine*) behavior and return arbitrary results to the *parameter server* (PS). This behavior may be attributed to a plethora of reasons including system failures and orchestrated attacks. Existing work suggests *robust aggregation* and/or computational *redundancy* to alleviate the effect of distorted gradients. However, most of these schemes are ineffective when an adversary knows the task assignment and can judiciously choose the attacked workers to induce maximal damage. Our proposed method *Aspis* assigns gradient computations to worker nodes using a subset-based assignment which allows for multiple consistency checks on the behavior of a worker node. Examination of the calculated gradients and post-processing (clique-finding in an appropriately constructed graph) by the central node allows for efficient detection and subsequent exclusion of adversaries from the training process. We prove the Byzantine resilience and detection guarantees of *Aspis* under weak and strong attacks and extensively evaluate the system on various large-scale training scenarios. The main metric for our experiments is the test accuracy for which we demonstrate significant improvement of about 30% compared to many state-of-the-art approaches on the CIFAR-10 dataset. The corresponding reduction of the fraction of corrupted gradients ranges from 16% to 98%.

## 1 Introduction

The increased sizes of data sets and associated model complexities have established distributed training setups as the de facto method for training models at scale. A typical setup consists of one central machine (*parameter server* or PS) and multiple worker machines. The PS coordinates the protocol by communicating parameters and maintaining the global copy of the model. The workers compute gradients of the loss with respect to the optimization parameters and transmit them to the PS. The PS then updates the model. This is an iterative process and is repeated until convergence. Several frameworks including MXNet [9], CNTK [27] and PyTorch [22] support this architecture.

Despite the speedup benefits of such distributed settings, they are prone to so-called *Byzantine* failures, i.e., when a set of workers return malicious or erroneous computations. Faulty workers may modify their portion of the data and/or models arbitrarily. This can happen on purpose due to adversarial

attacks or inadvertently due to hardware or software failures, such as bit-flipping in the memory or communication media. As a result, the distorted gradients can derail the optimization and lead to low test accuracy. Devising training algorithms that are resilient to such failures and which can efficiently *aggregate* the received gradients has inspired a series of contributions in the area [1, 10, 4, 15].

In order to alleviate the malicious effects, some existing papers use majority voting and median-based defenses [11, 35, 34, 31, 4, 10]. Other works replicate the gradient tasks across the cluster [36, 23, 8, 12, 19]. Finally, some schemes try to rank and/or detect the Byzantines [8, 33, 25].

### 1.1 Contributions

In this work, we present *Aspis*, our detection-based Byzantine resilience scheme for distributed training; no potentially negative societal impacts are expected from our work (*cf.* Appendix Section A.5.6). *Aspis* uses a combination of *redundancy* and *robust aggregation*. Unlike previous methods, the redundant subset-based assignment scheme for gradient computations is judiciously chosen such that the PS can perform global *consistency checks* on the behavior the worker nodes by examining the returned gradients. By performing clique-finding in appropriate graphs, the PS can perform *detection* as a first line of defense to exclude adversaries from further being considered during the aggregation.

*Aspis* has the following salient features: (i) Under weak attacks where the Byzantines act independently (without significant collusion) they will always be detected by the proposed novel clique-based algorithm. (ii) Furthermore, *Aspis* is resilient to stronger attacks where adversaries collude in optimal ways; these are much stronger than those considered in prior work. In particular, instead of simulating a random set of adversaries [23, 8], we have crafted a non-trivial attack specifically designed for our system such that the adversaries can evade detection and potentially corrupt more gradient values.

For both weak and strong attacks we provide theoretical guarantees on the fraction of corrupted gradients for *Aspis*. Comparisons with other methods indicate reductions ranging from 16% to 98%.

Finally, we present exhaustive top-1 classification accuracy results on the CIFAR-10 dataset for a variety of gradient distortion attacks coupled with choice/behavior patterns of the adversarial nodes. Our results indicate an average 30% accuracy increase on large scale image classification training (CIFAR-10) [20] under the most sophisticated attacks. In summary, the performance gap between *Aspis* and other methods is especially stark in the strong attack scenario.

### 1.2 Related work

Existing papers consider a wide range of assumptions with respect to the maximum number of adversaries, their ability to collude and their possession of knowledge involving the data assignment and the defense protocol and whether the adversarial machines are chosen at random or systematically. In this work, we will assume strong adversarial models as in prior work [2, 31, 34, 13, 28, 19].

The first category of prior work in the area is called *robust aggregation* [11, 35, 34, 31, 4, 10]. These methods provide robustness guarantees up to a constant fraction of the nodes being adversarial. However, this fraction is usually very small and the guarantees are limited (*e.g.*, only establishing convergence in the limit [4, 13]) which compromises their effectiveness in practical systems. Also, they require significant asymptotic complexity [31] and strict assumptions such as convexity of the loss that need to be adjusted for each individual training algorithm. Some popular robust aggregators are based on *trimmed mean* [31, 34, 13] and return a subset of the values which are closest to the median on an element-wise basis. *Auror* in [28] runs *k-means* clustering on the gradients and outputs the mean of the largest cluster. In *signSGD* [3], workers retain only the sign information and the PS uses majority voting aggregation. *Krum* in [4] determines a single honest worker whose gradient minimizes the distance to its  $k \in \mathbb{N}$  nearest neighbors. *Krum* may converge to

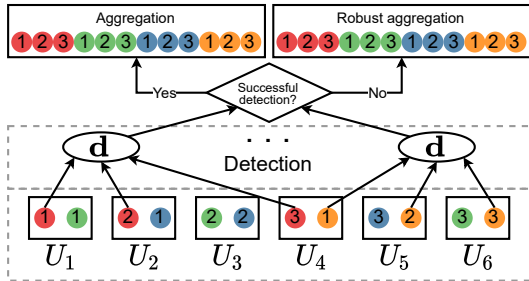


Figure 1: Aggregation of gradients on a cluster.

an *ineffectual* model in non-convex high dimensional problems and *Bulyan* [13] is proposed as an alternative. However, Bulyan is designed to defend only up to a small fraction of corrupted workers.

The second category is based on *redundancy* and it establishes resilience by assigning each gradient task to more than one node [36, 23, 8, 12, 19]. Existing techniques are sometimes combined with robust aggregation [23]. Fundamentally, these methods require higher computation load per worker but they come with stronger guarantees of correcting the erroneous gradients. It is important to note that most schemes in this category can be made to fail by a powerful *omniscient* adversary which can mount judicious attacks [19]. *DRACO* [8] builds on [24, 29] and uses majority vote and Fourier decoders to filter out the adversarial effects. With  $q$  Byzantines, it guarantees recovery as if the system had no adversaries, when each task is replicated  $r \geq 2q + 1$  times; it is not applicable if this bound is violated. Nonetheless, this requirement is very restrictive for the common assumption that up to half of the workers can be Byzantine. *DETOX* in [23] extends DRACO and performs multiple stages of gradient filtering requiring smaller redundancy. However, its resilience guarantees depend heavily on a “random choice” of the adversaries. Subsequent work in [19] has crafted simple attacks to make this aggregator fail under a more careful choice of adversaries. *ByzShield* in [19] borrows ideas from design theory to redundantly assign the tasks in an effort to minimize the distortion fraction.

A third category focuses on *ranking* and/or *detection* [8, 33, 25]. *Zeno* in [33] ranks each worker using a score which depends on the estimated loss reduction and the magnitude of the update. *Zeno* requires strict assumptions on the smoothness of the loss and the variance of the gradient estimates in order to tolerate an adversarial majority in the cluster. Similarly, *ByGARS* [25] computes reputation scores for the computing nodes based on an auxiliary dataset; these scores are used to weigh the contribution of each gradient to the model update. Its convergence result depend on the assumption of a strongly convex loss. In contrast, our method does not rely on such restrictive assumptions.

