High-Accuracy Low-Precision Training

Christopher R. Aberger[‡]   Christopher De Sa[†]   Megan Leszczynski[‡]   Alana Marzoev[†]
Kunle Olukotun[‡]   Christopher Ré[‡]   Jian Zhang[‡]

[†]Department of Computer Science, Cornell University, Ithaca, NY, USA

[‡]Department of Computer Science, Stanford University, Stanford, CA, USA

caberger@stanford.edu,cdesa@cs.cornell.edu, mleszczy@stanford.edu, mam655@cornell.edu

kunle@stanford.edu, chrismre@cs.stanford.edu, zjian@stanford.edu

Abstract

There is currently an arms race to design low-precision hardware accelerators capable of training machine learning models. This is because purpose-built, low-precision hardware accelerators can lower both the time and energy needed to complete a task. In contrast, traditional hardware architectures are over-provisioned, in terms of numerical precision, for machine learning tasks. Unfortunately, the statistical effects of low-precision computation *during training* are still not well understood. As a result, it is difficult to reach the statistical accuracies of traditional architectures on these new accelerators which have limited support for higher precision computation. This is due to a tradeoff with standard low-precision training algorithms: as the number of bits is decreased, noise that limits statistical accuracy is increased. In this paper we argue that one can reap the hardware benefits of low-precision accelerators while maintaining the statistical accuracies of traditional, higher-precision data formats. To do this we introduce a training algorithm called High-Accuracy Low-Precision (HALP). HALP is a low-precision stochastic gradient descent variant that uses entirely low-precision computation in its inner loop while infrequently recentering this computation with higher-precision computation done in an outer loop. HALP uses three techniques to reduce noise: (1) a known variance reduction method based on stochastic variance-reduced gradient (SVRG); (2) a novel bit centering technique that uses infrequent high-precision computation to reduce quantization noise; and (3) a novel dynamic bias adjustment technique to prevent overflow and underflow. On strongly convex problems, we show both theoretically and empirically that HALP converges at the same linear rate as full-precision SVRG. Inspired by these results, we show on two neural network applications (CNN and LSTM) that HALP can empirically compete with higher-precision training algorithms.

1. Introduction

Low-precision (less than 32-bit) hardware architectures are already here (Jouppi et al., 2017; Burger, 2017; Caulfield et al., 2017; Micikevicius, 2017) and future hardware roadmaps (Corporation, 2018; Micikevicius et al., 2017b) suggest that these lower-precision hardware architectures will soon be commonplace. Such architectures promise to improve the speed and efficiency of machine learning tasks when compared to traditional hardware architectures that only support more costly higher precision (32-bit) computation. Still, the effects of low-precision arithmetic on *training algorithms* are not yet well understood. As a result, if *only* low-precision computation is used throughout training, it is often very difficult to match the statistical accuracies of traditional, higher-precision hardware architectures. Even worse, these new, low-precision hardware accelerators often have limited capacity for higher-

precision computation, either from the host or due to the cost of reduced throughput. In this work, we argue that one can match the statistical accuracies of the traditional, high-precision hardware architectures by running the majority of the computation in lower-precision on an accelerator while infrequently (orders of magnitude less often) performing some full-precision computation.

It is well-accepted that the systems benefits of low-precision (LP) arithmetic commonly come with a cost. The round-off or *quantization error* that results from converting numbers into a low-precision representation introduces noise that can affect the convergence rate and accuracy of SGD. While quantization has been studied in communication (Tang et al., 2018), its effect when used on the updates in the *innermost loop* of training algorithms is much less well understood and more complex to understand due to error accumulation. Explicitly, we say that an algorithm has a *low-precision inner loop* if all the numeric computation in the inner loop is done in low precision. It is therefore unsurprising that algorithms where innermost loops are in low-precision tradeoff the number-of-bits against the resulting the statistical accuracy—the fewer bits used, the worse the solution will become. Theoretical upper bounds on the performance of low-precision SGD (De Sa et al., 2015; Li et al., 2017) and empirical observations of implemented low-precision algorithms (Courbariaux et al., 2014; Gupta et al., 2015; De Sa et al., 2017; Zhang et al., 2017) further confirm that current algorithms are limited by this precision-accuracy tradeoff.[1]

In this paper, we show that it is possible to get *high-accuracy solutions from an algorithm with a low-precision inner loop*, as long as the problem is sufficiently well-conditioned. We do this by introducing a technique called *bit centering*, in which each number is represented as the sum of a high-precision *offset* term, which is modified only infrequently, and a low-precision *delta* term, which is modified every iteration, as follows:

$$x = \underbrace{o_x}_{\text{high-precision offset}} + \underbrace{\delta_x}_{\text{low-precision delta.}}$$

The key insight behind bit centering is that floating point errors are relative to the magnitude of the number, and so a smaller low-precision delta means a smaller error due to the low-precision compute. That is, if we are able to choose o_x so that it is close to the value x will take on, then δ_x will be small. Therefore we can compute it using low-precision computations with much lower error than we would get if we had represented x directly in a low-precision format. Using bit centering, we develop an algorithm called HALP which transcends the accuracy limitations of ordinary low-precision SGD by reducing the three sources of errors that are present in standard low-precision SGD: gradient variance, quantization noise, and overflow/underflow. We address error from gradient variance using a known technique called SVRG, stochastic variance-reduced gradient (Johnson and Zhang, 2013). To address noise from quantization, we use bit centering, assigning our offsets o_x using the outer-loop full-gradient computation already present in the SVRG algorithm. Finally, we prevent overflow and underflow from occurring by dynamically rescaling the range of represented numbers for the delta term.

---

1. A simple way to avoid this and make an algorithm of arbitrary accuracy would be to increase the number of bits of precision as the algorithm converges. However, this is unsatisfying as it increases the cost of computation, and we want to be able to run on specialized low-precision accelerators that have a fixed bit width.

We prove that, for strongly convex problems, HALP can produce arbitrarily accurate solutions with the same linear asymptotic convergence rate as full-precision SVRG, while its inner loop uses entirely low-precision computation with a fixed number of bits.

We experimentally validate HALP on both convex and non-convex applications with varying levels of floating point precision. To accomplish this, we built three simulators, each serving a distinct experimental purpose: one in Julia capable of running arbitrary floating point precisions, one in PyTorch capable of running `float16` (10 bit mantissa, 5 bit exponent) on Nvidia cuDNN, and one in C++ that runs high-performance `bfloat16` (8 bit mantissa, 7 bit exponent) simulation on CPUs. On a linear regression model we show that 8-bit HALP matches the convergence trajectory of SVRG (as the theory suggests it should for convex applications). On multi-class logistic regression, we show that `float16` HALP beats 32-bit SGD and matches 32-bit SVRG in terms of both test accuracy and training loss. Further, on a CNN we show that `float16` HALP matches 32-bit training algorithms in terms of both test accuracy and training loss. Finally, on a LSTM we show that `bfloat16` beats 16-bit SVRG in terms of both test accuracy and training loss. Others (Jia et al., 2018) were inspired by earlier versions of this manuscript and went on to show that similar techniques can lead to state-of-the-art ImageNet training in only four minutes. We are excited to see where else these can techniques go they become more mainstream.

Our contributions and an outline are as follows:

- In Section 3, we introduce the technique of *bit centering*, and analyze its performance on a dot-product example. This simple examples illustrates the core mechanics of bit centering and its benefits in an important setting that is highly prevalent in machine learning applications.

- In Section 4, we describe how bit centering can be used in combination with SGD and SVRG to train a machine learning model, using linear regression as an example. We describe how overflow and underflow can be avoided by dynamically rescaling the bias of our floating point numbers. Using this, we introduce HALP, our *High-Accuracy Low-Precision* algorithm, for linear regression, and we show empirically that it outperforms both non-bit-centered algorithms and bit-centered algorithms that do not use re-scaling. Our theory exposes a novel relationship between the amount of precision that is needed and the condition number of the problem, and we validate this relationship empirically.

- In Section 5 we show how bit centering can be used to compute arbitrary Lipschitz-differentiable loss functions and their gradients, and we prove bounds on the error of these computations.

- In Section 6, we present a general version of HALP that can be applied to arbitrary loss functions using our methods from Section 5. We describe conditions under which we can prove that HALP converges down to solutions with error limited only by the high-precision offset format—and that HALP can produce *arbitrarily accurate* solutions if we assume the high-precision format has zero error.

- In Section 7, on a convex problem (multi-class logistic regression) and two non-convex problems (CNN and LSTM), we validate that an implementation of HALP beats the test accuracies and training loss of 16-bit SVRG. To run these experiments we built

two simulators with 16-bit precision: one in PyTorch using `float16` and one in C++ using `bfloat16`. Using these simulators we show that HALP not only can be useful on strongly convex problems (as the theory suggests), but can also be useful on popular non-convex problems such as deep learning applications.

2. Related work

Substantial previous research on low-precision machine learning has focused on evaluating and guaranteeing the effectiveness of low-precision training. Researchers have gathered empirical evidence for low-precision training in specific settings, although these results have typically not produced empirical support for 8-bit training (Savich and Moussa, 2011; Seide et al., 2014; Courbariaux et al., 2014; Gupta et al., 2015; Strom, 2015; Micikevicius et al., 2017b). Researchers have also proven bounds on the error that results from using low-precision computation on convex problems and non-convex matrix recovery problems (De Sa et al., 2015; Li et al., 2017). Recently, Zhang et al. (2017) has developed techniques called double sampling and optimal quantization which enable users to quantize the training dataset with provable guarantees on the accuracy for convex linear models. Similarly, De Sa et al. (2017), outlined how quantizing in different ways has different effects on accuracy and throughput when SGD is made low-precision. These works showed that low-precision training has many benefits but their accuracy degrades as precision is decreased. Our goal in this work is to show that this need not always be true.

While SGD is a very popular algorithm, the number of iterations it requires to achieve an objective gap of ε for a strongly convex problem is $O(1/\varepsilon)$. In comparison, ordinary gradient descent (GD) has an asymptotic rate of $O(\log(1/\varepsilon))$, which is known as a *linear rate*[2], and is asymptotically much faster than SGD. There has been work on modifications to SGD that preserve its high computational throughput while also recovering the linear rate of SGD (Roux et al., 2012; Shalev-Shwartz and Zhang, 2013). SVRG is one such method, which recovers the linear rate of gradient descent for strongly convex optimization, while still using stochastic iterations (Johnson and Zhang, 2013); recently it has been analyzed in the non-convex case as well and shown to be effective in some settings (Allen-Zhu and Hazan, 2016; Reddi et al., 2016). These *variance-reduced* methods are interesting because they preserve the simple hardware-efficient updates of SGD, while recovering the statistically-efficient linear convergence rate of the more expensive gradient descent algorithm. While we are not the first to present theoretical results combining low-precision with SVRG (Alistarh et al. (2017) previously studied using low-precision numbers for communication among workers in parallel SGD and SVRG), to the best of our knowledge we are the first to present empirical results on low-precision SVRG and to propose the additional bit centering technique.

Exciting results in the deep learning community suggest that neural network training can be done effectively by mixing 32-bit floating point arithmetic with lower precision formats (Nvidia, 2018; Corporation, 2018; Micikevicius et al., 2017a). Much of this work focuses on 16-bit training while performing 32-bit accumulations and storing a master copies of the networks weights in 32-bit precision. Researchers have also recently released mixed-precision training techniques that are capable performing 8-bit floating point computations for much

---

2. This is called a linear rate because the iterations required is linear in the number of significant figures of output precision desired.

of training (Wang et al., 2018). At the farthest extreme, there has been work on training networks with binarized weights from scratch, although these networks are still trained with full precision updates (Hubara et al., 2016; Rastegari et al., 2016). Empirically HALP is similar in spirit to much of this work, but theoretically it stands on its own by providing strong convergence guarantees. Because of this we are primarily focused on the strongly convex cases that we can reason about, of which deep learning is not. Still, we believe (and provide evidence in Section 7) that our methods can complement the exciting deep learning techniques that exist today. In fact, there is already empirical work, motivated by earlier versions of this manuscript, which has shown that similar techniques can lead to state-of-the-art results on ImageNet training (Jia et al., 2018).

3. Bit Centering

In this section, we introduce the *bit centering* technique and describe how it can reduce errors due to quantization. We illustrate the benefits of bit centering by analyzing its error on a simple dot-product example.

The technical inspiration for this paper is the recent resurgence of hardware accelerators for machine learning and other applications.[3] In such a scenario, a hardware accelerator designed for machine learning runs inside a host device with a larger memory and general purpose processor (e.g., a CPU). The accelerator typically has many more computational units of lower precision and smaller memory compared with the general purpose host. This motivates the development of algorithms that use low-precision arithmetic, and much past work has been devoted to studying such algorithms. However, standard analysis of low-precision algorithms often *ignores the host*, and analyzes a fully-low-precision or *accelerator-only* model of computation. This motivates us to ask whether performance can be improved by also using the host, which we call a *host-accelerator* model of computation (Figure 1). In this section, we describe a principled way to do this. We develop a novel technique called *bit centering* which greatly improves the accuracy of low-precision computation on the accelerator by leveraging infrequent full-precision computation on the host (orders of magnitude less frequent than the computation on the accelerator).

As we will see, the conventional wisdom is that training algorithms using low-precision iterations are limited to low accuracy. Our goal is to formalize situations in which adding a modest amount of host memory and computation allows us to achieve dramatically better error with minimal overhead.