## 2 Distributed training formulation

The formulation we discuss is standard in distributed deep learning. Assume a loss function  $l_i(\mathbf{w})$  for the  $i^{\text{th}}$  sample where  $\mathbf{w} \in \mathbb{R}^d$  is the parameter set of the model. We use *mini-batch Stochastic Gradient Descent* (SGD) to minimize the loss over the entire dataset, i.e.,

$$\min_{\mathbf{w}} L(\mathbf{w}) = \min_{\mathbf{w}} \frac{1}{n} \sum_{i=1}^n l_i(\mathbf{w})$$

where  $n$  is the dataset size. Initially,  $\mathbf{w}$  is randomly set to  $\mathbf{w}_0$  ( $\mathbf{w}_t$  is the model state at the end of iteration  $t$ ). A random *batch*  $B_t$  of  $b$  samples is chosen to perform the update in the  $t^{\text{th}}$  iteration. Thus,

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \frac{1}{|B_t|} \sum_{i \in B_t} \nabla l_i(\mathbf{w}_t) \quad (1)$$

where  $\eta_t$  is the learning rate of the  $t^{\text{th}}$  iteration. The workers, denoted  $U_1, U_2, \dots, U_K$ , compute gradients on subsets of the batch. The training is *synchronous*, i.e., the PS waits for all  $K$  workers to return before performing an update. It stores the dataset and the model and coordinates the protocol.

**Task assignment:** Each batch  $B_t$  is split into  $f$  disjoint files  $\{B_{t,i}\}_{i=0}^{f-1}$ . These files are then assigned to the workers according to our placement policy. Computational redundancy is introduced by assigning a given file to  $r > 1$  workers; we will occasionally use the term *group* (of the assigned workers) to refer to a file. It follows that each worker is responsible for  $l = fr/K$  files ( $l$  is the *computational load*). We let  $\mathcal{N}(U_j)$  to be the set of files assigned to worker  $U_j$  and  $\mathcal{N}(B_{t,i})$  to be the group of workers that are assigned file  $B_{t,i}$ . Our placement scheme is such that  $\mathcal{N}(B_{t,i})$  uniquely identifies the file  $B_{t,i}$ ; thus, we will sometimes refer to the file  $B_{t,i}$  by its worker assignment,  $\mathcal{N}(B_{t,i})$ . The actual placement algorithm will be presented in Section 3.

**Adversary model:** We assume that at most  $q$  workers can be adversarial, i.e., they can return arbitrary values to the PS. The workers know the data assignment of all nodes, the parameters  $\mathbf{w}_t$  and the defense at *every* iteration (*omniscient* attack); they can also collude. The adversarial machines may change at every single iteration. We will suppose that  $q < K/2$ . We emphasize that our attack setting is more powerful than random failures considered in related redundancy-based work [23, 8]. For each

---

**Algorithm 1:** Aspis subset-based file assignment.

---

**Input:** Batch size  $b$ , computation load  $l$ , redundancy  $r$  and worker set  $\mathcal{U}$ ,  $|\mathcal{U}| = K$ .

- 1 PS partitions batch  $B_t$  into  $f = \binom{K}{r}$  disjoint files of  $b/f$  samples each  
$$B_t = \{B_{t,i} : i = 0, 1, \dots, f - 1\}.$$
  - 2 PS constructs all subsets  $S_0, S_1, \dots, S_{\binom{K}{r}}$  of  $\mathcal{U} = \{U_1, U_2, \dots, U_K\}$  such that  $\forall i, |S_i| = r$ .
  - 3 **for**  $i = 0$  **to**  $f - 1$  **do**
  - 4     PS identifies all workers in group  $S_i = \{U_{j_1}, U_{j_2}, \dots, U_{j_r}\}$  and assigns the file indexed with  $i$  in  $B_t$  to all of them. Formally,  $\mathcal{N}(U_j) = \mathcal{N}(U_j) \cup \{B_{t,i}\}$  for  $j \in \{j_1, j_2, \dots, j_r\}$ .
  - 5 **end**
- 

assigned file  $B_{t,i}$  a worker  $U_j$  will return the value  $\hat{\mathbf{g}}_{t,i}^{(j)}$  to the PS. Then,

$$\hat{\mathbf{g}}_{t,i}^{(j)} = \begin{cases} \mathbf{g}_{t,i} & \text{if } U_j \text{ is honest,} \\ * & \text{otherwise,} \end{cases} \quad (2)$$

where  $\mathbf{g}_{t,i}$  is the sum of the loss gradients on all samples in file  $B_{t,i}$ , i.e.,

$$\mathbf{g}_{t,i} = \sum_{j \in B_{t,i}} \nabla l_j(\mathbf{w}_t)$$

and  $*$  is any arbitrary vector in  $\mathbb{R}^d$ . Within this setup, we examine multiple adversarial scenarios.

**Training:** We will refer to Figure 1 for this exposition. There are  $K = 6$  machines and  $f = 4$  distinct files (represented by colored circles) replicated  $r = 3$  times<sup>1</sup>. Each of the workers is assigned to  $l = 2$  files and computes the sum of gradients (or an distorted value) on each of them. The “**d**” ellipses refer to detection operations the PS performs immediately after receiving all the gradients.

The algorithm starts with the assignment of files to workers. Subsequently, each worker  $U_i$  will compute all  $l$  file gradients that involve its assigned files  $\mathcal{N}(U_j)$  and return them to the PS. In every iteration, the PS will initially run our detection algorithm in an effort to identify the  $q$  adversaries. Then, the PS will act differently depending on the result of detection as follows.

- *Case 1: Successful detection.* The PS will ignore all detected faulty machines and keep only the gradients from the remaining workers. Assume that  $h$  workers  $U_{i_1}, U_{i_2}, \dots, U_{i_h}$  have been identified as honest. For each of the  $f$  files if there is at least one honest worker that processed it, the PS will pick one of the “honest” gradient values. The chosen gradients are then averaged for the update (cf. Eq. (1)). For instance, in Figure 1, assume that  $U_1, U_2$  and  $U_4$  have been identified as faulty. During aggregation, the PS will ignore the red file as all 3 copies have been compromised. For the orange file, it will pick either the gradient computed by  $U_5$  or  $U_6$  as both of them are honest values.

- *Case 2: Unsuccessful detection.* During aggregation, the PS will perform a majority vote across the computations of each file. Recall that each file has been processed by  $r$  workers. For each such file  $B_{t,i}$ , the PS decides a majority value  $\mathbf{m}_i$

$$\mathbf{m}_i := \text{majority} \left\{ \hat{\mathbf{g}}_{t,i}^{(j)} : U_j \in \mathcal{N}(B_{t,i}) \right\}. \quad (3)$$

Assume that  $r$  is odd and let  $r' = \frac{r+1}{2}$ . Under the rule in Eq. (3), the gradient on a file is distorted only if at least  $r'$  of the computations are performed by Byzantines. Following the majority vote, we will further filter the gradients coordinate-wise median and will be referring to the combination of these two steps as *robust aggregation*; a similar setup was considered in [23, 19]. For example, in Figure 1, all returned values for the red file will be evaluated by a majority vote function on the PS which decides a single output value; a similar voting is done for the other 3 files. After the voting process, Aspis applies coordinate-wise median on the “winning” gradients  $\mathbf{m}_i, i = 0, 1, \dots, f - 1$ .

The details of this procedure are described in Algorithm 3 in Appendix Section A.1. In both cases, we ensure that there are no floating point precision issues in our implementation, i.e., all honest workers assigned to  $B_{t,i}$  will return the exact same gradient. In general, even if such issues arise they can easily be handled. In Case 1, if the honest gradients for a file were different the PS could average them. In Case 2, we could cluster the gradients and compute the average of the largest cluster.

---

<sup>1</sup>Some arrows and ellipses have been omitted from Figure 1; however, all files will be going through detection.

---

**Algorithm 2:** Proposed Aspis graph-based detection.

---

**Input:** Computed gradients  $\hat{\mathbf{g}}_{i,i}^{(j)}$ ,  $i = 0, 1, \dots, f - 1$ ,  $j = 1, 2, \dots, K$ , redundancy  $r$  and empty graph  $\mathbf{G}$  with worker vertices  $\mathcal{U}$ .

- 1 **for** each pair  $(U_{j_1}, U_{j_2})$ ,  $j_1 \neq j_2$  of workers **do**
- 2     PS computes the number of agreements  $\alpha^{(j_1, j_2)}$  of the pair  $U_{j_1}, U_{j_2}$  on the gradient value.
- 3     **if**  $\alpha^{(j_1, j_2)} = \binom{K-2}{r-2}$  **then**
- 4         | Connect vertex  $U_{j_1}$  to vertex  $U_{j_2}$  in  $\mathbf{G}$ .
- 5     **end**
- 6 **end**
- 7 PS enumerates all  $k$  maximum cliques  $M_{\mathbf{G}}^{(1)}, M_{\mathbf{G}}^{(2)}, \dots, M_{\mathbf{G}}^{(k)}$  in  $\mathbf{G}$ .
- 8 **if** there is a unique maximum clique  $M_{\mathbf{G}}$  ( $k = 1$ ) **then**
- 9     | PS determines the honest workers  $H = M_{\mathbf{G}}$  and the adversarial machines  $\hat{A} = \mathcal{U} - M_{\mathbf{G}}$ .
- 10 **end**
- 11 **else**
- 12     | PS declares unsuccessful detection.
- 13 **end**

---

### 3 Task assignment

In this section, we propose our technique which determines the allocation of gradient tasks to workers in Aspis. Let  $\mathcal{U}$  be the set of workers. Our scheme has  $|\mathcal{U}| \leq f$  (i.e., fewer workers than files).

To allocate the batch of an iteration,  $B_t$ , to the  $K$  workers, first, we will partition  $B_t$  into  $f = \binom{K}{r}$  disjoint files  $B_{t,0}, B_{t,1}, \dots, B_{t,f-1}$ ; recall that  $r$  is the redundancy. Following this, we associate each file with exactly one of the subsets  $S_0, S_1, \dots, S_{\binom{K}{r}}$  of  $\{U_1, U_2, \dots, U_K\}$  each of cardinality  $r$ , essentially using a bijection. Each file contains  $b/f$  samples. The details are specified in Algorithm 1. The following example showcases the proposed protocol of Algorithm 1.

**Example 1.** Consider  $K = 7$  workers  $U_1, U_2, \dots, U_7$  and  $r = 3$ . Based on our protocol, the  $f = \binom{7}{3} = 35$  files of each batch  $B_t$  would be associated one-to-one with 3-subsets of  $\mathcal{U}$ , e.g., the subset  $S_0 = \{U_1, U_2, U_3\}$  corresponds to file  $B_{t,0}$  and will be processed by  $U_1, U_2$  and  $U_3$ .

**Remark 1.** Our proposed task assignment ensures that every pair of workers processes  $\binom{K-2}{r-2}$  files. Moreover, the number of adversaries is  $q < K/2$ . Thus, upon receiving the gradients from the workers, the master node can examine them for consistency and flag certain nodes as adversarial if their computed gradients differ from a *lot* of the other nodes. We use this intuition to detect and mitigate the adversarial effects and to obtain an upper bound on the fraction of corrupted files.

### 4 Adversarial detection

The PS will run our detection method in every iteration as our model assumes that the adversaries can be different across different steps. The following analysis applies to any iteration; thus, the iteration index  $t$  will be omitted from most of the notation used in this section. Let the current set of adversaries be  $A \subset \{U_1, U_2, \dots, U_K\}$  with  $|A| = q$ ; also, let  $H$  be the honest worker set. The set  $A$  is unknown but our goal is to provide an estimate  $\hat{A}$  of it. Ideally, the two sets should be identical. In general, depending on the adversarial behavior, we will be able provide a set  $\hat{A}$  such that  $\hat{A} \subseteq A$ . For each file, there is a group of  $r$  workers which have processed it and there are  $\binom{r}{2}$  pairs of workers in each group. Each such pair may or may not agree on the gradient value for the file. For an iteration, let us encode the agreement of workers  $U_{j_1}, U_{j_2}$  on a common file  $i$  of them as

$$\alpha_i^{(j_1, j_2)} := \begin{cases} 1 & \text{if } \hat{\mathbf{g}}_i^{(j_1)} = \hat{\mathbf{g}}_i^{(j_2)}, \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

where the iteration subscript  $t$  is skipped for the computed gradients. Then, across all files, let us denote the total number of agreements between a pair of workers by

$$\alpha^{(j_1, j_2)} := \sum_{i \in \mathcal{N}(U_{j_1}) \cap \mathcal{N}(U_{j_2})} \alpha_i^{(j_1, j_2)}. \quad (5)$$



(a) Unique max-clique, detection succeeds.

(b) Two max-cliques, detection fails.

Figure 2: Detection graph  $\mathbf{G}$  for  $K = 7$  workers among which  $U_1, U_2$  and  $U_3$  are the adversaries.

Since the placement is known, the PS can always perform the above computation. Next, we form an undirected graph  $\mathbf{G}$  whose vertices correspond to all workers  $\{U_1, U_2, \dots, U_K\}$ . An edge  $(U_{j_1}, U_{j_2})$  exists in  $\mathbf{G}$  only if the computed gradients of  $U_{j_1}$  and  $U_{j_2}$  match in all their  $\binom{K-2}{r-2}$  common groups.

A *clique* in an undirected graph is defined as a subset of vertices in which there is an edge between any pair of them. A *maximal clique* is one that cannot be enlarged by adding additional vertices to it. A *maximum clique* is a clique such that there is no clique with more vertices in the given graph. We note that the set of honest workers  $H$  will pair-wise agree everywhere. In particular, this implies that the subset  $H$  forms a clique (of size  $K - q$ ) within  $\mathbf{G}$ . The clique containing the honest workers may not be maximal. However, it will have a size at least  $K - q$ . Let the maximum clique on  $\mathbf{G}$  be  $M_{\mathbf{G}}$ . Any worker  $U_j$  with  $\deg(U_j) < K - q - 1$  will not belong to a maximum clique and can straight away be eliminated as a “detected” adversary.

The essential idea of our detection is to run a clique-finding algorithm on  $\mathbf{G}$  (summarized in Algorithm 2). If we find a unique maximum clique, we declare it to be the set of honest workers; the gradients from the detected adversaries are ignored. On the other hand, if there is more than one maximum clique, we resort to the robust aggregation technique discussed in Section 2. Let us denote the number of distorted tasks upon Aspis detection and aggregation by  $c^{(q)}$  and its maximum value (under the worst-case attack) by  $c_{\max}^{(q)}$ . The *distortion fraction* is  $\epsilon := c^{(q)} / f$ . Clique-finding is well-known to be an NP-complete problem [18]. Nevertheless, there are fast practical algorithms that have excellent performance on graphs even up to hundreds of nodes [6, 14, 7]. Specifically, the authors of [14] have shown that their proposed algorithm which enumerates all maximal cliques has similar complexity as other methods [30, 26] which are used to find a single maximum clique. We utilize this algorithm which is proven to be optimal. Our extensive experimental evidence suggests that clique-finding is not a computation bottleneck for the size and structure of the graphs that Aspis uses. We have experimented with clique-finding on a graph of  $K = 100$  workers and  $r = 5$  for different values of  $q$ ; in all cases, enumerating all maximal cliques took no more than 15 milliseconds. These experiments and asymptotic complexity of entire protocol are addressed in Appendix Section A.2.

#### 4.1 Weak adversarial strategy

We first consider a class of weak attacks where the Byzantine nodes attempt to distort the gradient on any file that they participate in. For instance, a node may try to return arbitrary gradients on all its files. A more sophisticated strategy (but still weak) for the node would be to match the result with other adversarial nodes that process the same file. In both cases, it is clear that the Byzantine node will disagree with at least  $K - q$  honest nodes and thus the degree of the node in  $\mathbf{G}$  will be at most  $q - 1 < K - q - 1$  and it will not be part of the maximum clique. Thus, each of the adversaries will be detected and their returned gradients will not be considered further. The algorithm declares the (unique) maximum clique as honest and proceeds to aggregation (*cf.* Section 2). The only files that can be distorted in this case are those that consist exclusively of adversarial nodes.

Figure 2a (corresponding to Example 1) shows an example where in a cluster of size  $K = 7$  the  $q = 3$  adversaries are  $A = \{U_1, U_2, U_3\}$  and the remaining workers are honest with  $H = \{U_4, U_5, U_6, U_7\}$ . In this case, the unique maximum clique is  $M_{\mathbf{G}} = H$  and detection is successful. Under this attack, the number of distorted tasks are those whose all copies have been compromised, i.e.,  $c^{(q)} = \binom{q}{r}$ .



**Remark 2.** We emphasize that even though we call this attack “weak” this is the attack model considered in several prior works [23, 8]. Furthermore, to our best knowledge most of them have not considered the adversarial problem from the lens of detection.