The core idea. Bit centering is a technique in which, instead of representing a number x as a single floating-point number, we represent x as the *sum* of a pre-computed full-precision *offset* and a low-precision *delta*:

$$x = o_x + \delta_x.$$

For example, we could store o_x as a 64-bit double-precision float and store δ_x as a 16-bit half-precision float. The benefit of bit centering, as compared with simply storing the number x itself in low-precision, comes from the fact that floating point computation has *relative*

---

3. Christopher De Sa, Christopher Aberger, Kunle Olukuton, and Christopher Ré have a financial interest in SambaNova Systems, which produces machine learning and big data platform including hardware accelerators and an accelerated software stack.
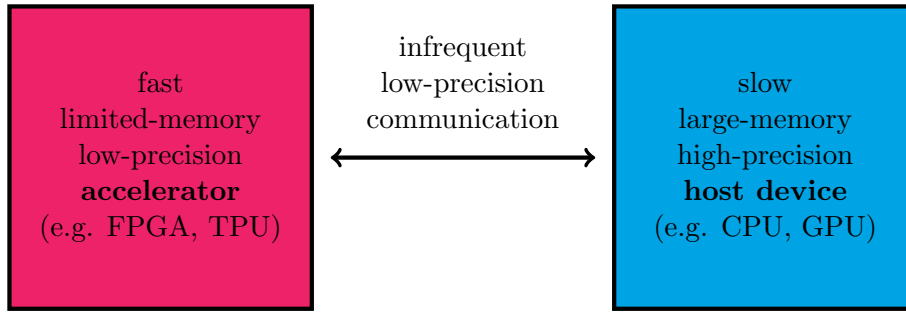
Figure 1: An illustration of the host-accelerator model of computation. In this model, a fast machine learning accelerator device capable of only low-precision computation is connected to a slower standard computation device that can perform high-precision arithmetic via an infrequently used communication channel.

error. That is, if we do some computation in low precision that is intended to approximate some number $x \in \mathbb{R}$, and the result of that computation is a floating point number $\tilde{x}$, then

$$|\tilde{x} - x| \leq \varepsilon |x|,$$

where ε, which is called the *machine epsilon*, is some fixed-but-small number that is a function of the number of bits used. In comparison, if we use bit centering to compute x, and we compute only the delta $\tilde{\delta}_x$ in low-precision (assuming negligible error in computing o_x), the error will look like

$$|\tilde{x} - x| = \left| \tilde{\delta}_x - \delta_x \right| \leq \varepsilon |\delta_x| = \varepsilon |x - o_x|.$$

So if we can choose o_x so that it will be close to x (equivalently, so that δ_x will be small), then bit centering can produce a much more accurate estimate of x using the same low-precision format. In this section, we will introduce our notation for bit centering and illustrate its benefits by analyzing a simple dot-product example, supposing that an offset o_x is already given. We will describe how to select the offset in the following sections.

3.1 Notation

As much as possible, we will follow the notation of Trefethen and Bau III (1997) and their error model. Namely, given a number x and a floating point number system with machine error $\varepsilon_{\text{machine}}$, then if $\mathsf{fl}(x)$ is the floating point representation of x, there will exist an ε such that $|\varepsilon| \leq \varepsilon_{\text{machine}}$ and:

$$\mathsf{fl}(x) = (1 + \varepsilon)x.$$

Moreover, if $*$ denotes any basic operation in scalar arithmetic (and also fused-multiply and add on two floating point representations), and $\circledast$ denotes its floating point analogue, then for some ε such that $|\varepsilon| \leq \varepsilon_{\text{machine}}$

$$x \circledast y = (1 + \varepsilon)(x * y).$$

Equivalently,

$$|x \circledast y - x * y| = \varepsilon_{\text{machine}}.$$

6

This is sometimes called the *fundamental axiom of floating point arithmetic.*[4]

Note that this model does not explicitly consider overflow or underflow, but it is an accurate model of floating point arithmetic in situations when overflow and underflow cannot occur (e.g. if a very large number of exponent bits are used). The point at which a number underflows is the *underflow threshold*, the smallest representable number, which we denote η_{machine}. In situations where a number could underflow (but not overflow), we represent the result of a floating point computation with a modified version of the fundamental axiom of floating point arithmetic: for some ε with $|\varepsilon| \leq \varepsilon_{\text{machine}}$ and some η with $|\eta| \leq \eta_{\text{machine}}$,

$$x \circledast y = (1 + \varepsilon)(x * y) + \eta.$$

Overflow, on the other hand, always results in an infinite value, and so its error cannot be bounded with a formula such as this. We let M_{machine} denote the maximum representable value of the floating point representation, above which overflow could occur. In our analysis that follows, when we do consider overflow, we will ensure that it is avoided. For the first part of this paper, we will consider an idealized version of floating point in which overflow and underflow do not occur, and then we will come back to overflow and underflow later in Section 4.3.

In this paper, we will deal with two floating point number systems: (1) a **low precision** number system with machine precision $\varepsilon_{\text{machine-lo}}$ and (2) a **high precision** number system with machine precision $\varepsilon_{\text{machine-hi}}$. We let the low-precision representation of a number x be denoted $\mathsf{fl}_{\text{lo}}(x)$ and the high-precision representation be denoted $\mathsf{fl}_{\text{hi}}(x)$. In the analysis that follows, we will use ε_{lo} to denote a value that satisfies $|\varepsilon_{\text{lo}}| \leq \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2)$ and ε_{hi} to denote a value that satisfies $|\varepsilon_{\text{hi}}| \leq \varepsilon_{\text{machine-hi}} + O(\varepsilon_{\text{machine-hi}}^2)$. This is a straightforward generalization of the $1 + \varepsilon$ model of floating point error described above to the case of mixed-precision algorithms. This also allows us to simplify the expression of factors that are quadratic in the machine epsilon using the big-O notation, which is typical in rounding error analysis. We use analogous notation for bounding overflow and underflow: we use $\eta_{\text{machine-lo}}$ to denote the underflow limit of the low-precision format, and let η_{lo} generally denote a value that satisfies $|\eta_{\text{lo}}| \leq \eta_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}} \cdot \eta_{\text{machine-lo}} + \eta_{\text{machine-lo}}^2)$, and similarly for the high-precision format. We also define the precision ratio k as

$$k = \frac{\varepsilon_{\text{machine-lo}}}{\varepsilon_{\text{machine-hi}}}.$$

To more easily denote arithmetic operations with floating point numbers of multiple precisions, we use the following conventions. We follow the notation of Trefethen and Bau III (1997) in using operators with a circle around them to denote a floating point operation, and we use an underset letter H or L to denote high-precision or low-precision respectively. For example,

$$\mathsf{fl}_{\text{hi}}(x \underset{\text{L}}{\odot} y) \underset{\text{H}}{\oplus} z$$

denotes the result of multiplying x and y in the low-precision format, converting that into high precision, and then adding it to z using high-precision computation. We extend this

---

4. Note that in some sources, δ is used for these sorts of expressions, rather than ϵ. Here we use ϵ both to follow Trefethen and Bau III (1997) and to keep the letter δ free for use as the delta in bit centering.

notation to allow floating point operators to also operate on vectors and matrices: for example, if A is a matrix, and x and y are vectors, then

$$y^T \underset{\text{H}}{\odot} \text{fl}_{\text{hi}}(A \underset{\text{L}}{\odot} x)$$

denotes the result of multiplying matrix A by vector x using low-precision arithmetic, converting the (vector) result into the high-precision format, and then taking its dot product in high-precision with the vector y. We also extend the notation for computing functions and sums. For example,

$$\underset{\text{L}}{\exp}(x)$$

denotes the result of computing the exponential function of x in low precision and

$$\sum_{i=1}^{N}{}_{\text{H}} x_i$$

denotes the result of summing up the x_i in high-precision arithmetic. Ordinary numeric operations without circles or subscripts denote those operations performed in exact arithmetic, as usual. Unless otherwise indicated, any norm $\|\cdot\|$ indicates the Euclidean ℓ_2 norm.

3.2 Case Study: Dot Product

In this subsection, we will describe bit centering and compare it with other quantization strategies when used to compute a simple dot product, $x^T y$, for x and y non-negative vectors of real numbers. We suppose that x and y are representable exactly in both the low-precision and high-precision formats, are available on both the host and the accelerator, and that we want to produce an output in the low-precision floating-point format. The quantization strategies are:

- **High-precision.** Perform the whole computation in high precision. This is the null quantization strategy, as it cannot use any of the low-precision capabilities of the accelerator. (It is also in a sense cheating, since it does not produce low-precision output.)

- **Quantize last.** Perform the computation in high-precision floating poin, and then quantize the output to low-precision floating point at the end. When this strategy is used to quantize a training algorithm, it is sometimes called *gradient quantization*, since the full gradient is computed using high-precision arithmetic, and then only the computed gradient itself is quantized.

- **Quantize first.** Quantize all the inputs to the computation into low-precision floating point, and then perform the computation in low-precision. This quantizes all operations, but can cause accuracy to suffer.

- **Bit centering.** Compute using bit-centered numbers with a pre-computed-on-the-host high-precision offset and a computed-on-the-accelerator low-precision floating-point delta, and then quantize the output into low-precision on the accelerator.

Strategy	# LP Ops	# HP Ops	Result
High-precision	—	n multiplies $n-1$ adds	$(x^T y)(1 + \varepsilon_{\mathrm{hi}} \cdot n)$
Quantize last	—	n multiplies $n-1$ adds	$(x^T y)(1 + \varepsilon_{\mathrm{hi}} \cdot (n + k))$
Quantize first	n multiplies $n-1$ adds	—	$(x^T y)(1 + \varepsilon_{\mathrm{hi}} \cdot kn)$
Bit centering	n multiplies n adds	n multiplies $n-1$ adds precomputed on host	$(x^T y)(1 + \varepsilon_{\mathrm{hi}} \cdot (n + k))$ $+ (\delta_x^T y) \cdot \varepsilon_{\mathrm{hi}} \cdot kn$ $+ (x^T y) \cdot \varepsilon_{\mathrm{hi}} \cdot k + (\delta_x^T y) \cdot (n + 3k)$

Table 1: Results of performing a dot product using various quantization strategies. We refer to eq. (2) for discussions on the count of Ops for bit centering.

In the derivations that follow, we will use the well-known fact (Bindel, 2012) that for a nonnegative floating-point dot product with machine error $\varepsilon_{\mathrm{machine}}$, the result of the computation will be, for some $|\varepsilon| \leq \varepsilon_{\mathrm{machine}}$,

$$x^T \odot y = x^T y + |x|^T |y| \, n\varepsilon$$

where $|x|$ denotes the absolute value of x applies entrywise, and where we are ignoring terms proportional to ε^2. (Note that this is for using naive summation; if we used pairwise summation the n would be replaced by a $\log n$ factor.) For non-negative x and y, this is equivalent to

$$x^T \odot y = x^T y \cdot (1 + n\varepsilon), \tag{1}$$

a fact we use in the subsequent derivation. Our results in this section are summarized in Table 1.

High-precision The result of computing a dot product in high precision follows immediately from eq. (1). We will have

$$\mathsf{fl}_{\mathrm{hi}}(x)^T \underset{\mathrm{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(y) = x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n).$$

Quantize last We compute the dot product first in high precision, producing an intermediate value which will satisfy

$$\mathsf{fl}_{\mathrm{hi}}(x)^T \underset{\mathrm{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(y) = x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n).$$

If we then quantize that high-precision intermediate value to low precision, we will get

$$
\begin{aligned}
\mathsf{fl}_{\mathrm{lo}}(\mathsf{fl}_{\mathrm{hi}}(x)^T \underset{\mathrm{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(y)) &= (x^T y)(1 + \varepsilon_{\mathrm{hi}} \cdot n)(1 + \varepsilon_{\mathrm{lo}}) \\
&= (x^T y)(1 + \varepsilon_{\mathrm{hi}} \cdot n + \varepsilon_{\mathrm{lo}}) \\
&= (x^T y)\left(1 + \varepsilon_{\mathrm{hi}} \cdot \left(n + \frac{\varepsilon_{\mathrm{machine\text{-}lo}}}{\varepsilon_{\mathrm{machine\text{-}hi}}}\right)\right) \\
&= (x^T y)\left(1 + \varepsilon_{\mathrm{hi}} \cdot (n + k)\right).
\end{aligned}
$$

Note again that we drop terms that are $O(\varepsilon^2)$, as is standard in Trefethen and Bau III (1997), and as is justified by our definitions of $\varepsilon_{\mathrm{lo}}$ and $\varepsilon_{\mathrm{hi}}$ in Section 3.1. While this is a small relative error (as long as $\varepsilon_{\mathrm{machine\text{-}hi}}$ is small enough that $\varepsilon_{\mathrm{machine\text{-}hi}} \cdot (n + k) \ll 1$), all of the computation we did here was in high precision.

Quantize first   If we first quantize to low precision and then do the multiplication, then

$$
\mathsf{fl}_{\mathrm{lo}}(x)^T \underset{\mathrm{L}}{\odot} \mathsf{fl}_{\mathrm{lo}}(y) = (x^T y)(1 + \varepsilon_{\mathrm{lo}} \cdot n) = (x^T y)(1 + \varepsilon_{\mathrm{hi}} \cdot kn)
$$

That is, we suffer the extra error *multiplicatively*, rather than additively, as we did in the quantize-last case.

Hybrid method: bit centering   In bit centering, some of the computation is done in high-precision and some is done in low-precision. Here, we suppose that $x = o_x + \delta_x$ for some high-precision precomputed vector o_x and some low-precision vector δ_x. We also suppose that y is stored in high-precision on the host. The dot product of these two terms is

$$
x^T y = (o_x + \delta_x)^T y = o_x^T y + \delta_x^T y.
$$

The first dot product term here, $o_x^T y$, depends only on the offset of x, not on the delta, so $o_x^T y$ can be precomputed on the host in full-precision, and then sent to the accelerator. While doing this, we can also precompute a low-precision version of y. The remaining terms can be computed in low-precision on the accelerator. Explicitly, this results in the following algorithm, Algorithm 1. In the remaining of this section, we will demonstrate that the error of this bit centered dot product interpolates those of quantize first and quantize last. At the same time, it can use less overall computation than quantize last if we evaluate the dot products with y for multiple different values of δ_x, with the same offset o_x.