## 4.2 Optimal adversarial strategy

Our second scenario is a strong one and involves adversaries who collude in the “best” way possible while knowing the full details of our detection algorithm. In the discussion below we provide an upper bound on the number of files that can be corrupted in this case and demonstrate a strategy that the adversarial workers can follow to achieve this upper bound.

Let us index the  $q$  adversaries in  $A = \{A_1, A_2, \dots, A_q\}$  and the honest workers in  $H$ . We say that two workers  $U_{j_1}$  and  $U_{j_2}$  disagree if there is no edge between them in the corresponding graph  $\mathbf{G}$ . The non-existence of an edge between  $U_{j_1}$  and  $U_{j_2}$  only means that they disagree in *at least one* of the  $\binom{K-2}{r-2}$  files that they jointly participate in. For corrupting the gradients, each adversary has to disagree on the computations with a subset of the honest workers. An adversary may also disagree with other adversaries. Let  $D_i$  denote the set of disagreement workers for adversary  $A_i$ ,  $i = 1, 2, \dots, q$ , where  $D_i$  can contain members from  $A$  and from  $H$ .

Upon formation of  $\mathbf{G}$  we know that a worker  $U_j$  will be flagged as adversarial if  $\deg(U_j) < K - q - 1$ . Therefore to avoid detection, a *necessary* condition is that  $|D_j| \leq q$ . We now upper bound the number of files that can be corrupted under any possible strategy employed by the adversaries. We fall back to robust aggregation in case of more than one maximum clique in  $\mathbf{G}$ . Then, a gradient can only be corrupted if a majority of the workers computing it are adversarial and agree on a wrong value.

For a given file  $F$ , let  $A' \subseteq A$  with  $|A'| \geq r'$  be the set of “active adversaries” in it, i.e.,  $A' \subseteq F$  consists of Byzantines that collude to create a majority that distorts the gradient on it. In this case, the remaining workers in  $F$  belong to  $\cap_{i \in A'} D_i$ , where we note that  $|\cap_{i \in A'} D_i| \leq q$ . Let  $X_j$ ,  $j = r', r' + 1, \dots, r$  denote the subset of files where the set of active adversaries is of size  $j$ ; note that  $X_j$  depends on the disagreement sets  $D_i$ ,  $i = 1, 2, \dots, q$ . Formally,

$$X_j = \{F : \exists A' \subseteq A \cap F, |A'| = j, \text{ and } \forall U_j \in F \setminus A', U_j \in \cap_{i \in A'} D_i\}. \quad (6)$$

Then, for a given choice of disagreement sets, the number of files that can be corrupted is given by  $|\cup_{j=r'}^r X_j|$ . We obtain an upper bound on the maximum number of corrupted files by maximizing this quantity with respect to the choice of  $D_i$ ,  $i = 1, 2, \dots, q$ , i.e.,

$$c_{\max}^{(q)} = \max_{D_i, |D_i| \leq q, i=1,2,\dots,q} |\cup_{j=r'}^r X_j| \quad (7)$$

where the maximization is over the choices of the disagreement sets  $D_1, D_2, \dots, D_q$ . An intuitive adversarial strategy based on Eq. (6) would be to maximize the set  $\bigcap_{i \in A' \cap F} D_i$  for every possible file  $F$ .

In order to achieve this, for all groups  $F$ , the adversaries will need to fix a subset of  $q$  non-adversaries, say  $D \subset H$ , which will be the set of workers with which all adversaries will disagree, i.e.,  $D_i = D$  for  $i = 1, 2, \dots, q$ . We present the following theorem (see Appendix Section A.3 for proof).

**Theorem 1.** Consider a training cluster of  $K$  workers with  $q$  adversaries using Algorithm 1 to assign the  $f = \binom{K}{r}$  files to workers and Algorithm 2 for adversary detection. Under an optimal adversary model the maximum number of files that can be corrupted is

$$c_{\max}^{(q)} = \frac{1}{2} \binom{2q}{r}. \quad (8)$$

Furthermore, if all adversaries fix a set  $D \subset H$  of honest workers with which they will consistently disagree on the gradient value (by distorting it), this upper bound can be achieved.

In particular, the proposed attack is optimal and there is no other attack that can corrupt more files under the Aspis algorithm. One such attack is carried out in Figure 2b for the setup of Example 1. The adversaries  $A = \{U_1, U_2, U_3\}$  consistently disagree with the workers in  $D = \{U_4, U_5, U_6\} \subset H$ . The ambiguity as to which of the two maximum cliques ( $\{U_1, U_2, U_3, U_7\}$  or  $\{U_4, U_5, U_6, U_7\}$ ) is the honest one makes an accurate detection impossible; robust aggregation will be performed instead.




Figure 3: *ALIE* distortion under optimal attack scenarios (CIFAR-10),  $K = 15$ .

## 5 Large-scale deep learning experiments

### 5.1 Experiment setup

We have evaluated the performance of our method and competing techniques in classification tasks on large datasets on Amazon EC2 clusters. The project is written in PyTorch [22] and uses the MPICH library for communication between the different nodes. We worked with the CIFAR-10 dataset [20] using the ResNet-18 [17] model. We used a cluster of  $K = 15$  workers and redundancy  $r = 3$ . This results in partitioning each batch into  $f = 455$  files. We simulated values of  $q = 2, 4, 6$  during training. Detailed information about the implementation can be found in Appendix Section A.5.

There are two different dimensions we experimentally evaluate with respect to the adversarial setup:

**1) Choice/orchestration of the adversaries:** This involves the different ways in which adversaries are chosen and can work together to inflict damage (*cf.* weak and optimal attacks in Section 4).

**2) Gradient distortion methods:** This dimension is concerned with the method an adversary uses to distort the gradient value (*cf.* Algorithm 3). We use a variety of state-of-the-art methods for distorting the computed gradients. *ALIE* [2] involves communication among the Byzantines in which they jointly estimate the mean  $\mu_i$  and standard deviation  $\sigma_i$  of the batch’s gradient for each dimension  $i$  and subsequently use them to construct a distorted gradient that attempts to distort the median of the results. Another powerful attack is *Fall of Empires (FoE)* [32] which performs “inner product manipulation” to make the inner product between the true gradient and the robust estimator to be negative even when their distance is upper bounded by a small value. *Reversed gradient* distortion returns  $-c\mathbf{g}$  for  $c > 0$ , to the PS instead of the true gradient  $\mathbf{g}$ . The *ALIE* algorithms is, to the best of our knowledge, the most sophisticated attack in literature.

**Competing methods:** We compare our work with *baseline* approaches which do not involve redundancy or majority voting. Their aggregation is applied directly to the  $K$  gradients returned by the workers ( $f = K$ ,  $c_{\max}^{(q)} = q$  and  $\epsilon = q/K$ ). We have tested median-of-means [21], Bulyan [13] and Multi-Krum [11]. If  $c_{\max}^{(q)}$  is the number of adversarial computations then Bulyan requires at least  $4c_{\max}^{(q)} + 3$  total number of computations while the same number for Multi-Krum is  $2c_{\max}^{(q)} + 3$ . These constraints make these methods inapplicable for larger values of  $q$  for which our method is robust. The second class of comparisons is with methods that use redundancy and specifically DETOX [23] for which we show that it can easily fail under malicious scenarios for large  $q$ .

Note that for a baseline scheme all choices of  $A$  are equivalent in terms of the value of  $\epsilon$ . In our comparisons with DETOX we will consider two attack scenarios with respect to the choice of the adversaries. For the *optimal* choice in DETOX, we will use the method proposed in [19] and compare with the attack introduced in Section 4.2. For the *weak* one, we will choose the adversaries such that  $\epsilon$  is minimized in DETOX and compare its performance with the scenario discussed in Section 4.1.

**Metrics:** We are interested in two main metrics, the fraction of distorted files,  $\epsilon$ , and the top-1 test accuracy of the final trained model. We evaluate these metrics for the various competing methods.

### 5.2 Results

We performed simulations on the exact value of  $\epsilon$  (*cf.* Appendix Tables 2 and 3). This analysis is motivated by the experiments and the fact that  $\epsilon$  is a useful indicator of a scheme’s test accuracy.