10

Algorithm 1 Simple Bit Centered Dot Product

$\triangleright$ *high-precision precomputation on host*

given: high-precision offset o_x and value y

compute in high precision: $o_{x^T y} \leftarrow o_x^T \underset{\text{H}}{\odot} y$

send to accelerator: $\mathsf{fl}_{\mathrm{lo}}(y)$ and $\mathsf{fl}_{\mathrm{lo}}(o_{x^T y})$

$\triangleright$ *low-precision computation on accelerator*

given: low-precision delta δ_x

receive from host: $\mathsf{fl}_{\mathrm{lo}}(y)$ and $\mathsf{fl}_{\mathrm{lo}}(o_{x^T y})$

compute in low precision: $\delta_{x^T y} \leftarrow \delta_x^T \underset{\text{L}}{\odot} \mathsf{fl}_{\mathrm{lo}}(y)$

output $\delta_{x^T y} \underset{\text{L}}{\oplus} \mathsf{fl}_{\mathrm{lo}}(o_{x^T y})$

Here, part of the dot product is computed in high precision (the offset part), and part is computed in low precision (the delta part). We expect that if the delta is small, this algorithm will have low error. Let's upper bound the error of this algorithm, assuming for simplicity that o_x, δ_x, and y are nonnegative.[5] First, for the offset, we immediately get

$$\mathsf{fl}_{\mathrm{lo}}(\mathsf{fl}_{\mathrm{hi}}(o_x)^T \underset{\text{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(y)) = \mathsf{fl}_{\mathrm{lo}}(o_x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n))$$
$$= o_x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n) \cdot (1 + \varepsilon_{\mathrm{lo}})$$
$$= o_x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n + \varepsilon_{\mathrm{lo}}),$$

where following convention we ignore $O(\varepsilon^2)$ terms. Second, for the delta, we have

$$\delta_x^T \underset{\text{L}}{\odot} \mathsf{fl}_{\mathrm{lo}}(y) = \delta_x^T \mathsf{fl}_{\mathrm{lo}}(y) \cdot (1 + \varepsilon_{\mathrm{lo}} \cdot n)$$
$$= \delta_x^T (y \cdot (1 + \varepsilon_{\mathrm{lo}})) \cdot (1 + \varepsilon_{\mathrm{lo}} \cdot n)$$
$$= \delta_x^T y \cdot (1 + \varepsilon_{\mathrm{lo}} \cdot (n + 1)).$$

Adding these together,

$$\mathsf{fl}_{\mathrm{lo}}(\mathsf{fl}_{\mathrm{hi}}(o_x)^T \underset{\text{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(y)) \underset{\text{L}}{\oplus} \left(\delta_x^T \underset{\text{L}}{\odot} \mathsf{fl}_{\mathrm{lo}}(y) \right)$$
$$= \left(\mathsf{fl}_{\mathrm{lo}}(\mathsf{fl}_{\mathrm{hi}}(o_x)^T \underset{\text{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(y)) + \left(\delta_x^T \underset{\text{L}}{\odot} \mathsf{fl}_{\mathrm{lo}}(y) \right) \right) \cdot (1 + \varepsilon_{\mathrm{lo}})$$
$$= \left(o_x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n + \varepsilon_{\mathrm{lo}}) + \delta_x^T y \cdot (1 + \varepsilon_{\mathrm{lo}} \cdot (n + 1)) \right) \cdot (1 + \varepsilon_{\mathrm{lo}}) \quad (2)$$
$$= \left((x - \delta_x)^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n + \varepsilon_{\mathrm{lo}}) + \delta_x^T y \cdot (1 + \varepsilon_{\mathrm{lo}} \cdot (n + 1)) \right) \cdot (1 + \varepsilon_{\mathrm{lo}})$$
$$= \left(x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n + \varepsilon_{\mathrm{lo}}) + \delta_x^T y \cdot (\varepsilon_{\mathrm{hi}} \cdot n + \varepsilon_{\mathrm{lo}} \cdot (n + 2)) \right) \cdot (1 + \varepsilon_{\mathrm{lo}})$$
$$= x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot n + 2\varepsilon_{\mathrm{lo}}) + \delta_x^T y \cdot (\varepsilon_{\mathrm{hi}} \cdot n + \varepsilon_{\mathrm{lo}} \cdot (n + 3))$$
$$= x^T y \cdot (1 + \varepsilon_{\mathrm{hi}} \cdot (n + k)) + \delta_x^T y \cdot \varepsilon_{\mathrm{hi}} \cdot kn + x^T y \cdot \varepsilon_{\mathrm{hi}} \cdot k + \delta_x^T y \cdot (n + 3k)$$

---

5. The non-negativity assumption here is not necessary to bound the error, but just used to simplify the illustration. Later, in Section 5 we will provide a full analysis of general loss functions that does not make any assumptions of non-negativity, and this analysis will apply to the dot product as a special case.

This formula interpolates between the two formulae for the quantize-first and quantize-last approaches. We've written this so that the first term is identical to the result of the quantize last strategy, the second term is the quantize first error, and the remaining terms measure a small extra error from doing the extra addition to add the offset and delta components together. Since the quantize last term is the best we could hope for, we can view the second and third terms as additional error. When δ_x is small enough this approach improves error—as we suspected.

A difference of this approach from quantize first is that we pre-compute $\mathsf{fl}_{\mathrm{lo}}(o_x^T \underset{\mathrm{H}}{\odot} y)$ on the host and then load it into memory. In terms of accelerator compute, it costs at worst a $2/n$ additional relative computation time over quantize first (since we need to do an extra add and an extra load from the host, compared with quantize first). This may actually save compute overall if we are evaluating dot products with y for multiple different values of δ_x, but the same offset o_x. In the next section we will show how we can arrange such a situation in the training of linear models using bit centering.

4. Learning Using Bit Centering: Linear Regression

In this section, we describe how we can train a machine learning model using bit centering, by focusing first on the simple case of linear regression; we will extend from linear models to general loss function in Section 5. A linear regression problem has a loss of the form

$$f(w) = \frac{1}{n}\sum_{i=1}^{n} f_i(w) = \frac{1}{2n}\sum_{i=1}^{n} \left(x_i^T w - y_i\right)^2,$$

where n is the number of training examples, $w \in \mathbb{R}^d$ is the model, each $x_i \in \mathbb{R}^d$ is a training example, and $y_i \in \mathbb{R}$ is the corresponding training label. We choose to introduce our algorithms using linear regression because it results in a particularly simple inner loop for SVRG, which has update step

$$\begin{aligned} w_{k,t} &= w_{k,t-1} - \alpha \left(\nabla f_i(w_{k,t-1}) - \nabla f_i(\tilde{w}_k) + \nabla f(\tilde{w}_k)\right) \\ &= w_{k,t-1} - \alpha \left((x_i^T w_{k,t-1} - y_i) \cdot x_i - (x_i^T \tilde{w}_k - y_i) \cdot x_i + \nabla f(\tilde{w}_k)\right). \end{aligned}$$

Since this update step can be computed with just dot products and basic vector arithmetic, we can bound its error using the analysis of dot products we developed in the previous section.

4.1 Learning Using Stochastic Gradient Descent

Stochastic gradient descent (SGD) is one of the most basic and fundamental training algorithms used in machine learning, and many more sophisticated algorithms are built on top of SGD. In this subsection, we will follow this trend by describing low-precision and bit-centered variants of SGD. Low-precision SGD (Algorithm 2) has been explored many times in prior work (De Sa et al., 2015; Li et al., 2017). This is just ordinary SGD computed entirely using the low-precision format, and corresponds to the "quantize first" strategy described in Section 3.2.

Algorithm 2 LP-SGD: Low-precision SGD on linear regression

1: **given:** N low-precision training examples x_i and labels y_i, number of steps T, low-precision step size α, initial low-precision iterate w_0.
2: **for** $t = 1$ **to** T **do**
3:     **sample** i uniformly from $\{1, \ldots, N\}$
4:     $w_t \leftarrow w_{t-1} \ominus_{\mathrm{L}} \alpha \odot_{\mathrm{L}} \left(\left(x_i^T \odot_{\mathrm{L}} w_{t-1} \ominus_{\mathrm{L}} y_i \right) \odot_{\mathrm{L}} x_i \right)$
5: **end for**
6: **return** w_T

Algorithm 3 BC-SGD: Bit-centered SGD for linear regression

---

given: N low-precision training examples x_i and labels y_i, number of epochs K, epoch length T, low-precision step size α, initial iterate w_0.
$o_{w,1} \leftarrow \mathsf{fl}_{\mathrm{hi}}(w_0)$
for $k = 1$ **to** K **do**

> ▷ *high-precision precomputation on host*
> **for** $i = 1$ **to** N **do** {This loop is fully parallelizable.}
>
> $\quad g_{k,i} \leftarrow \left(\mathsf{fl}_{\mathrm{hi}}(x_i)^T \odot_{\mathrm{H}} o_{t-1} \ominus_{\mathrm{H}} \mathsf{fl}_{\mathrm{hi}}(y_i) \right) \odot_{\mathrm{H}} \mathsf{fl}_{\mathrm{hi}}(x_i)$
> $\quad h_{k,i} \leftarrow \mathsf{fl}_{\mathrm{lo}}(g_{k,i})$
> **end for**

    $\delta_{w,k,0} \leftarrow \mathsf{fl}_{\mathrm{lo}}(0)$

> ▷ *low-precision computation on accelerator*
> **for** $t = 1$ **to** T **do**
>   **sample** i uniformly from $\{1, \ldots, N\}$
> $\delta_{w,k,t} \leftarrow \delta_{w,k,t-1} \ominus_{\mathrm{L}} \alpha \odot_{\mathrm{L}} \left(\left(x_i^T \odot_{\mathrm{L}} \delta_{w,k,t-1} \right) \odot_{\mathrm{L}} x_i \oplus_{\mathrm{L}} h_{k,i} \right)$
> **end for**

    $o_{w,k+1} \leftarrow o_{w,k} \oplus_{\mathrm{H}} \mathsf{fl}_{\mathrm{hi}}(\delta_{w,k,T})$
end for
return $o_{w,K+1}$

A bit-centered version of SGD is described in Algorithm 3. BC-SGD computes the gradient samples used for SGD using bit centering. Each outer loop iteration is composed of a phase of high-precision parallelizable precomputation (highlighted in blue) followed by an inner loop which runs a bit-centered version of the SGD update loop using only low-precision operations (highlighted in red). It is straightforward to verify that in the absence of any numerical error (i.e. if both the "low-precision" and "high-precision" formats are exact arithmetic), BC-SGD and LP-SGD are both identical to ordinary SGD. In other words, we should expect these algorithms to differ only in their interaction with quantization error. From our results in Section 3.2, we should expect a bit-centered computation to have less

numerical error than one based on the "quantize-first" low precision strategy, so we should expect BC-SGD to perform better than LP-SGD.

Empirical validation. To validate this intuition, we ran an experiment comparing the convergence of these two algorithms. To do this, we generated a synthetic linear regression dataset with $n = 1024$ examples in dimension $d = 256$. We generated training examples by first sampling a weight vector w_{gen} at random and then sampling examples according to

$$w_{\text{gen}} \sim \mathcal{N}(0, I_d), \qquad x_i \sim \mathcal{N}\left(0, \frac{1}{\sqrt{d}}I_d\right), \qquad y_i \sim \mathcal{N}\left(x_i^T w_{\text{gen}}, \frac{1}{100}\right),$$

where $\mathcal{N}(\mu, \Sigma)$ denotes the normal distribution with mean μ and covariance matrix Σ. We evaluated both SGD and BC-SGD on this synthetic training set for three different low-precision formats: IEEE 32-bit floats (8-bit exponent and 23-bit mantissa), IEEE 16-bit floats (5-bit exponent and 10-bit mantissa), and the *bfloat16* 16-bit float format (8-bit exponent and 7-bit mantissa). We used 64-bit floats for the high-precision representation, and randomized rounding was used for all the low-precision formats. We ran each algorithm for $K = 100$ epochs of $T = 24 \times 1024$ iterations each. For each algorithm-precision pair, we initialized with $w_0 = 0$ and ran for step sizes α in $\{1.0, 0.3, 0.1, 0.03, 0.01, 0.003, 0.001, 0.0003, 0.0001\}$ and plotted the convergence trajectory for the step size α that resulted in the lowest value of the ℓ_2 norm of the full gradient $\|\nabla f(w)\|$ after 100 epochs.

Our results are presented in Figure 2. These results validate our intuition that bit-centered versions of SGD should be more accurate. A clear example of the difference produced by bit centering can be seen in the trajectories for bfloat formats: while the LP-SGD bfloat trajectory (blue circles) is limited by numerical precision and has the worst eventual gradient norm of all formats, the BC-SGD bfloat trajectory (red circles) eventually converges down to an error that is nearly that achieved by the 32-bit float format.

While the results in Figure 2 indicate that bit centering can decrease the numerical error caused by using low-precision numbers in SGD, even BC-SGD still converges to a limited level of accuracy—sometimes called a *noise ball*. This is because there is still a substantial source of noise in the algorithm: the variance of the gradient samples. While this variance can be reduced by lowering the step size, we can observe empirically that a lower step size also increases the negative effect from using low-precision arithmetic. This is illustrated in Figure 3, which compares the loss gradient after 100 epochs in the same experiment as Figure 2 for different values of the learning rate α. From the figure, we can see that additional error caused by low-precision becomes more of an issue as the learning rate is decreased: while all the algorithms behave roughly the same when $\alpha = 1$, as α is decreased most of the lower-precision (16-bit) methods start to perform worse compared with the 32-bit SGD baseline, and they only match its performance again when the step size is too small for even the 32-bit baseline to make much progress. As a result of this effect, lowering the step size is not a panacea for lowering the error in low-precision SGD as it is in SGD when numerical issues are not considered. This motivates us to explore other methods for improving the accuracy of bit-centered learning algorithms.

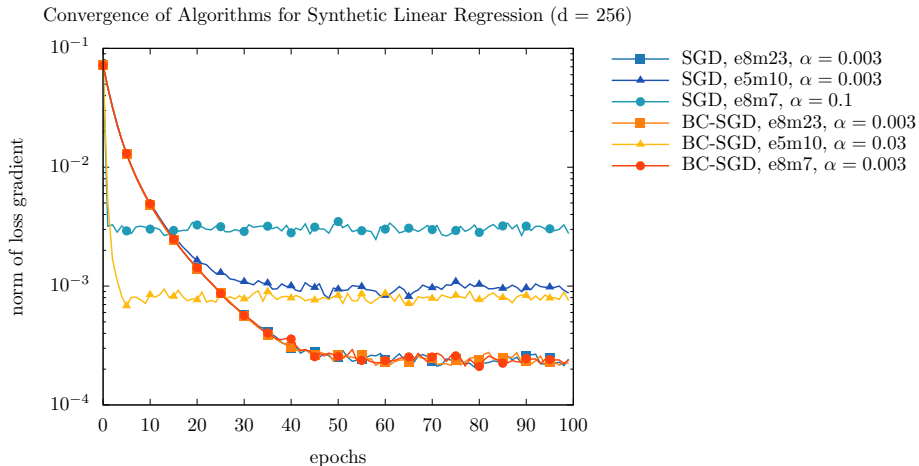Convergence of Algorithms for Synthetic Linear Regression (d = 256)

Figure 2: A comparison of the convergence of SGD and bit-centered SGD (BC-SGD) for a variety of low-precision numerical formats on a synthetic linear regression dataset. The notation e*x*m*y* indicates that *x* bits are used for the exponent and *y* bits are used for the mantissa.

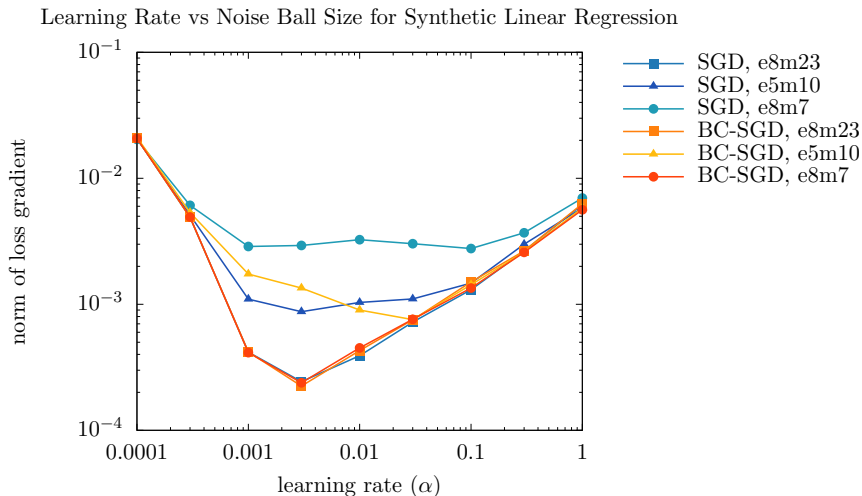Learning Rate vs Noise Ball Size for Synthetic Linear Regression

Figure 3: A comparison of the convergence of SGD and bit-centered SGD (BC-SGD) for a variety of numerical formats on a synthetic linear regression dataset.

4.2 Learning with Variance Reduction: SVRG

Since we cannot reduce error in the iterations of LP-SGD and BC-SGD by making the step size arbitrarily small, we will look at other ways to reduce this error. One way to do this is to use *variance reduction*, which lowers the error of the updates of SGD by lowering the variance of the gradient samples it uses. For the remainder of this section, we will focus on a variance reduction technique called *stochastic variance-reduced gradient* (SVRG) (Johnson and Zhang, 2013). Compared with standard stochastic gradient descent, SVRG is able to

---

Algorithm 4 LP-SVRG: Low-precision SVRG on linear regression

1: **given:** N low-precision training examples x_i and labels y_i, number of epochs K, epoch length T, low-precision step size α, initial low-precision iterate $\tilde{w}_1$.

2: **for** $k = 1$ **to** K **do**

3: $\quad \tilde{g}_k \underset{\text{H}}{\leftarrow} \underset{\text{H}}{\nabla} f(\mathsf{fl}_{\text{hi}}(\tilde{w}_k)) = \left(\sum_{i=1}^{N}{}_{\text{H}} \left(\mathsf{fl}_{\text{hi}}(x_i)^T \underset{\text{H}}{\odot} \mathsf{fl}_{\text{hi}}(\tilde{w}_k) \underset{\text{H}}{\ominus} \mathsf{fl}_{\text{hi}}(y_i) \right) \underset{\text{H}}{\odot} \mathsf{fl}_{\text{hi}}(x_i) \right) \underset{\text{H}}{\oslash} N$

4: $\quad \tilde{h}_k \leftarrow \mathsf{fl}_{\text{lo}}(\tilde{g}_k)$

5: $\quad w_{k,0} \leftarrow \tilde{w}_k$

6: $\quad$ **for** $t = 1$ **to** T **do**

7: $\qquad$ **sample** i uniformly from $\{1, \ldots, N\}$

8: $\qquad w_{k,t} \leftarrow w_{k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\left(x_i^T \underset{\text{L}}{\odot} w_{k,t-1} \right) \underset{\text{L}}{\odot} x_i \underset{\text{L}}{\ominus} \left(x_i^T \underset{\text{L}}{\odot} \tilde{w}_k \right) \underset{\text{L}}{\odot} x_i \underset{\text{L}}{\oplus} \tilde{h}_k \right)$

9: $\quad$ **end for**

10: $\quad$ **option I:** set $\tilde{w}_{k+1} \leftarrow w_{k,T}$

11: $\quad$ **option II: sample** t uniformly from $\{0, \ldots, T-1\}$, then set $\tilde{w}_{k+1} \leftarrow w_{k,t}$

12: **end for**

13: **return** $\tilde{w}_{K+1}$

---

converge at a linear rate because it periodically uses full gradients $\tilde{g}_k$ to reduce the variance of its stochastic gradient estimators.

As a warmup, we derive low-precision SVRG (LP-SVRG), which combines low-precision computation with variance reduction (but without bit centering). This will provide a baseline which we can use to evaluate the additional benefits of using bit centering. To construct LP-SVRG in our mixed precision setting, we simply modify the standard SVRG algorithm (see Appendix B.1) so that its entire inner loop is computed in low-precision: this is effectively using the "quantize first" strategy that we analyzed in Section 3.2. For simplicity, we assume that all the inputs to the algorithm (the training examples, the labels, the initial iterate, and the step size) are given in the low-precision format (and representable without error in the high-precision format as well). The LP-SVRG algorithm is shown in Algorithm 4. In prior work (Harikandeh et al., 2015), it has been standard to use option II for the theoretical analysis (as it simplifies the derivation) while using option I for all empirical experiments; we follow this convention here.

We can immediately see that LP-SVRG will not converge asymptotically at a linear rate, as it will be limited to producing outputs representable in the low-precision format $\mathsf{fl}_{\text{lo}}(\cdot)$: once it gets as close as possible to the solution in this representation, it can get no closer, and convergence will stop. The next-best thing we can hope for is that LP-SVRG will converge at a linear rate until it reaches this limit, at which point it will stop converging—in fact, we can prove that this is what happens. First, we will state the assumptions that we need to prove this.

Assumption 1 *We require that the objective f is μ-strongly convex:*

$$(w - v)^T (\nabla f(w) - \nabla f(v)) \geq \mu \|w - v\|^2 \, ;$$

and the gradients ∇f_i are all L-Lipschitz continuous:

$$\|\nabla f_i(w) - \nabla f_i(v)\| = \left| x_i^T w - x_i^T v \right| \cdot \|x_i\| \leq L \|w - v\| \, .$$

These assumptions are standard in the analysis of optimization algorithms, and they are the same ones used for the analysis of SVRG (Johnson and Zhang, 2013). In terms of these parameters, the *condition number* of the problem is defined as $\kappa = L/\mu$.

Assumption 2 *We also assume, for simplicity, that the high-precision numbers used have zero error (i.e. the high-precision arithmetic is actual real number arithmetic).*

This assumption, while not strictly necessary to prove convergence, simplifies the analysis greatly. It means that the result of theorems proven with this assumption will only depend on the error of the low-precision representation, and so this assumption lets us cleanly capture what the error *just due to low-precision computation* is. For simplicity, we will also assume that no overflow or underflow occurs in the computation. Under these conditions we can provide convergence guarantees for LP-SVRG.

Theorem 3 *Suppose that we run LP-SVRG for linear models (Algorithm 4) under Assumptions 1 and 2, using option II for the epoch update. Suppose that no overflow or underflow occurs during the computation. Let d denote the dimension of the model, and let $w^* \in \mathbb{R}^d$ be the global optimum. Suppose that $d > 7$. For any constant $0 < \gamma < 1$ (a parameter which controls how often we take full gradients), if we set our step size and epoch lengths to be*

$$\alpha = \frac{\gamma}{4L(1+\gamma)} \qquad T \geq \frac{8\kappa(1+\gamma)}{\gamma^2},$$

which is the same as in the original analysis of SVRG in Johnson and Zhang (2013), then the outer iterates of LP-SVRG will converge to an accuracy limit at a linear rate:

$$\mathbf{E}\left[\tilde{f}(w_{K+1}) - f(w^*)\right] \leq \left(\gamma + \frac{288\kappa \cdot d \cdot \varepsilon_{\text{machine-lo}}}{\gamma}\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$

$$+ \frac{72L \cdot \|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}}}{\gamma} + O(\varepsilon_{\text{machine-lo}}{}^2).$$

Because it is long, we defer the proof of this theorem to the appendix. This result is tight in the sense that we can recover the original convergence theorem for SVRG by setting $\varepsilon_{\text{machine-lo}} = 0$ in the above expression. As a consequence, LP-SVRG will initially converge at a linear rate, just like SVRG—but only down to an error that is limited by the precision used. We can observe a tradeoff between bits-of-precision and accuracy here: as the number of bits becomes smaller, $\varepsilon_{\text{machine-lo}}$ will become larger, and so the accuracy limit will become worse. Also interesting is the fact that, for LP-SVRG to contract at each iteration by roughly the same factor as the base SVRG algorithm, we need

$$\frac{288\kappa \cdot d \cdot \varepsilon_{\text{machine-lo}}}{\gamma} \ll 1,$$

which requires that $\varepsilon_{\text{machine-lo}} \ll (\kappa d)^{-1}$. This suggests a relationship between the machine epsilon, and therefore the number of bits, that is needed and the condition number of the problem: for problems with a larger condition number (more poorly conditioned), we need more bits.

Bit-centered SVRG. While LP-SVRG converges at a linear rate, it only converges down to a level of accuracy proportional to the machine epsilon of the low-precision format used. In fact, this is a fundamental limitation of algorithms like LP-SVRG and LP-SGD: we cannot produce a solution that is asymptotically more accurate than the *most accurate solution representable* in the low-precision format we have chosen. Since the low-precision format in the previous section is fixed throughout the algorithm, this accuracy limitation is impossible to overcome. For the same reason, this limitation will still hold even if we use a non-floating-point low-precision representation, such as fixed-point arithmetic. While the use of SVRG allowed us to reach this minimum level of accuracy more quickly (at a linear rate, in fact) compared with low-precision SGD (De Sa et al., 2015; Li et al., 2017), it has not let us surpass this minimum.

In this section we introduce (for linear models) a bit-centered version of SVRG that *can* surpass this minimum level of accuracy and converge to arbitrarily accurate solutions (limited only by the high-precision machine epsilon) while still using low-precision arithmetic in the inner loop. To construct our bit-centered SVRG variant, we simply modify the standard SVRG algorithm to compute its inner loop using bit-centered computation, using the outer loop iterate $\tilde{w}_k$ as the offset. Recall that the inner loop of SVRG for linear regression has update step

$$w_{k,t} = w_{k,t-1} - \alpha \left((x_i^T w_{k,t-1} - y_i) \cdot x_i - (x_i^T \tilde{w}_k - y_i) \cdot x_i + \nabla f(\tilde{w}_k) \right).$$

If we re-write $w_{k,t}$ using bit centering with an offset of $\tilde{w}_k$, then we get the following update step for the delta,

$$
\begin{aligned}
\delta_{w,k,t} &= w_{k,t} - \tilde{w}_k \\
&= w_{k,t-1} - \tilde{w}_k - \alpha \left((x_i^T w_{k,t-1} - y_i) \cdot x_i - (x_i^T \tilde{w}_k - y_i) \cdot x_i + \nabla f(\tilde{w}_k) \right) \\
&= \delta_{w,k,t-1} - \alpha \left((x_i^T (\delta_{w,k,t-1} + \tilde{w}_k) - y_i) \cdot x_i - (x_i^T \tilde{w}_k - y_i) \cdot x_i + \nabla f(\tilde{w}_k) \right) \\
&= \delta_{w,k,t-1} - \alpha \left((x_i^T \delta_{w,k,t-1}) \cdot x_i + \nabla f(\tilde{w}_k) \right).
\end{aligned}
$$

We can use this bit-centered update step in place of the ordinary update step in SVRG. Doing this results in Algorithm 5, BC-SVRG.

Compared with LP-SVRG, BC-SVRG returns a full-precision vector $\tilde{w}_{K+1}$. As a result, it is not fundamentally limited to the outputs representable in the low-precision format, like LP-SVRG is. Also unlike LP-SVRG, the numerical error in the inner loop of BC-SVRG *decreases as the algorithm converges*. This is because all the numbers in the inner loop, including $\delta_{w,k,t}$ and $\tilde{h}_k$, are proportional to the distance of the outer loop iterate to the optimum $\|\tilde{w}_k - w^*\|$. So as $\tilde{w}_k$ gets closer to w^* (or, equivalently, as the delta gets smaller), the error of the floating point computation, which is relative to the magnitude of the numbers, also goes down. In the following theorem, we show that this is enough to allow BC-SVRG to converge down to solutions of accuracy limited only by machine epsilon of the high-precision format. More precisely, we show that if the high-precision format has zero error (i.e. is exact arithmetic) then BC-SVRG can output solutions arbitrarily close to the optimum w^*.