Figure 4: *Reversed gradient* distortion under optimal attack scenarios (CIFAR-10),  $K = 15$ .




Figure 5: *FoE* optimal attack and median-based defenses (CIFAR-10),  $K = 15$




Figure 6: *Reversed gradient* weak attack and median-based defenses (CIFAR-10),  $K = 15$ .

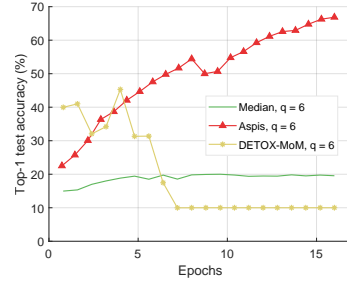


Figure 7: *ALIE* weak attack and median-based defenses (CIFAR-10),  $K = 15$ .

Also, Aspis can achieve values of  $\epsilon$  that are reduced by up to 98% on optimal attacks, compared to prior methods. More details are in the Appendix Section A.4.

**Comparison under optimal attacks:** We compare the different defense algorithms under optimal attack scenarios. Figure 3a compares our scheme Aspis with the baseline implementation of coordinate-wise median ( $\epsilon = 0.133, 0.267$  for  $q = 2, 4$ , respectively) and DETOX with median-of-means ( $\epsilon = 0.2, 0.4$  for  $q = 2, 4$ , respectively) under the ALIE attack. Aspis converges faster and achieves at least a 35% average accuracy overhead (at the end of the training) for both values of  $q$  ( $\epsilon^{Aspis} = 0.004, 0.062$  for  $q = 2, 4$ , respectively). In Figures 3b, 3c we observe similar trends in our experiments with Bulyan and Multi-Krum where Aspis significantly outperforms these techniques. For the current setup, Bulyan is not applicable for  $q = 4$  since  $K = 15 < 4c_{\max}^{(q)} + 3 = 4q + 3 = 19$ . Also, neither Bulyan nor Multi-Krum can be paired with DETOX for  $q \geq 1$  since the inequalities  $f \geq 4c_{\max}^{(q)} + 3$  and  $f \geq 2c_{\max}^{(q)} + 3$  cannot be satisfied. Please refer to Section 5.1 and Appendix Section A.4 for more details on these requirements and the values of  $c_{\max}^{(q)}$  and  $\epsilon$  each scheme can achieve. The results for the reversed gradient attack are shown in Figures 4a, 4b and 4c. Given that this is a much weaker attack [23, 19] all of the schemes considered, including the baseline methods, are expected to perform well; indeed in most of the cases, the model converges to approximately 80% accuracy. However, DETOX fails to converge to high accuracy for  $q = 4$  as in the case of ALIE; one explanation is that  $\epsilon^{DETOX} = 0.4$  for  $q = 4$ . Under the Fall of Empires (FoE) distortion (*cf.* Figure 5) our method still enjoys an accuracy advantage over the baseline and DETOX schemes; this accuracy gap becomes more important as the number of Byzantines in the cluster increases.

**Comparison under weak attacks:** For baseline schemes, the discussion of weak versus optimal choice of the adversaries is not very relevant as any choice of the  $q$  Byzantines can overall distort at most  $q$  out of the  $K$  gradients. Hence, for weak scenarios, we chose to compare mostly with DETOX. The accuracy is reported on Figures 6 and 7 according to which Aspis shows an improvement under attacks on the more challenging end of the spectrum (ALIE). According to Appendix Table 3, Aspis enjoys a fraction  $\epsilon^{Aspis} = 0.044$  while  $\epsilon^{Baseline} = 0.4$  and  $\epsilon^{DETOX} = 0.2$  for  $q = 6$ .

## 6 Conclusions and future work

In this work we have presented Aspis, a Byzantine-resilient distributed scheme that utilizes redundancy and robust aggregation in novel ways to detect adversarial behavior by the worker nodes. Our theoretical analysis and numerical experiments clearly indicate the superior performance of Aspis as compared to prior work. Our numerical experiments indicate that Aspis requires increased computation and communication time as compared to prior work, e.g., note that each worker has to transmit  $l$  gradients instead of 1 in related work [8, 23] (see Appendix Section A.5.4 for a detailed discussion). There are opportunities for reducing the time-overhead. For instance, utilizing GPUs and communication-related algorithmic improvements are worth exploring.

## Acknowledgments and disclosure of funding

This work was supported in part by the National Science Foundation (NSF) under grant CCF-1718470 and grant CCF-1910840.

## References

- [1] Dan Alistarh, Zeyuan Allen-Zhu, and Jerry Li. Byzantine stochastic gradient descent. In *Advances in Neural Information Processing Systems*, December 2018.
- [2] Gilad Baruch, Moran Baruch, and Yoav Goldberg. A little is enough: Circumventing defenses for distributed learning. In *Advances in Neural Information Processing Systems*, pages 8635–8645, December 2019.
- [3] Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, and Anima Anandkumar. signSGD with majority vote is communication efficient and fault tolerant. *arXiv*, 2019. 1810.05291.
- [4] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. In *Advances in Neural Information Processing Systems*, pages 119–129, December 2017.
- [5] Robert S. Boyer and J. Strother Moore. *MJRTY—A Fast Majority Vote Algorithm*. Springer Netherlands, Dordrecht, 1991.
- [6] Coen Bron and Joep Kerbosch. Algorithm 457: Finding all cliques of an undirected graph. *Commun. ACM*, 16(9):575—577, September 1973.
- [7] F. Cazals and C. Karande. A note on the problem of reporting maximal cliques. *Theoretical Computer Science*, 407(1):564–568, November 2008.
- [8] Lingjiao Chen, Hongyi Wang, Zachary Charles, and Dimitris Papailiopoulos. DRACO: Byzantine-resilient distributed training via redundant gradients. In *Proceedings of the 35th International Conference on Machine Learning*, pages 903–912, July 2018.
- [9] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. MXNet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv*, 2015. 1512.01274.
- [10] Yudong Chen, Lili Su, and Jiaming Xu. Distributed statistical machine learning in adversarial settings: Byzantine gradient descent. *Proc. ACM Meas. Anal. Comput. Syst.*, 1(2), December 2017.
- [11] Georgios Damaskinos, El Mahdi El Mhamdi, Rachid Guerraoui, Arsany Hany Abdelmessih Guirguis, and Sébastien Louis Alexandre Rouault. Aggregathor: Byzantine machine learning via robust gradient aggregation. In *Conference on Systems and Machine Learning (SysML) 2019*, page 19, March 2019.
- [12] D. Data, L. Song, and S. Diggavi. Data encoding for Byzantine-resilient distributed gradient descent. In *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 863–870, October 2018.
- [13] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. The hidden vulnerability of distributed learning in Byzantium. In *Proceedings of the 35th International Conference on Machine Learning*, pages 3521–3530, July 2018.
- [14] Tomita Etsuji, Tanaka Akira, and Takahashi Haruhisa. The worst-case time complexity for generating all maximal cliques and computational experiments. *Theoretical Computer Science*, 363(1):28–42, October 2006.
- [15] Nirupam Gupta and Nitin H. Vaidya. Byzantine fault-tolerant parallelized stochastic gradient descent for linear regression. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 415–420, September 2019.
- [16] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference*, pages 11–15, August 2008.
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, June 2016.