Theorem 4 *Suppose that we run BC-SVRG (Algorithm 5) under Assumptions 1 (strong convexity and Lipschitz continuity) and 2 (zero-error high-precision arithmetic), using option*

Algorithm 5 BC-SVRG: Bit-centered SVRG for linear regression

---

given: N low-precision training examples x_i and labels y_i, number of epochs K, epoch length T, low-precision step size α, initial low-precision iterate $\tilde{w}_1$.
for $k = 1$ **to** K **do**

> ▷ **_high-precision precomputation on host_**
>
> $\tilde{g}_k \leftarrow \underset{\text{H}}{\nabla} f(\mathsf{fl}_{\mathrm{hi}}(\tilde{w}_k)) = \left(\underset{i=1}{\overset{N}{\underset{\text{H}}{\sum}}} \left(\mathsf{fl}_{\mathrm{hi}}(x_i)^T \underset{\text{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(\tilde{w}_k) \underset{\text{H}}{\ominus} \mathsf{fl}_{\mathrm{hi}}(y_i) \right) \underset{\text{H}}{\odot} \mathsf{fl}_{\mathrm{hi}}(x_i) \right) \underset{\text{H}}{\oslash} N$
>
> $\tilde{h}_k \leftarrow \mathsf{fl}_{\mathrm{lo}}(\tilde{g}_k)$

$\delta_{w,k,0} \leftarrow \mathsf{fl}_{\mathrm{lo}}(0)$

> ▷ **_low-precision computation on accelerator_**
> **for** $t = 1$ **to** T **do**
>   **sample** i uniformly from $\{1, \ldots, N\}$
> $\delta_{w,k,t} \leftarrow \delta_{w,k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\text{L}}{\odot} x_i \underset{\text{L}}{\oplus} \tilde{h}_k \right)$
> **end for**

option I: set $\tilde{w}_{k+1} \leftarrow \tilde{w}_k \underset{\text{H}}{\oplus} \mathsf{fl}_{\mathrm{hi}}(\delta_{w,k,T})$
option II: sample t uniformly from $\{0, \ldots, T-1\}$, then set $\tilde{w}_{k+1} \leftarrow \tilde{w}_k \underset{\text{H}}{\oplus} \mathsf{fl}_{\mathrm{hi}}(\delta_{w,k,t})$
end for
return $\tilde{w}_{K+1}$

II for the epoch update. Also assume that no overflow or underflow occurs in the computation, and that $d \geq 4$. For any constant $0 < \gamma < 1$, if we set our step size and epoch lengths to be

$$\alpha = \frac{\gamma}{4L(1+\gamma)}, \qquad T \geq \frac{8\kappa(1+\gamma)}{\gamma^2},$$

then the iterates of BC-SVRG will satisfy

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\gamma + \frac{192\kappa d}{\gamma} \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2)) \right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].$$

This theorem shows that we can achieve a linear asymptotic convergence rate even using constant-bit-width low-precision computation in the inner loop. It also describes an interesting tradeoff between the precision and the condition number. As the condition number becomes larger while the precision stays fixed, we need to use longer epochs (T becomes larger) and the algorithm converges at a slower rate, until eventually (for κ large enough to make the contraction factor ≥ 1) the algorithm might stop working altogether. This suggests that low-precision training should be combined with techniques to improve the condition number, such as preconditioning, to achieve the best performance.

4.3 HALP: Handling overflow and underflow

From the result of Theorem 4, BC-SVRG for linear models had the remarkable property that it was able to converge to produce solutions of high accuracy, even though it used

entirely low-precision compute in its inner loop. However, this theorem had one unrealistic assumption: that no overflow or underflow would occur during the computation. When using real floating point representations, this assumption certainly will not hold. For example, imagine that we use 16-bit floating point numbers for the low-precision format, and some big-float arithmetic with very many bits for the high-precision format, such that we can assume that it behaves like exact arithmetic. As BC-SVRG converges, if it follows the convergence rate described by Theorem 4, eventually $\tilde{w}_k$ will become very close to w^*. But when this happens, $\tilde{g}_k$ will also become very small, and its low-precision version $\tilde{h}_k$ will underflow the 16-bit float format. Similar underflows will occur to $\delta_{w,k,t}$. As a result of these underflows, the algorithm may still have its accuracy limited by the 16-bit float format.

To address this, we introduce a technique called *dynamic bias adjustment* which prevents overflow and underflow from occurring by altering the *exponent bias* of the low-precision format dynamically. A floating point number with sign bit s, exponent e, and mantissa bits $m_1, m_2, \ldots$ has value

$$(-1)^s \cdot 2^{e-\text{bias}} \cdot 1.m_1 m_2 \ldots$$

where "bias" here denotes some integer value that is usually fixed for each floating point representation. For example, for 16-bit IEEE floats, the exponent bias is 15. Computing with a different exponent bias changes the range of the floating point representation. By increasing this bias as BC-SVRG converges, we can make the representable range smaller and so prevent underflow from occurring. To do this, we add an extra bias term to the exponent,

$$(-1)^s \cdot 2^{e-\text{bias}+\text{extra bias}} \cdot 1.m_1 m_2 \ldots.$$

To be consistent with our previous notation, we let $\text{fl}_{\text{lo-bias}(B)}(x)$ denote the result of converting x into a biased low-precision format with extra exponent bias B, and introduce analogous notation for computations with the biased-exponent format, which we describe in detail in Appendix B.2. Note that using an extra bias of B has the effect of changing the overflow and underflow thresholds to

$$M_{\text{lo-bias}(B)} = 2^B \cdot M_{\text{lo}} \qquad \eta_{\text{lo-bias}(B)} = 2^B \cdot \eta_{\text{lo}}.$$

This is just scaling both of them by the same factor, while the overall machine epsilon of the format (which is determined by size of the mantissa) remains the same, since the total number of bits is not changed when we use dynamic bias.

Setting the bias How do we decide how to set the bias? One simple heuristic is to use the full gradient we are already computing as part of SVRG. Since, as BC-SVRG converges, the magnitude of all the numbers in the inner loop becomes small proportional to the distance to the optimum, it will also become small proportional to the magnitude of $\tilde{g}$, which is within a constant factor of the distance to the optimum (specifically, $\mu \|\tilde{w}_k - \tilde{w}^*\| \le \|\tilde{g}_k\| \le L\mu \|\tilde{w}_k - \tilde{w}^*\|$). As a result, we can reduce the bias without overflowing anywhere if we set it such that $2^{-\text{bias}}$ is proportional to $\|\tilde{g}_k\|$. If we do this, then any underflow that occurs will also be small proportional to $\|\tilde{g}_k\|$, which means that it will be small proportional to the distance to the optimum. This, in turn, means that the numerical error due to underflow should become small at the same rate as the more ordinary numerical error due to the low-precision machine epsilon (that is, proportional to the distance to the optimum), so

---

Algorithm 6 HALP: Bit-centered, dynamic-bias-adjusted SVRG for linear models

---

    given: N low-precision training examples x_i and labels y_i, number of epochs K, epoch length T, low-precision step size α, initial low-precision iterate $\tilde{w}_1$, and bias factor $\zeta > 0$.

    for $k = 1$ **to** K **do**

      $\triangleright$ *high-precision precomputation on host*

$$\tilde{g}_k \leftarrow \underset{\text{H}}{\nabla} f(\mathsf{fl}_{\text{hi}}(\tilde{w}_k)) = \left(\underset{i=1}{\overset{N}{\sum_{\text{H}}}} \left(\mathsf{fl}_{\text{hi}}(x_i)^T \underset{\text{H}}{\odot} \mathsf{fl}_{\text{hi}}(\tilde{w}_k) \underset{\text{H}}{\ominus} \mathsf{fl}_{\text{hi}}(y_i) \right) \underset{\text{H}}{\odot} \mathsf{fl}_{\text{hi}}(x_i) \right) \underset{\text{H}}{\oslash} N$$

      compute bias: $B \leftarrow \left\lfloor \underset{\text{H}}{\log_2} \left(\zeta \cdot \|\tilde{g}_k\|_{\text{H}} \right) \right\rfloor$

      $\tilde{h}_k \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(\tilde{g}_k)$

      $\tilde{s}_k \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(2 \underset{\text{H}}{\odot} \|\tilde{g}_k\|_{\text{H}} \underset{\text{H}}{\oslash} \mu)$

    $\delta_{w,k,0} \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(0)$

      $\triangleright$ *low-precision computation on accelerator*

      for $t = 1$ **to** T **do**

        sample i uniformly from $\{1, \dots, N\}$

        $\delta_{w,k,t} \leftarrow \delta_{w,k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\text{L}}{\odot} x_i \underset{\text{L}}{\oplus} \tilde{h}_k \right)$

        if $\|\delta_{w,k,t}\|_{\text{L}} > \tilde{s}_k$ **then**

          $\delta_{w,k,t} \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(0)$

        end if

      end for

    option I: set $\tilde{w}_{k+1} \leftarrow \tilde{w}_k \underset{\text{H}}{\oplus} \mathsf{fl}_{\text{hi-from}(B)}(\delta_{w,k,T})$

    option II: sample t uniformly from $\{0, \dots, T-1\}$,

               then set $\tilde{w}_{k+1} \leftarrow \tilde{w}_k \underset{\text{H}}{\oplus} \mathsf{fl}_{\text{hi-from}(B)}(\delta_{w,k,t})$

  end for

  return $\tilde{w}_{K+1}$

---

we should be able to get a linear rate of convergence. Using this heuristic to assign the exponent bias results in Algorithm 6, which we call *high-accuracy low-precision* or HALP.

Compared with BC-SVRG, in addition to the dynamic biasing, the only change in Algorithm 6 is the extra `if` statement comparing with $\tilde{s}_k$. This statement acts to prevent another kind of overflow that *could* still occur when running HALP, if this if statement were not present. Without this statement, the updates to $\delta_{w,k,t}$ could accumulate until it becomes large enough that computations like $x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}$ overflow. The if statement guards against this happening by preventing $\delta_{w,k,t}$ from becoming too large. The resetting of $\delta_{w,k,t}$ to 0 if it fails this check is justified by the fact that by strong convexity,

$$\|\tilde{w}_k - w^*\| \leq \frac{1}{\mu} \|\nabla f(\tilde{w}_k) - \nabla f(w^*)\| = \frac{1}{\mu} \|\tilde{g}_k\| = \frac{\tilde{s}_k}{2},$$
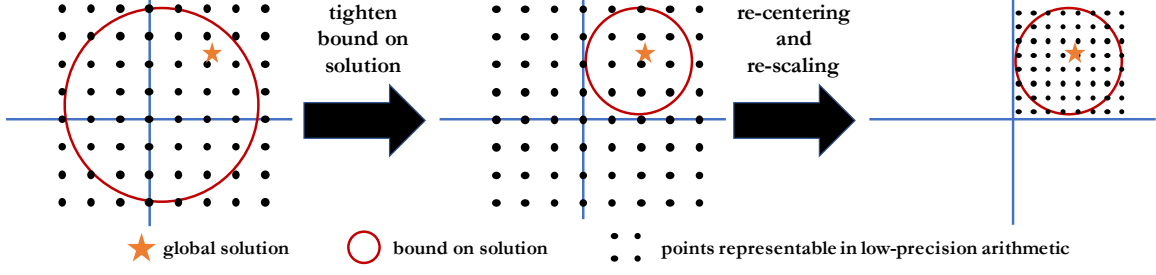
Figure 4: A diagram of the bit centering and dynamic biasing operation in HALP. As the algorithm converges, we are able to bound the solution within a smaller and smaller ball. Periodically, we re-center the points that our low-precision model can represent so they are centered on this ball, and we re-scale the points so that more of them are inside the ball.

and so if $\delta_{w,k,t}$ has norm larger than $\tilde{s}_k$, it follows that

$$\|w_{k,t} - w^*\| = \|\delta_{w,k,t} + \tilde{w}_k - w^*\| \geq \|\delta_{w,k,t}\| - \|\tilde{w}_k - w^*\| > \tilde{s}_k - \frac{\tilde{s}_k}{2} = \frac{\tilde{s}_k}{2} \geq \|\tilde{w}_k - w^*\| .$$

This means that by resetting the delta to zero, we are actually moving closer to the global optimum. Equivalently, we can think about the whole process as producing some bound on where the solution could be (using strong convexity) and then remapping our representable low-precision numbers, using both bit centering with an offset and re-scaling with dynamic bias adjustment, to cover the region that we know our solution is in. This process is illustrated in Figure 4. Using this, we can prove the following theorem, which shows that HALP converges to produce arbitrarily accurate solutions, even when using low-precision floating point numbers for which overflow and underflow are possible.

Theorem 5 *Suppose that we run HALP for linear regression (Algorithm 6) under Assumptions 1 (strong convexity and Lipschitz continuity) and 2 (zero-error high-precision arithmetic), using option II for the epoch update. Suppose that $d > 16$, that our underflow threshold $\eta_{\text{machine-lo}}$ is small enough that*

$$(L+1) \cdot \zeta \cdot \eta_{\text{machine-lo}} \leq \varepsilon_{\text{machine-lo}},$$

and that our overflow threshold is large enough that

$$\frac{4\kappa + 2}{\zeta} \cdot \max\left(1, L^{-1}\right) \cdot (1 + O(\varepsilon_{\text{machine-lo}})) \leq M_{\text{machine-lo}},$$

For any constant $0 < \gamma < 1$, if we set our step size and epoch lengths to be

$$\alpha = \frac{\gamma}{4L(1+\gamma)}, \qquad T \geq \frac{8\kappa(1+\gamma)}{\gamma^2},$$

then the iterates of HALP will satisfy

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\gamma + \frac{192\kappa d}{\gamma} \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].$$
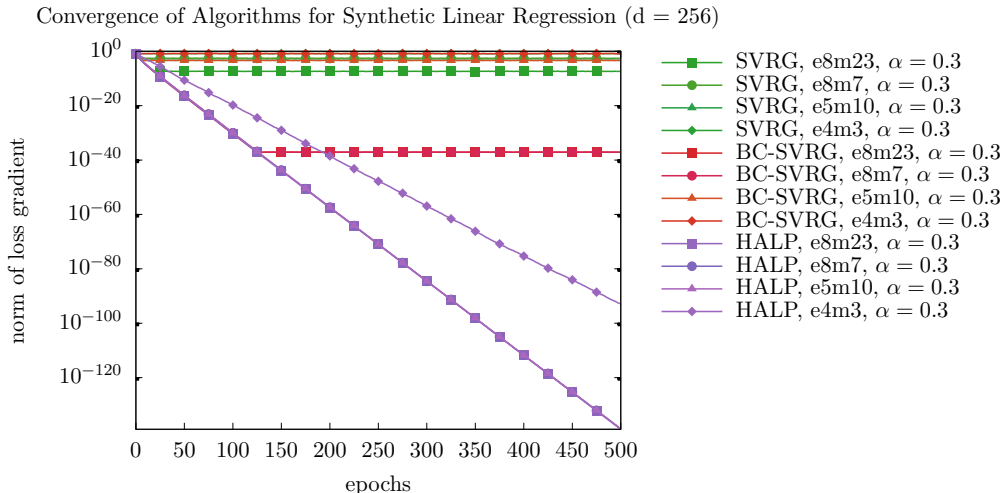
Figure 5: A comparison of the convergence of LP-SVRG and BC-SVRG (without dynamic biasing) and HALP for a variety of low-precision numerical formats on a synthetic linear regression dataset.

Experimental validation. Theorem 5 predicts that we can get highly accurate solutions to linear regression problems using only low-precision compute in the inner loop—and *real* low-precision compute, not just an idealized version with no overflow or underflow. We validate this result empirically in Figure 5 where we ran low-precision SVRG, bit-centered SVRG, and HALP for a variety of low-precision formats on the same synthetic dataset that was used for Figure 2. Here, we used a learning rate of $\alpha = 0.3$ for all experiments, we used randomized rounding for the low-precision formats, and we used a 1024-bit BigFloat format for the high-precision numbers used in these algorithms in order to match the setting of Theorem 5, in which Assumption 2 holds, as closely as possible. Figure 5 illustrates that HALP is able to converge down to solutions of arbitrarily low accuracy (in this figure, ranging down to 10^{-120} which is even smaller than the machine epsilon of 64-bit floating point arithmetic) even while using only low-precision computation in its inner loop. Even an 8-bit floating point representation with 4 exponent bits and 3 mantissa bits is able to achieve this feat, albeit at a somewhat slower convergence rate.

Effect of the condition number. Another very interesting effect predicted by our theorems in this section is a relationship between the number of bits needed and the condition number κ of the problem. To study this effect, we generated a series of simple two-dimensional linear regression problems with different condition numbers. For each condition number κ we sought to evaluate, we generated 512 synthetic examples of the form

$$x_i = [1, 0]^T \qquad y_i \sim \mathcal{N}(0, 1),$$

and combined them with 512 synthetic examples of the form

$$x_i = \left[0, (\kappa/2)^{-1/2}\right]^T \qquad y_i \sim \mathcal{N}\left(0, (\kappa/2)^{1/2}\right).$$

23

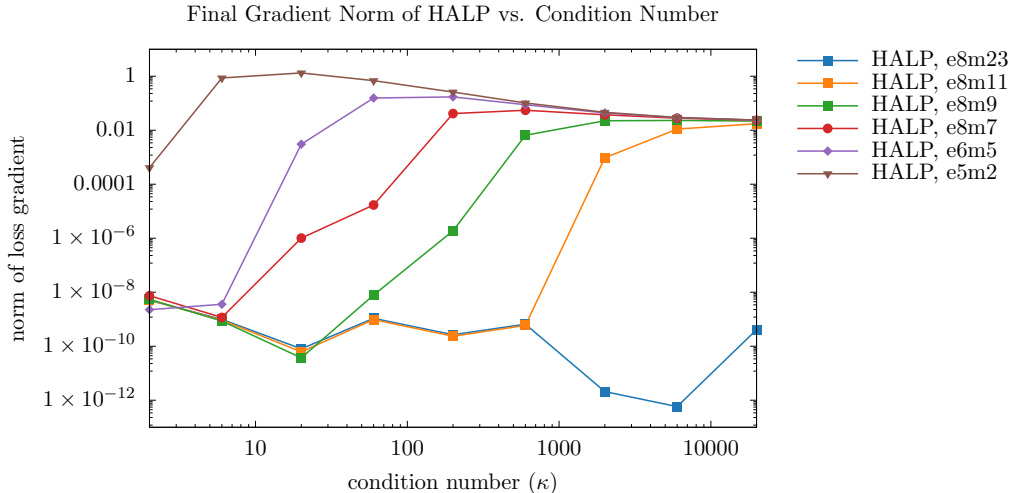Final Gradient Norm of HALP vs. Condition Number

Figure 6: A comparison of the convergence of HALP on synthetic linear regression tasks as the condition number is changed. Notice that as the condition number increases, the performance of the lower-precision experiments becomes worse than that of the 32-bit HALP baseline.

It is straightforward to verify that this example indeed has condition number κ. For each condition number in the set

$$\kappa \in \{2, 6, 20, 60, 200, 600, 2000, 6000, 20000\}$$

and for several different low-precision formats ranging from 8-bits to 32-bits, we ran HALP for $K = 10$ epochs each of length $T = 30\kappa$ using a learning rate of $\alpha = 0.1$, which is what the theory in Theorem 3 predicts we should use for a setting of $\gamma = 0.667$. We used nearest-neighbor rounding for all the quantization done in this experiment.[6] The final gradient norm after 10 epochs is reported in Figure 6.

Notice that for low-precision HALP, there is a threshold where the performance of the algorithm starts to degrade, above which the performance is worse than the 32-bit HALP baseline. This validates our theory, which predicts that such a phenomenon will happen, and places a bound on the maximum machine epsilon (or equivalently, the minimum number of bits) that will suffice for solving problems of a particular condition number.

5. Bit Centering for General Loss Functions

In the previous sections, we described how to learn using bit centering for linear regression, which has a loss function of particularly simple form involving only a dot product and scalar-vector computation. In this section, we describe how bit centering can be used on arbitrary numerical computations. The key idea is that, in the linear regression and dot product case, we were able to produce a bound on the error of the delta of the gradients (or

---

6. We reported results for nearest-neighbor rounding here because when randomized rounding is used for the same experiment, other effects dominate the convergence and the relationship between condition number and precision is less visible.

of the dot product) of the form

$$\text{error}(\delta) = O\left(\|\delta_w\| \cdot \varepsilon_{\text{machine-lo}}\right).$$

This bound allowed us to show that the errors are getting smaller as the algorithm converges. Next, we will show how we can develop bit-centered versions of all common numerical operations, and thereby get a bound of the same form on the delta error of the gradients for *arbitrary* loss functions which can be computed via a computation graph. Each bit-centered operation proceeds in three steps:

1. During high-precision precomputation, compute the (high-precision) offset of the output of the operation, as a function of the (high-precision) offsets of the inputs.

2. During high-precision precomputation, compute (in high-precision) and *cache* (in low-precision) any extra values that will be needed to compute the deltas later.

3. During low-precision computation, compute the (low-precision) delta of the output of the operation, as a function of the (low-precision) deltas of the inputs and the cached values (if any).

We introduce the following notation for the computations done in each of these steps:

$$o_{\underset{\text{C}}{\circledast}}(o_x, o_y)$$

which computes (usually in high precision on the host device) the offset of the output of the $*$ operation, given the offset of the inputs,

$$\text{cache}_{\underset{\text{C}}{\circledast}}(o_x, o_y)$$

which computes (usually in high precision on the host device) any low-precision values that will be needed later for computation of the deltas, and

$$\delta_{\underset{\text{C}}{\circledast}}(\delta_x, \delta_y, \text{cache})$$

which computes (usually in low precision on the accelerator) the delta of the offset of the operation, given the low-precision deltas of the inputs and the low-precision values precomputed in cache. To represent an entire bit-centered operation, we let (o_x, δ_x) denote the value x represented using bit centering (that is, we represent it as $x = o_x + \delta_x$ using a high-precision-floating-point offset o_x and a low-precision-floating-point delta δ_x, stored in a biased low-precision format with exponent bias B[7]) and let

$$(o_x, \delta_x) \underset{\text{C}}{\circledast} (o_y, \delta_y)$$

denote the entire bit-centered computation of the $*$ operation, just as we have used similar notation to denote high-precision and low-precision computation.

In this section, we will explicitly define how to perform these bit-centered operations for most numerical operations used in machine learning. Then we compute the error of these bit-centered operations just like we would compute the error of floating point operations. Our results in this vein are summarized in Table 2.

---

7. This handles the case of unbiased low-precision deltas as a special case with $B = 0$.

Operation	# LP Ops	# HP Ops	Result error ($\cdot \varepsilon_{\mathrm{lo}}$)
Add	1 add	1 add	$\lvert \delta_x + \delta_y \rvert$
Subtract	1 sub	1 sub	$\lvert \delta_x - \delta_y \rvert$
Multiply x by constant y	1 multiply	1 multiply	$\lvert \delta_x \cdot y \rvert$
General multiply	2 extra loads 2 multiplies 2 adds	1 multiply 2 extra stores	$4 \cdot (\lvert o_x \cdot \delta_y \rvert + \lvert \delta_x \cdot o_y \rvert + \lvert \delta_x \cdot \delta_y \rvert)$
Divide	2 extra loads 1 multiply 1 divide 1 add 1 subtract	1 divide 2 extra stores	$7 \cdot \frac{\lvert \delta_x \rvert + \lvert o_x / o_y \rvert \cdot \lvert \delta_y \rvert}{y^2} \cdot (\lvert y \rvert + \lvert \delta_y \rvert)$
Exp	1 extra load 1 multiply 1 expm1	1 exp 1 extra store	$4 \cdot \exp(o_x) \cdot (\exp(\delta_x) - 1)$
Function eval	1 extra load 1 eval of Δf	1 eval of f 1 extra store	$\lvert f'(o_x) \rvert \cdot \lvert \delta_x \rvert + L_f \cdot \lvert \delta_x \rvert \cdot (\lvert o_x \rvert + \lvert \delta_x \rvert)$

Table 2: A table summarizing our results on bounding the error of bit-centered operations in Section 5. For simplicity, we suppose in this table that $\varepsilon_{\mathrm{machine\text{-}hi}} = 0$ and $\eta_{\mathrm{machine\text{-}lo}} = 0$; i.e. we are looking only at the error caused by the low-precision format and we are not considering underflow.

Addition Suppose that we want to do addition using bit centering, i.e. to compute $x \underset{\text{C}}{\oplus} y$ where $x = o_x + \delta_x$ and $y = o_y + \delta_y$. The natural way to do this is to let

$$x \underset{\text{C}}{\oplus} y \overset{\text{def}}{=} ((o_x \underset{\text{H}}{\oplus} o_y), (\delta_x \underset{\text{L}}{\oplus} \delta_y)),$$

where more explicitly the offset of the bit-centered result is the first term $o_x \underset{\text{H}}{\oplus} o_y$ and the delta is the second term $\delta_x \underset{\text{L}}{\oplus} \delta_y$. Explicitly, this decomposes into

$$o_{\underset{\text{C}}{\oplus}}(o_x, o_y) \overset{\text{def}}{=} o_x \underset{\text{H}}{\oplus} o_y$$

$$\text{cache}_{\underset{\text{C}}{\oplus}}(o_x, o_y) \overset{\text{def}}{=} \emptyset$$

$$\delta_{\underset{\text{C}}{\oplus}}(\delta_x, \delta_y, \emptyset) \overset{\text{def}}{=} \delta_x \underset{\text{L}}{\oplus} \delta_y.$$

Note that we say $\text{cache}_{\underset{\text{C}}{\oplus}}(o_x, o_y) \overset{\text{def}}{=} \emptyset$ to denote that *nothing* is pre-computed and stored by this function for an add operation. This will have error

$$\left| x \underset{\text{C}}{\oplus} y - (x + y) \right| = \left| (o_x \underset{\text{H}}{\oplus} o_y) + (\delta_x \underset{\text{L}}{\oplus} \delta_y) - (o_x + o_y + \delta_x + \delta_y) \right|$$

$$\leq \left| (o_x \underset{\text{H}}{\oplus} o_y) - (o_x + o_y) \right| + \left| (\delta_x \underset{\text{L}}{\oplus} \delta_y) - (\delta_x + \delta_y) \right|$$

$$= |o_x + o_y| \cdot \varepsilon_{\text{hi}} + \eta_{\text{hi}} + |\delta_x + \delta_y| \cdot \varepsilon_{\text{lo}} + \eta_{\text{lo-bias}(B)}$$

$$= |x + y| \cdot \varepsilon_{\text{hi}} + |\delta_x + \delta_y| \cdot (\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + \eta_{\text{lo-bias}(B)} + \eta_{\text{hi}}.$$

It will also be the case that the magnitude of the delta that results from this operation will be bounded by

$$\left| \delta_{x \underset{\text{C}}{\oplus} y} \right| = \left| \delta_x \underset{\text{L}}{\oplus} \delta_y \right| \leq |\delta_x + \delta_y| \cdot \varepsilon_{\text{lo}} + \eta_{\text{lo-bias}(B)}.$$

Subtraction We define

$$x \underset{\text{C}}{\ominus} y \overset{\text{def}}{=} ((o_x \underset{\text{H}}{\ominus} o_y), (\delta_x \underset{\text{L}}{\ominus} \delta_y)).$$

The analysis of this is the same as for addition, so we defer it to Appendix B.3 for brevity.