- [18] Richard M. Karp. *Reducibility among Combinatorial Problems*. Springer US, Boston, MA, 1972.
- [19] Konstantinos Konstantinidis and Aditya Ramamoorthy. ByzShield: An efficient and robust system for distributed training. In *Machine Learning and Systems 3 (MLSys 2021)*, April 2021.
- [20] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- [21] Stanislav Minsker. Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21(4):2308–2335, November 2015.
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, December 2019.
- [23] Shashank Rajput, Hongyi Wang, Zachary Charles, and Dimitris Papailiopoulos. DETOX: A redundancy-based framework for faster and more robust gradient aggregation. In *Advances in Neural Information Processing Systems*, pages 10320–10330, December 2019.
- [24] Netanel Raviv, Rashish Tandon, Alex Dimakis, and Itzhak Tamo. Gradient coding from cyclic MDS codes and expander graphs. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4302–4310, July 2018.
- [25] Jayanth Regatti, Hao Chen, and Abhishek Gupta. ByGARS: Byzantine sgd with arbitrary number of attackers. *arXiv*, 2020. 2006.13421.
- [26] J.M Robson. Algorithms for maximum independent sets. *Journal of Algorithms*, 7(3):425–440, September 1986.
- [27] Frank Seide and Amit Agarwal. CNTK: Microsoft’s open-source deep-learning toolkit. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 2135, August 2016.
- [28] Shiqi Shen, Shruti Tople, and Prateek Saxena. Auror: Defending against poisoning attacks in collaborative deep learning systems. In *Proceedings of the 32nd Annual Conference on Computer Security Applications*, pages 508—519, December 2016.
- [29] Rashish Tandon, Qi Lei, Alexandros G. Dimakis, and Nikos Karampatziakis. Gradient coding: Avoiding stragglers in distributed learning. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3368–3376, August 2017.
- [30] Robert Endre Tarjan and Anthony E. Trojanowski. Finding a maximum independent set. *SIAM Journal on Computing*, 6(3):537–546, 1977.
- [31] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Generalized Byzantine-tolerant SGD. *arXiv*, 2018. 1802.10116.
- [32] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Fall of empires: Breaking byzantine-tolerant sgd by inner product manipulation. In *35th Conference on Uncertainty in Artificial Intelligence, UAI 2019*, pages 6893–6901, July 2019.
- [33] Cong Xie, Sanmi Koyejo, and Indranil Gupta. Zeno: Distributed stochastic gradient descent with suspicion-based fault-tolerance. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6893–6901, June 2019.
- [34] Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5650–5659, July 2018.
- [35] Dong Yin, Yudong Chen, Kannan Ramchandran, and Peter L. Bartlett. Defending against saddle point attack in Byzantine-robust distributed learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 7074–7084, June 2019.
- [36] Qian Yu, Songze Li, Netanel Raviv, Seyed Mohammadreza Mousavi Kalan, Mahdi Soltanolkotabi, and Salman Avestimehr. Lagrange coded computing: Optimal design for resiliency, security and privacy. *arXiv*, 2019. 1806.00939.

## A Appendix

### A.1 Algorithms

---

**Algorithm 3:** Proposed Aspis aggregation algorithm to alleviate Byzantine effects.

---

**Input:** Dataset of  $n$  samples, batchsize  $b$ , computation load  $l$ , redundancy  $r$ , number of files  $f$ , maximum iterations  $T$ , file assignments  $\{\mathcal{N}(U_i)\}_{i=1}^K$ .

- 1 The PS randomly initializes model's parameters to  $\mathbf{w}_0$ .
- 2 **for**  $t = 1$  to  $T$  **do**
- 3     PS chooses a random batch  $B_t \subseteq \{1, 2, \dots, n\}$  of  $b$  samples, partitions it into  $f$  files  $\{B_{t,i}\}_{i=0}^{f-1}$  and assigns them to workers according to  $\{\mathcal{N}(U_i)\}_{i=1}^K$ . It then transmits  $\mathbf{w}_t$  to all workers.
- 4     **for each worker**  $U_j$  **do**
- 5         **if**  $U_j$  is honest **then**
- 6             **for each file**  $i \in \mathcal{N}(U_j)$  **do**
- 7                  $U_j$  computes the sum of gradients
 
$$\hat{\mathbf{g}}_{t,i}^{(j)} = \sum_{k \in B_{t,i}} \nabla l_k(\mathbf{w}_t).$$
- 8                 **end**
- 9             **else**
- 10                  $U_j$  constructs  $l$  adversarial vectors
 
$$\hat{\mathbf{g}}_{t,i_1}^{(j)}, \hat{\mathbf{g}}_{t,i_2}^{(j)}, \dots, \hat{\mathbf{g}}_{t,i_l}^{(j)}.$$
- 11                 **end**
- 12              $U_j$  returns  $\hat{\mathbf{g}}_{t,i_1}^{(j)}, \hat{\mathbf{g}}_{t,i_2}^{(j)}, \dots, \hat{\mathbf{g}}_{t,i_l}^{(j)}$  to the PS.
- 13     **end**
- 14     PS runs a detection algorithm to identify the adversaries.
- 15     **if detection is successful then**
- 16         Let  $H$  be the detected honest workers. Initialize a non-corrupted gradient set as  $\mathcal{G} = \emptyset$ .
- 17         **for each file in**  $\{B_{t,i}\}_{i=0}^{f-1}$  **do**
- 18             PS chooses the gradient of a worker in  $\mathcal{N}(B_{t,i}) \cap H$  (if non-empty) and adds it to  $\mathcal{G}$ .
- 19         **end**
- 20         
$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \frac{1}{|\mathcal{G}|} \sum_{\mathbf{g} \in \mathcal{G}} \mathbf{g}.$$
- 21     **end**
- 22     **else**
- 23         **for each file in**  $\{B_{t,i}\}_{i=0}^{f-1}$  **do**
- 24             PS determines the  $r$  workers in  $\mathcal{N}(B_{t,i})$  which have processed  $B_{t,i}$  and computes
 
$$\mathbf{m}_i = \text{majority} \left\{ \hat{\mathbf{g}}_{t,i}^{(j)} : j \in \mathcal{N}(B_{t,i}) \right\}.$$
- 25             **end**
- 26             PS updates the model via
 
$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \times \text{median} \{ \mathbf{m}_i : i = 0, 1, \dots, f-1 \}.$$
- 27         **end**
- 28 **end**

---

### A.2 Asymptotic Complexity

If the gradient computation has linear complexity (assuming  $\mathcal{O}(1)$  cost for the gradient computation with respect to one model parameter) and since each worker is assigned to  $l$  files of  $b/f$  samples each, the gradient computation cost at the worker level is  $\mathcal{O}((lb/f)d)$  ( $K$  such computations in parallel). In our schemes, however,  $b$  is a constant multiple of  $f$  and in general  $r < l = \binom{K-1}{r-1}$ ; hence, the complexity becomes  $\mathcal{O}(ld)$  which is similar to other redundancy-based schemes in [23, 19, 8]. The clique-finding problem that follows as part of our detection is NP-complete. However, our

experimental evidence suggests that for the kind of graphs we construct this computation takes a infinitesimal fraction of the execution time. The NetworkX package [16] which we use for enumerating all maximal cliques is based on the algorithm of [14] and has asymptotic complexity  $\mathcal{O}(3^{K/3})$ . We provide extensive simulations of the clique enumeration time under the Aspis file assignment for  $K = 100$  and redundancy  $r = 5$  (cf. Tables 1(a), 1(b) for weak and optimal attack as introduced in Section 4, respectively). We emphasize that this value of  $K$  exceeds by far the typical values of  $K$  of prior work and the number of servers would suffice for the majority of challenging training tasks. Even in this case, the cost of enumerating all cliques is negligible. For this experiment, we used an EC2 instance of type `i3.16xlarge`. The complexity of robust aggregation varies significantly depending on the operator. For example, majority voting can be done in time which scales linearly with the number of votes using *MJRTY* proposed in [5]. In our case, this is  $\mathcal{O}(Kd)$  as the PS needs to use the  $d$ -dimensional input from all  $K$  machines. Krum [4], Multi-Krum [11] and Bulyan [13] are applied to all  $K$  workers by default and require  $\mathcal{O}(K^2(d + \log K))$ .

Table 1: Clique enumeration time in Aspis graph of  $K = 100$  vertices and redundancy  $r = 5$ .

$q$	Time (milliseconds)	$q$	Time (milliseconds)
5	9	5	11
15	7	15	11
25	5	25	9
35	5	35	8
45	5	45	6

1(a) Adversaries carry out weak attack.

1(b) Adversaries carry out optimal attack.

### A.3 Proof of Theorem 1

With  $X_j$  given in (6), assuming  $q \geq r'$ , the number of distorted files is upper bounded by

$$|\cup_{j=r'}^r X_j| \leq \sum_{j=r'}^r |X_j| \text{ (by the union bound).} \quad (9)$$

Note that  $X_j$  consists of those files where the active adversaries  $A'$  are of size  $j$ ; these can be chosen in  $\binom{q}{j}$  ways. The remaining workers in the file belong to  $\cap_{i \in A'} D_i$  where  $|\cap_{i \in A'} D_i| \leq q$ . Thus, the remaining workers can be chosen in at most  $\binom{q}{r-j}$  ways. It follows that

$$|X_j| \leq \binom{q}{j} \binom{q}{r-j}. \quad (10)$$

Therefore,

$$\begin{aligned} c_{\max}^{(q)} &\leq \binom{q}{r'} \binom{q}{r-r'} + \binom{q}{r'+1} \binom{q}{r-(r'+1)} + \dots \\ &\quad + \binom{q}{r-1} \binom{q}{r-(r-1)} + \binom{q}{r} \end{aligned} \quad (11)$$

$$= \sum_{i=r'}^q \binom{q}{i} \binom{q}{r-i} \quad (12)$$

$$= \sum_{i=0}^q \binom{q}{i} \binom{q}{r-i} - \sum_{i=0}^{r'-1} \binom{q}{i} \binom{q}{r-i} \quad (13)$$

$$= \frac{1}{2} \binom{2q}{r}. \quad (14)$$

Eq. (12) follows from the convention that  $\binom{n}{k} = 0$  when  $k > n$  or  $k < 0$ . Eq. (14) follows from Eq. (13) using the following observations



- $\sum_{i=0}^q \binom{q}{i} \binom{q}{r-i} = \sum_{i=0}^r \binom{q}{i} \binom{q}{r-i} = \binom{2q}{r}$  in which the first equality is straightforward to show by taking all possible cases:  $q < r$ ,  $q = r$  and  $q > r$ .
- By symmetry,  $\sum_{i=0}^{r'-1} \binom{q}{i} \binom{q}{r-i} = \sum_{i=r'}^q \binom{q}{i} \binom{q}{r-i} = \frac{1}{2} \binom{2q}{r}$ .

The upper bound in Eq. (11) is met with equality when all adversaries choose the same disagreement set which is a  $q$ -sized subset of the honest workers, i.e.,  $D_i = D \subset H$  for  $i = 1, \dots, q$ . In this case, it can be seen that the sets  $X_j, j = r', \dots, r$  are disjoint, so that (9) is met with equality. Moreover, (10) is also an equality. This finally implies that (11) is also an equality, i.e., this choice of disagreement sets saturates the upper bound.

It can also be seen that in this case the adversarial strategy yields a graph  $\mathbf{G}$  with multiple maximum cliques. To see this, we note that the adversaries in  $A$  agree with all the computed gradients in  $H \setminus D$ . Thus, they form of a clique of  $M_{\mathbf{G}}^{(1)}$  of size  $K - q$  in  $\mathbf{G}$ . Furthermore, the honest workers in  $H$  form another clique  $M_{\mathbf{G}}^{(2)}$  which is also of size  $K - q$ . Thus, the detection algorithm cannot select one over the other and the adversaries will evade detection; and the fallback robust aggregation strategy will apply.

#### A.4 Distortion fraction comparisons with prior work

We have performed simulations of the fraction of distorted files (defined as  $\epsilon = c^{(q)}/f$ ) incurred by Aspis and other competing aggregators. The main motivation of this analysis is that our deep learning experiments (cf. Section 5.2) as well as prior work [19] show that  $\epsilon$  serves as a surrogate of the model’s convergence with respect to accuracy. In addition, our simulations show that Aspis enjoys values of  $\epsilon$  which are as much as 72% lower for the same  $q$  compared to other techniques and this attests to our theoretical robustness guarantees. This comparison involves our work and state-of-the-art schemes under the best-case and worst-case choice of the  $q$  adversaries with respect to the achievable value of  $\epsilon$ .

Let us first discuss the scenario of an *optimal* attack. For Aspis, we used the proposed attack from Section 4.2 and the corresponding computation of  $c^{(q),Aspis}$  of Theorem 1. *DETOX* in [23] employs a redundant assignment followed by majority voting and offers robustness guarantees which crucially rely on a “random assignment” of tasks to workers and on “random choice” of the Byzantines. The authors of [19] have demonstrated the importance of a careful task assignment and observed that redundancy is not sufficient to yield Byzantine resilience dividends. They proposed an optimal choice of the  $q$  Byzantines that maximizes  $\epsilon^{DETOX}$  which we used in our experiments. In short, *DETOX* splits the  $K$  workers into  $K/r$  groups. All workers within a group process the same subset of the batch, specifically containing  $br/K$  samples. This phase is followed by majority voting on a group-by-group basis. The authors of [19] suggested choosing the Byzantines such that at least  $r'$  workers in each group are adversarial in order to distort the corresponding gradients. In this case,  $c^{(q),DETOX} = \lfloor \frac{q}{r'} \rfloor$  and  $\epsilon^{DETOX} = \lfloor \frac{q}{r'} \rfloor \times r/K$ . We also compare with the distortion fraction incurred by *ByzShield* [19] under a worst-case scenario. For this scheme, there is no known optimal attack and its authors performed an exhaustive combinatorial search to find the  $q$  adversaries that maximize  $\epsilon^{ByzShield}$  among all possible options; we follow the same process here to simulate the distortion fraction computation and utilize their scheme based on *mutually orthogonal Latin squares*. The reader can refer to Tables 2(a) and 2(b) for our results. Aspis is able to achieve major reductions in  $\epsilon$ ; for instance,  $\epsilon^{Aspis}$  is reduced by up to 97% compared to  $\epsilon^{Baseline}$ , by up to 98% compared to  $\epsilon^{DETOX}$  and by up to 90% compared to  $\epsilon^{ByzShield}$  in Table 2(a).

We will next introduce the utilized *weak* attacks. For our scheme, we will make an arbitrary choice of  $q$  adversaries which carry out the method introduced in Section 4.1, i.e., they will distort all files and a successful detection is possible. As discussed, the fraction of corrupted gradients is  $\epsilon^{Aspis} = \binom{q}{r} / \binom{K}{r}$  in that case. For *DETOX*, a simple benign attack is used. To that end, let the  $K/r$  groups (files) be  $B_{t,0}, B_{t,1}, \dots, B_{t,K/r-1}$ . Initialize  $A = \emptyset$  and choose the  $q$  Byzantines as follows: for  $i = 0, 1, \dots, q - 1$ , among the remaining workers in  $\{U_1, U_2, \dots, U_K\} - A$  add a worker from the group  $B_{t,i \bmod K/r}$  to the adversarial set  $A$ . In this manner,

$$c^{(q),DETOX} = \begin{cases} q - \frac{K}{r}(r' - 1) & \text{if } q > \frac{K}{r}(r' - 1), \\ 0 & \text{otherwise.} \end{cases}$$

The results on this scenario are on Table 3.

Table 2: Distortion fraction of optimal attacks and comparison.

$q$	$\epsilon^{Aspis}$	$\epsilon^{Baseline}$	$\epsilon^{DETOX}$	$\epsilon^{ByzShield}$
2	0.004	0.133	0.2	0.04
3	0.022	0.2	0.2	0.12
4	0.062	0.267	0.4	0.2
5	0.132	0.333	0.4	0.32
6	0.242	0.4	0.6	0.48
7	0.4	0.467	0.6	0.56

2(a)  $(K, f, l, r) = (15, 455, 91, 3)$ .

$q$	$\epsilon^{Aspis}$	$\epsilon^{Baseline}$	$\epsilon^{DETOX}$	$\epsilon^{ByzShield}$
2	9.8e-04	0.083	0.125	0.031
3	0.005	0.125	0.125	0.063
4	0.014	0.167	0.25	0.125
5	0.03	0.208	0.25	0.188
6	0.054	0.25	0.375	0.281
7	0.09	0.292	0.375	0.375
8	0.138	0.333	0.5	0.5
9	0.202	0.375	0.5	0.5
10	0.282	0.417	0.625	0.531
11	0.38	0.458	0.625	0.625

2(b)  $(K, f, l, r) = (24, 2024, 253, 3)$ .

Table 3: Distortion fraction of weak attacks and comparison.

$q$	$\epsilon^{Aspis}$	$\epsilon^{Baseline}$	$\epsilon^{DETOX}$
2	0	0.13	0
3	0.0022	0.2	0
4	0.0088	0.267	0
5	0.022	0.333	0
6	0.044	0.4	0.2
7	0.0769	0.467	0.4

3(a)  $(K, f, l, r) = (15, 455, 91, 3)$ .