Constant multiplication Suppose that we want to do multiplication *by a constant* using bit centering, to compute $x \underset{\text{C}}{\odot} y$ where $x = o_x + \delta_x$ and y is a constant representable in low-precision (but is not itself represented using bit centering). This, it turns out, is easier than doing general bit-centered multiplication, which we will consider next. The natural way to do constant multiplication is

$$x \underset{\text{C}}{\odot} y \overset{\text{def}}{=} \left(o_x \underset{\text{H}}{\odot} \text{fl}_{\text{hi}}(y), \delta_x \underset{\text{L}}{\odot} \text{fl}_{\text{lo}}(y) \right).$$

Explicitly, this decomposes into

$$o_{\underset{C}{\odot}}(o_x, y) \stackrel{\text{def}}{=} o_x \underset{H}{\odot} y$$

$$\text{cache}_{\underset{C}{\odot}}(o_x, y) \stackrel{\text{def}}{=} \mathsf{fl}_{\text{lo}}(y)$$

$$\delta_{\underset{C}{\odot}}(\delta_x, (y_{\text{lo}})) \stackrel{\text{def}}{=} \delta_x \underset{L}{\odot} y_{\text{lo}}.$$

This will have error

$$\left| x \underset{C}{\odot} y - (x \cdot y) \right| = \left| (o_x \underset{H}{\odot} \mathsf{fl}_{\text{hi}}(y)) + (\delta_x \underset{L}{\odot} \mathsf{fl}_{\text{lo}}(y)) - (o_x \cdot y + \delta_x \cdot y) \right|$$

$$\leq \left| (o_x \underset{H}{\odot} y) - (o_x \cdot y) \right| + \left| (\delta_x \underset{L}{\odot} y) - (\delta_x \cdot y) \right|$$

$$= |o_x \cdot y| \cdot \varepsilon_{\text{hi}} + \eta_{\text{hi}} + |\delta_x \cdot y| \cdot \varepsilon_{\text{lo}} + \eta_{\text{lo-bias}(B)}$$

$$= |x \cdot y| \cdot \varepsilon_{\text{hi}} + |\delta_x \cdot y| \cdot (\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + \eta_{\text{lo-bias}(B)} + \eta_{\text{hi}}.$$

The magnitude of the delta that results from this operation will be bounded by

$$\left| \delta_{x \odot y} \right| = \left| \delta_x \underset{L}{\odot} y \right| \leq |\delta_x \cdot y| \cdot \varepsilon_{\text{lo}} + \eta_{\text{lo-bias}(B)}.$$

Multiplication   Suppose that we want to do multiplication using bit centering, to compute $x \underset{C}{\odot} y$ where $x = o_x + \delta_x$ and $y = o_y + \delta_y$. The natural way to do this is

$$x \underset{C}{\odot} y \stackrel{\text{def}}{=} \left((o_x \underset{H}{\odot} o_y), \mathsf{fl}_{\text{lo}}(o_x) \underset{L}{\odot} \delta_y \underset{L}{\oplus} \delta_x \underset{L}{\odot} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{L}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right) \right).$$

Explicitly, this decomposes into

$$o_{\underset{C}{\odot}}(o_x, o_y) \stackrel{\text{def}}{=} o_x \underset{H}{\odot} o_y$$

$$\text{cache}_{\underset{C}{\odot}}(o_x, o_y) \stackrel{\text{def}}{=} \left(\mathsf{fl}_{\text{lo}}(o_x), \mathsf{fl}_{\text{lo}}(o_y) \right)$$

$$\delta_{\underset{C}{\odot}}(\delta_x, \delta_y, (\sigma_x, \sigma_y)) \stackrel{\text{def}}{=} \sigma_x \underset{L}{\odot} \delta_y \underset{L}{\oplus} \delta_x \underset{L}{\odot} \left(\sigma_y \underset{L}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right).$$

Here, σ_x and σ_y denote the low-precision numbers stored in the cache and used later by the low-precision delta computation. Note that since δ_y is stored in low-precision with a biased exponent, we explicitly write $\mathsf{fl}_{\text{lo}}(\delta_y)$ here to indicate that we are converting it to the unbiased-exponent low-precision format. In order to compute this delta, we now need to additionally quantize and load low-precision version of o_x and o_y; this will require additional

28

loads on the accelerator. The error of this will be given by

$$
\begin{aligned}
x \underset{\text{C}}{\odot} y &= (o_x \underset{\text{H}}{\odot} o_y) + \mathsf{fl}_{\text{lo}}(o_x) \underset{\text{L}}{\odot} \delta_y \underset{\text{L}}{\oplus} \delta_x \underset{\text{L}}{\odot} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{\text{L}}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right) \\
&= (o_x \cdot o_y) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}} \\
&\quad + \left(\mathsf{fl}_{\text{lo}}(o_x) \underset{\text{L}}{\odot} \delta_y + \delta_x \underset{\text{L}}{\odot} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{\text{L}}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right) \right) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)} \\
&= (o_x \cdot o_y) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}} \\
&\quad + \Bigg((o_x \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}}) \cdot \delta_y \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)} \\
&\qquad + \delta_x \cdot ((o_y \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}} + \delta_y \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}}) \cdot (1 + \varepsilon_{\text{lo}})) \cdot (1 + \varepsilon_{\text{lo}}) \\
&\qquad + \eta_{\text{lo-bias}(B)} \Bigg) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)} \\
&= (o_x \cdot o_y) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}} + o_x \cdot \delta_y \cdot (1 + 3\varepsilon_{\text{lo}}) + \delta_y \cdot \eta_{\text{lo}} \\
&\quad + \delta_x \cdot (o_y + \delta_y) \cdot (1 + 4\varepsilon_{\text{lo}}) + \delta_x \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)} \\
&= o_x \cdot o_y + o_x \cdot \delta_y + \delta_x \cdot (o_y + \delta_y) \\
&\quad + (o_x \cdot o_y) \cdot \varepsilon_{\text{hi}} + \eta_{\text{hi}} + o_x \cdot \delta_y \cdot 3\varepsilon_{\text{lo}} + \delta_y \cdot \eta_{\text{lo}} \\
&\quad + \delta_x \cdot (o_y + \delta_y) \cdot 4\varepsilon_{\text{lo}} + \delta_x \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)} \\
&= xy + ((x - \delta_x) \cdot (y - \delta_y)) \cdot \varepsilon_{\text{hi}} + \eta_{\text{hi}} + (x - \delta_x) \cdot \delta_y \cdot 3\varepsilon_{\text{lo}} + \delta_y \cdot \eta_{\text{lo}} \\
&\quad + \delta_x \cdot y \cdot 4\varepsilon_{\text{lo}} + \delta_x \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)} \\
&= xy + xy \cdot \varepsilon_{\text{hi}} + \delta_x y \cdot (4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + x\delta_y \cdot (3\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + \delta_x \delta_y \cdot (3\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) \\
&\quad + \eta_{\text{hi}} + \delta_y \cdot \eta_{\text{lo}} + \delta_x \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)} \\
&= xy + xy \cdot \varepsilon_{\text{hi}} + (|\delta_x y| + |x\delta_y| + |\delta_x \delta_y|) \cdot (4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) \\
&\quad + \eta_{\text{hi}} + (|\delta_y| + |\delta_x|) \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)}.
\end{aligned}
$$

It follows that the error will be bounded by

$$
\begin{aligned}
\left| x \underset{\text{C}}{\odot} y - (x \cdot y) \right| &= xy \cdot \varepsilon_{\text{hi}} + (|\delta_x y| + |x\delta_y| + |\delta_x \delta_y|) \cdot (4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) \\
&\quad + \eta_{\text{hi}} + (|\delta_y| + |\delta_x|) \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)}.
\end{aligned}
$$

The magnitude of the delta that results from this operation will be bounded by

$$
\begin{aligned}
\left| \delta_{x \underset{\text{C}}{\odot} y} \right| &= \left| \mathsf{fl}_{\text{lo}}(o_x) \underset{\text{L}}{\odot} \delta_y \underset{\text{L}}{\oplus} \delta_x \underset{\text{L}}{\odot} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{\text{L}}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right) \right| \\
&= \left| o_x \cdot \delta_y \cdot (1 + 3\varepsilon_{\text{lo}}) + \delta_y \cdot \eta_{\text{lo}} + \delta_x \cdot (o_y + \delta_y) \cdot (1 + 4\varepsilon_{\text{lo}}) + \delta_x \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)} \right| \\
&= |\delta_y| \cdot |o_x| \cdot (1 + 3\varepsilon_{\text{lo}}) + |\delta_y| \cdot \eta_{\text{lo}} + |\delta_x| \cdot y \cdot (1 + 4\varepsilon_{\text{lo}}) + |\delta_x| \cdot 2\eta_{\text{lo}} + 3\eta_{\text{lo-bias}(B)}.
\end{aligned}
$$

Division   Suppose that we want to do division using bit centering, to compute $x \underset{\text{C}}{\oslash} y$ where $x = o_x + \delta_x$ and $y = o_y + \delta_y$. The natural way to do this is

$$
x \underset{\text{C}}{\oslash} y \overset{\text{def}}{=} \left((o_x \underset{\text{H}}{\oslash} o_y), \left(\delta_x \underset{\text{L}}{\ominus} (\mathsf{fl}_{\text{lo}}(o_x \underset{\text{H}}{\oslash} o_y) \underset{\text{L}}{\odot} \delta_y) \right) \underset{\text{L}}{\oslash} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{\text{L}}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right) \right).
$$

Explicitly, this decomposes into

$$o_{\underset{C}{\oslash}}(o_x, o_y) \overset{\text{def}}{=} o_x \underset{H}{\oslash} o_y$$

$$\text{cache}_{\underset{C}{\oslash}}(o_x, o_y) \overset{\text{def}}{=} (\mathsf{fl}_{\mathrm{lo}}(o_x \underset{H}{\oslash} o_y), \mathsf{fl}_{\mathrm{lo}}(o_y))$$

$$\delta_{\underset{C}{\odot}}(\delta_x, \delta_y, (\sigma_{x/y}, \sigma_y)) \overset{\text{def}}{=} \left(\delta_x \underset{L}{\ominus} (\sigma_{x/y} \underset{L}{\odot} \delta_y) \right) \underset{L}{\oslash} \left(\sigma_y \underset{L}{\oplus} \mathsf{fl}_{\mathrm{lo}}(\delta_y) \right).$$

Note that, as was the case for multiplication, in order to do this we will need to quantize and load low-precision versions of o_y and $o_x \underset{H}{\oslash} o_y$, which will result in additional loads on the accelerator. The error analysis for this operation is a bit longer than for the previous operations, so we defer it to Appendix B.3.

Exponential function Now that we have analyzed the four primary arithmetic operations, we can start looking at special functions. The most basic one of these is the exponential function. Given $x = o_x + \delta_x$, we define

$$\underset{C}{\exp}(x) \overset{\text{def}}{=} \left(\underset{H}{\exp}(o_x), \mathsf{fl}_{\mathrm{lo}}(\underset{H}{\exp}(o_x)) \underset{L}{\odot} \underset{L}{\mathrm{expm1}}(\delta_x) \right).$$

Here, $\mathrm{expm1}(\cdot)$ denotes the function

$$\mathrm{expm1}(z) = \exp(z) - 1$$

computed as a single floating-point operation following the fundamental axiom of floating point arithmetic. This exp-minus-one operation is recommended in the IEEE standard (IEEE, 2008) and is usually included in floating point libraries, where it provides more precision for small values of z than would result from actually performing the subtraction.

Explicitly, $\underset{C}{\exp}(x)$ decomposes into

$$o_{\underset{C}{\exp}}(o_x) \overset{\text{def}}{=} \underset{H}{\exp}(o_x)$$

$$\text{cache}_{\underset{C}{\exp}}(o_x) \overset{\text{def}}{=} \mathsf{fl}_{\mathrm{lo}}(\underset{H}{\exp}(o_x))$$

$$\delta_{\underset{C}{\exp}}(\delta_x, \delta_y, (\sigma)) \overset{\text{def}}{=} \sigma \underset{L}{\odot} \underset{L}{\mathrm{expm1}}(\delta_x).$$

30

Using this decomposition, we will have an error analysis of

$$
\begin{aligned}
\underset{\text{C}}{\exp}(x) &= \underset{\text{H}}{\exp}(o_x) + \mathsf{fl}_{\text{lo}}(\underset{\text{H}}{\exp}(o_x)) \underset{\text{L}}{\odot} \underset{\text{L}}{\text{expm1}}(\delta_x) \\
&= \exp(o_x) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}} + \mathsf{fl}_{\text{lo}}(\underset{\text{H}}{\exp}(o_x)) \cdot \underset{\text{L}}{\text{expm1}}(\delta_x) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)} \\
&= \exp(o_x) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}} \\
&\quad + ((\exp(o_x) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}}) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}}) \\
&\qquad \cdot \big(\text{expm1}(\delta_x) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)}\big) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)} \\
&= \exp(o_x) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}} \\
&\quad + (\exp(o_x) \cdot (1 + 2\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + \eta_{\text{lo}} + \eta_{\text{hi}}) \\
&\qquad \cdot \big(\text{expm1}(\delta_x) \cdot (1 + 2\varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)}\big) + \eta_{\text{lo-bias}(B)} \\
&= \exp(o_x) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}} + \exp(o_x) \, \text{expm1}(\delta_x) \cdot (1 + 4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) \\
&\quad + \exp(o_x) \cdot \eta_{\text{lo-bias}(B)} + \text{expm1}(\delta_x) \cdot (\eta_{\text{lo}} + \eta_{\text{hi}}) + \eta_{\text{lo-bias}(B)} \\
&= \exp(x) + \exp(o_x) \cdot \varepsilon_{\text{hi}} + \eta_{\text{hi}} + \exp(o_x) \, \text{expm1}(\delta_x) \cdot (4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) \\
&\quad + \exp(o_x) \cdot \eta_{\text{lo-bias}(B)} + \text{expm1}(\delta_x) \cdot (\eta_{\text{lo}} + \eta_{\text{hi}}) + \eta_{\text{lo-bias}(B)}.
\end{aligned}
$$

So, the error of this computation will be bounded by

$$
\begin{aligned}
\left| \underset{\text{C}}{\exp}(x) - \exp(x) \right| &\leq \exp(o_x) \cdot \varepsilon_{\text{hi}} + \exp(o_x) \cdot \text{expm1}(\delta_x) \cdot (4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) \\
&\quad + \exp(o_x) \cdot \eta_{\text{lo-bias}(B)} + \text{expm1}(\delta_x) \cdot (\eta_{\text{lo}} + \eta_{\text{hi}}) + \eta_{\text{lo-bias}(B)} + \eta_{\text{hi}}.
\end{aligned}
$$

As usual, the parts of this expression that depend on the low-precision format will become small as δ_x becomes small (because $\text{expm1}(\delta_x)$ becomes small as δ_x becomes small). The magnitude of the delta of the result of this operator will be bounded by

$$
\begin{aligned}
\left| \underset{\text{C}}{\delta_{exp(x)}} \right| &= \left| \mathsf{fl}_{\text{lo}}(\underset{\text{H}}{\exp}(o_x)) \underset{\text{L}}{\odot} \underset{\text{L}}{\text{expm1}}(\delta_x) \right| \\
&= \Big| \exp(o_x) \, \text{expm1}(\delta_x) \cdot (1 + 4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) \\
&\qquad + \exp(o_x) \cdot \eta_{\text{lo-bias}(B)} + \text{expm1}(\delta_x) \cdot (\eta_{\text{lo}} + \eta_{\text{hi}}) + \eta_{\text{lo-bias}(B)} \Big| \\
&= |\delta_x| \cdot \left| \exp(o_x) \cdot \frac{\text{expm1}(\delta_x)}{\delta_x} \cdot (1 + 4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + \frac{\text{expm1}(\delta_x)}{\delta_x} \cdot (\eta_{\text{lo}} + \eta_{\text{hi}}) \right| \\
&\quad + \left| \eta_{\text{lo-bias}(B)} \right| + \left| \exp(o_x) \cdot \eta_{\text{lo-bias}(B)} \right|.
\end{aligned}
$$

Just as for the previous operators, this is becoming small as δ_x becomes small.