$q$	$\epsilon^{Aspis}$	$\epsilon^{Baseline}$	$\epsilon^{DETOX}$
2	0	0.083	0
3	4.9e-04	0.125	0
4	0.002	0.167	0
5	0.0049	0.208	0
6	0.0099	0.25	0
7	0.0173	0.292	0
8	0.0277	0.333	0
9	0.0415	0.375	0.125
10	0.0593	0.417	0.25
11	0.0815	0.458	0.375

3(b)  $(K, f, l, r) = (24, 2024, 253, 3)$ .

Note that for baseline schemes, there is no notion of “weak” or “optimal” attack with respect to the choice of the  $q$  Byzantines; hence we can choose any subset of them achieving  $\epsilon^{Baseline} = q/K$ .

## A.5 Experiment setup details

### A.5.1 Cluster setup

We used clusters of  $K = 15$  workers arranged in various setups within Amazon EC2. Initially, we used a PS of type `i3.16xlarge` and several workers of type `c5.4xlarge` to setup a distributed cluster. However, this requires training data to be transmitted from the PS to every single machine based on our current implementation; an alternative approach one can follow is to setup shared storage space accessible by all machines to store the training data. Also, some instances were automatically terminated by AWS per the AWS *spot instance* policy limitations<sup>2</sup>; this incurred some delays in resuming the experiments that were stopped. In order to facilitate our evaluation and to avoid these issues we decided to simulate the PS and the workers for the rest of the experiments on a single instance of type `x1.16xlarge`. We emphasize that the choice of the EC2 setup does not affect any of the numerical results in this paper since in all cases we used a single virtual machine image with the same dependencies.

### A.5.2 Dataset preprocessing and hyperparameter tuning

The images have been normalized using standard values of mean and standard deviation for the dataset. The value used for momentum (for gradient descent) was set to 0.9 and we trained for 16 epochs in all experiments. The number of epochs is precisely the invariant we maintain across all experiments, i.e., all schemes process the training data the same number of times. The batch size and the learning rate are chosen independently for each method; the number of iterations are adjusted

<sup>2</sup><https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-interruptions.html>




Figure 8: *ALIE* optimal attack and median-based defenses (CIFAR-10),  $K = 15$  with different random seeds.

accordingly to account for the number of epochs. We followed the advice of the authors of DETOX and chose  $b = 480$  for the DETOX and baseline schemes. For AspIs, we used  $b = 14560$  (32 samples per file) for the *ALIE* experiments and  $b = 1365$  (3 samples per file) for the remaining experiments except for the FoE optimal attack  $q = 4$  (cf. Figure 5) for which  $b = 14560$  performed better. In Table 4, a learning rate schedule is denoted by  $(x, y)$ ; this notation signifies the fact that we start with a rate equal to  $x$  and every  $z$  iterations we set the rate equal to  $x \times y^{t/z}$ , where  $t$  is the index of current iteration and  $z$  is set to be the number of iterations occurring between two consecutive checkpoints in which we store the model (points in the accuracy figures). We will also index the schemes in order of appearance in the corresponding figure’s legend. Experiments which appear in multiple figures are not repeated in Table 4 (we ran those training processes once). In order to pick the optimal hyperparameters for each scheme, we performed an extensive grid search involving different combinations of  $(x, y)$ . In particular, the values of  $x$  we tested are 0.1, 0.01 and 0.001 and for  $y$  we tried 1, 0.975, 0.95, 0.7 and 0.5. For each method, we ran 3 epochs for each such combination and chose the one which was giving the lowest value of average cross-entropy loss (principal criterion) and the highest value of top-1 accuracy (secondary criterion).


### A.5.3 Error bars

In order to examine whether the choice of the random seed affects the accuracy of the trained model we have performed the experiments for the *ALIE* distortion for two different seeds for the values  $q = 2, 4$  for every scheme; we used 428 and 50 as random seeds. In Figure 8a, for a given method we report the minimum accuracy, the maximum accuracy and their average for each evaluation point. We repeat the same process in Figures 9a and 10a when comparing with Bulyan and Multi-Krum, respectively. The corresponding experiments for  $q = 4$  are shown in the Figures 8b, 9b and 10b.

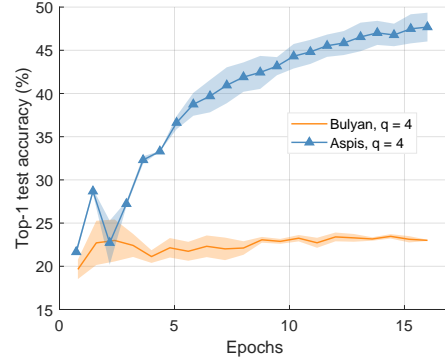
Given the fact that these experiments take a significant amount of time and that they are computationally expensive, we chose to perform this consistency check for a subset of our experiments. Nevertheless, these results indicate that prior schemes [13, 23, 11] are very sensitive to the choice of the random seed and demonstrate an unstable behavior in terms of convergence. In all of these cases the achieved value of accuracy at the end of the 16 epochs of training is small compared to AspIs. On the other hand, the accuracy results for AspIs are almost identical for both choices of the random seed.

### A.5.4 Computation and communication overhead

Our scheme provides robustness under powerful attacks and sophisticated distortion methods at the expense of increased computation and communication time. Note that each worker has to perform  $l$  forward/backward propagation computations and transmit  $l$  gradients per iteration. In related baseline [13, 4] and redundancy-based methods [8, 23] each worker is responsible for a single such computation. Experimentally, we have observed that AspIs needs up to  $5 \times$  overall training time compared to other schemes to complete the same number of training epochs. We emphasize that the training time incurred by each scheme depends on a wide range of parameters including the utilized defense, the batch size and the number of iterations and can vary significantly. We




(a)  $q = 2$  adversaries.




(b)  $q = 4$  adversaries.

Figure 9: *ALIE* optimal attack and Bulyan-based defenses (CIFAR-10),  $K = 15$  with different random seeds.



(a)  $q = 2$  adversaries.



(b)  $q = 4$  adversaries.

Figure 10: *ALIE* optimal attack and Multi-Krum-based defenses (CIFAR-10),  $K = 15$  with different random seeds.

believe that implementation-related improvements such as utilizing GPUs can alleviate some of the overhead. Communication-related algorithmic improvements are also worth exploring. Finally, our implementation natively supports resuming from a checkpoint (trained model) and hence, when new data becomes available we can only use that data to perform more epochs of training.

Table 4: Parameters used for training.

Figure	Schemes	Learning rate schedule
3a	1,2,5,6	(0.01, 0.7)
3a	3,4	(0.1, 0.95)
3b	1	(0.001, 0.95)
3c	1,2	(0.01, 0.7)
4a	1,2	(0.1, 0.7)
4a	3,4	(0.1, 0.95)
4a	5,6	(0.01, 0.7)
4b	1	(0.1, 0.7)
4c	1,2	(0.01, 0.975)
5	1,2	(0.1, 0.7)
5	3,4	(0.1, 0.95)
5	5,6	(0.01, 0.95)
6	1	(0.1, 0.95)
6	2	(0.01, 0.7)
7	1	(0.01, 0.7)
7	2	(0.1, 0.95)
7	3	(0.01, 0.7)

### A.5.5 Software

Our implementation of the Aspis algorithm used for the experiments builds on the ByzShield’s [19] PyTorch skeleton (obtained with the consent of the authors) and has been provided along with dependency information and instructions. The implementation of ByzShield is available at <https://github.com/kkonstantinidis/ByzShield> and uses the standard Github license. We utilized the NetworkX package [16] for the clique-finding; its license is 3-clause BSD. The CIFAR-10 dataset [20] comes with the MIT license; we have cited its technical report, as required. The data we are using does not contain personally identifiable information or offensive content.

### A.5.6 Broader impact

Our work provides improved robustness in large-scale training compared to state-of-the-art defenses and is useful for security-critical applications at the expense of increased computation and communication time. The proposed scheme can provide benefits in environments where safety is of utmost importance including healthcare, self-driving vehicles and fraudulent financial activity. Furthermore, it is directly applicable to a plethora of optimization algorithms without requiring strict assumptions. Finally, we have provided rigorous proofs that our training process is reliable and fail-safe and we have computed the exact fraction of distorted computations one can expect in different malicious scenarios. We do not anticipate any negative societal impacts stemming from our work.