General function application To go beyond the above analysis of the exponential function, we now want to produce bit-centered versions of differentiable and Lipschitz continuous scalar functions. Suppose that we want to apply a function f to a value stored using bit centering, to compute $f(x)$ where $x = o_x + \delta_x$. Assume that our floating-point library allows us to compute f in high-precision arithmetic, following the fundamental axiom. There are a couple of ways we could use bit centering to compute f. The simplest of these

is to suppose that our low-precision floating-point library allows us to directly compute a function Δf that satisfies

$$\Delta f(x, y) = f(x + y) - f(x)$$

following the fundamental axiom (and for possibly exponent-biased y). If our library supports this, then we can define

$$\underset{\text{C}}{f}(x) \overset{\text{def}}{=} \left(\underset{\text{H}}{f}(o_x), \underset{\text{L}}{\Delta f}(\mathsf{fl}_{\mathsf{lo}}(o_x), \delta_x) \right).$$

Explicitly, we can write this as

$$\underset{\text{C}}{o_f}(o_x) \overset{\text{def}}{=} \underset{\text{H}}{f}(o_x)$$

$$\underset{\text{C}}{\mathrm{cache}_f}(o_x) \overset{\text{def}}{=} \mathsf{fl}_{\mathsf{lo}}(o_x)$$

$$\underset{\text{C}}{\delta_f}(\delta_x, (\sigma_x)) \overset{\text{def}}{=} \underset{\text{L}}{\Delta f}(\sigma_x, \delta_x).$$

If our function f is differentiable and f' is Lipschitz continuous with parameter L_f, then we can characterize the error with

$$\underset{\text{C}}{f}(x) = \underset{\text{H}}{f}(o_x) + \underset{\text{L}}{\Delta f}(\mathsf{fl}_{\mathsf{lo}}(o_x), \delta_x)$$

$$= f(o_x) \cdot (1 + \varepsilon_{\mathrm{hi}}) + \eta_{\mathrm{hi}} + \Delta f(o_x \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo}}, \delta_x) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)}.$$

Now, since f' is Lipschitz continuous, for any z, ϵ, and δ, it will hold by the mean value theorem that for some u between $z + \delta$ and $z + \delta + \epsilon$ and some v between z and $z + \epsilon$,

$$\begin{aligned}
|\Delta f(z + \epsilon, \delta) &- \Delta f(z, \delta)| \\
&= |f(z + \epsilon + \delta) - f(z + \epsilon) - f(z + \delta) + f(z)| \\
&= \left| \epsilon f'(u) - \epsilon f'(v) \right| \\
&\leq \epsilon \cdot L_f \cdot |u - v| \\
&\leq \epsilon \cdot L_f \cdot (|\delta| + |\epsilon|) \\
&= L_f \cdot |\delta| \cdot \left(|\epsilon| + O(\epsilon^2) \right).
\end{aligned}$$

So, applying this to our error analysis above,

$$\begin{aligned}
\underset{\text{C}}{f}(x) &= f(o_x) \cdot (1 + \varepsilon_{\mathrm{hi}}) + \eta_{\mathrm{hi}} + (\Delta f(o_x, \delta_x) + L_f \cdot \delta_x \cdot (o_x \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo}})) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= f(x) + f(o_x) \cdot \varepsilon_{\mathrm{hi}} + \eta_{\mathrm{hi}} + \Delta f(o_x, \delta_x) \cdot \varepsilon_{\mathrm{lo}} + L_f \cdot \delta_x \cdot (o_x \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)}.
\end{aligned}$$

By the mean value theorem again, for some z between o_x and $o_x + \delta_x$,

$$\Delta f(o_x, \delta_x) = f(o_x + \delta_x) - f(o_x) = f'(z) \cdot \delta_x = (f'(o_x) - f'(z)) \cdot \delta_x + f'(o_x) \cdot \delta_x,$$

and since $|f'(o_x) - f'(z)| \leq L_f \cdot |\delta_x|$, it follows that

$$\left| \Delta f(o_x, \delta_x) - f'(o_x) \cdot \delta_x \right| \leq L_f \cdot \delta_x^2.$$

Substituting this into our bound gives us

$$\underset{\mathrm{C}}{f}(x) = f(x) + f(o_x) \cdot \varepsilon_{\mathrm{hi}} + \eta_{\mathrm{hi}} + f'(o_x) \cdot \delta_x \cdot \varepsilon_{\mathrm{lo}} + L_f \cdot \delta_x^2 \cdot \varepsilon_{\mathrm{lo}}$$

$$+ L_f \cdot \delta_x \cdot (o_x \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)}$$

$$= f(x) + f(o_x) \cdot \varepsilon_{\mathrm{hi}} + \eta_{\mathrm{hi}} + f'(o_x) \cdot \delta_x \cdot \varepsilon_{\mathrm{lo}} + L_f \cdot \delta_x \cdot (o_x \cdot \varepsilon_{\mathrm{lo}} + \delta_x \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)}.$$

So, the error will be bounded by

$$\left| \underset{\mathrm{C}}{f}(x) - f(x) \right| = f(o_x) \cdot \varepsilon_{\mathrm{hi}} + \eta_{\mathrm{hi}} + f'(o_x) \cdot \delta_x \cdot \varepsilon_{\mathrm{lo}} + L_f \cdot \delta_x \cdot (o_x \cdot \varepsilon_{\mathrm{lo}} + \delta_x \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)}.$$

The magnitude of the delta of this operation will be bounded by

$$\left| \underset{\mathrm{C}}{\delta_{f(x)}} \right| = \left| \underset{\mathrm{L}}{\Delta} f(\mathrm{fl}_{\mathrm{lo}}(o_x), \delta_x) \right|$$

$$= \left| f'(o_x) \cdot \delta_x \cdot \varepsilon_{\mathrm{lo}} + L_f \cdot \delta_x \cdot (o_x \cdot \varepsilon_{\mathrm{lo}} + \delta_x \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \right|$$

$$= |\delta_x| \cdot \left| f'(o_x) \cdot \varepsilon_{\mathrm{lo}} + L_f \cdot \delta_x \cdot \varepsilon_{\mathrm{lo}} + L_f \cdot (o_x \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo}}) \right| + \left| \eta_{\mathrm{lo\text{-}bias}(B)} \right|.$$

Just as before, this gets small as δ gets small. Now, of course this is only one possible way of doing bit centering for general functions. If the function Δf is not available for computing in the low-precision representation, other arrangements must be made. One possibility is to follow the technique we used for exp by computing Δf via other functions that do exist in our low-precision floating point library. For example, for $f(x) = \tanh(x)$, we could decompose Δf with

$$\Delta f(x, y) = \tanh(x + y) - \tanh(x)$$

$$= \frac{\tanh(x) + \tanh(y)}{1 + \tanh(x) \tanh(y)} - \tanh(x)$$

$$= \frac{\tanh(y) - \tanh^2(x) \tanh(y)}{1 + \tanh(x) \tanh(y)}$$

$$= \frac{\mathrm{sech}^2(x)}{1 + \tanh(x) \tanh(y)} \cdot \tanh(y)$$

It is readily apparent that computing this expression in low-precision will result in an error that is low as y becomes small, which is what we need.

Another way of computing a bit-centered function is to approximate the function $\Delta f(o_x, \delta)$ for each o_x with a polynomial in δ (e.g. via a Chebyshev polynomial interpolation) on the host device, and then proceed to compute that polynomial in low-precision arithmetic whenever $\Delta f(o_x, \delta)$ is needed on the accelerator. When approximating Δf in this way, care must be taken to ensure that the numerical error from computing $\Delta f(o_x, \delta)$ becomes small as δ becomes small.

Although the bounds we got for the bit-centered computations in the previous section are all different, they all have one important property in common: both the delta of the outputs and the error become small as the magnitude of the δ of the inputs become small. We can formalize this notion into a bit-centered analogue of the fundamental axiom of floating point arithmetic.

Definition 6 (Fundamental axiom of bit-centered arithmetic) *Let f be a function or operator we are computing using bit centering, and let $\underset{C}{f}$ denote the result of that bit-centered computation. Suppose that f takes m arguments $f(x_1, x_2, \ldots, x_m)$. Then we say that f satisfies the* fundamental axiom of bit-centered arithmetic *if there exist continuous non-negative functions A_{hi}, A_{lo}, C_{hi}, C_{lo}, $C_{\mathrm{lo\text{-}bias}}$, R_{delta}, and $R_{\mathrm{lo\text{-}bias}}$, such that for any bit-centered inputs $x_1, x_2, \ldots, x_m$ with $x_i = o_{x,i} + \delta_{x,i}$, if overflow cannot occur (but underflow possibly can), then the error of the bit-centered computation is bounded by*

$$\left| \underset{C}{f}(x_1, x_2, \ldots, x_m) - f(x_1, x_2, \ldots, x_m) \right| \leq A_{\mathrm{hi}} \cdot \varepsilon_{\mathrm{machine\text{-}hi}} + C_{\mathrm{hi}} \cdot \eta_{\mathrm{machine\text{-}hi}}$$

$$+ A_{\mathrm{lo}} \cdot \left(\sum_{i=1}^{m} |\delta_{x,i}| \right) \cdot \varepsilon_{\mathrm{machine\text{-}lo}}$$

$$+ C_{\mathrm{lo}} \cdot \left(\sum_{i=1}^{m} |\delta_{x,i}| \right) \cdot \eta_{\mathrm{machine\text{-}lo}}$$

$$+ C_{\mathrm{lo\text{-}bias}} \cdot \eta_{\mathrm{machine\text{-}lo\text{-}bias}(B)}$$

and the magnitude of the delta of the bit-centered computation is bounded by

$$\left| \underset{C}{\delta_f}(\delta_{x,1}, \ldots, \delta_{x,m}, \underset{C}{\mathrm{cache}}_f(o_{x,1}, \ldots, o_{x,m})) \right| \leq R_{\mathrm{delta}} \cdot \left(\sum_{i=1}^{m} |\delta_{x,i}| \right) + R_{\mathrm{lo\text{-}bias}} \cdot \eta_{\mathrm{machine\text{-}lo\text{-}bias}(B)}$$

where each of the functions A_{hi} et cetera is a function of x_1, $o_{x,1}$, x_2, $o_{x,2}$, $\ldots$, x_m, $o_{x,m}$, $\varepsilon_{\mathrm{machine\text{-}hi}}$, $\eta_{\mathrm{machine\text{-}hi}}$, $\varepsilon_{\mathrm{machine\text{-}lo}}$, $\eta_{\mathrm{machine\text{-}lo}}$, and $\eta_{\mathrm{machine\text{-}lo\text{-}bias}(B)}$, and each is continuous over all values of x_i and $o_{x,i}$ in $\mathbb{R}$ and over all non-negative values of the various ε and η terms.

We can readily see that all the operations we analyzed above satisfy this axiom. For example, addition satisfies this axiom with parameters

$$A_{\mathrm{hi}}(x, o_x, y, o_y, \ldots) = |o_x + o_x|$$
$$C_{\mathrm{hi}}(x, o_x, y, o_y, \ldots) = 1$$
$$A_{\mathrm{lo}}(x, o_x, y, o_y, \ldots) = 1$$
$$C_{\mathrm{lo}}(x, o_x, y, o_y, \ldots) = 0$$
$$C_{\mathrm{lo\text{-}bias}}(x, o_x, y, o_y, \ldots) = 1$$
$$R(x, o_x, y, o_y, \ldots) = |(x - o_x) + (y - o_y)|$$
$$S(x, o_x, y, o_y, \ldots) = 1.$$

Definition 6 is particularly useful because of the following property: compositions of continuous functions that satisfy the axiom also satisfy the axiom.

Theorem 7 (Bit-centered arithmetic for composite functions) *Compositions of continuous functions that satisfy the fundamental axiom of bit-centered arithmetic also satisfy that axiom. That is, if $f : \mathbb{R}^m \to \mathbb{R}$ and $g_1, g_2, \ldots, g_m : \mathbb{R}^n \to \mathbb{R}$ are continuous functions all*

of which have implementations that satisfy the fundamental axiom of bit-centered arithmetic, then so does the function $h : \mathbb{R}^n \to \mathbb{R}$ defined by

$$h(x_1, \ldots, x_n) = f(g_1(x_1, \ldots, x_n), g_2(x_1, \ldots, x_n), \ldots, g_m(x_1, \ldots, x_n))$$

and with bit-centered implementation

$$\underset{C}{h}(x_1, \ldots, x_n) = \underset{C}{f}(\underset{C}{g_1}(x_1, \ldots, x_n), \underset{C}{g_2}(x_1, \ldots, x_n), \ldots, \underset{C}{g_m}(x_1, \ldots, x_n)).$$

We defer the proof of this theorem to the appendix. As a result of this theorem, we now have guarantees on the error of bit-centered arithmetic for *any* function that is a composition of primitives that we can prove satisfy the fundamental axiom. For example, this applies to matrix multiplication, which is just a composition of addition and multiplication.

6. HALP for General Loss Functions

Now that we have a way of computing bit centering for arbitrary functions, we can finally state a universal version of HALP. Just like our HALP for linear regression (Algorithm 6), the more general HALP uses bit centering to reduce the error from low-precision arithmetic. We describe HALP explicitly in Algorithm 7.

In prose, HALP does the following, in the outer loop. First, just like regular SVRG, it computes the full gradient at the outer iterate $\tilde{w}$ in high precision. Next, just as we did for the linear regression version of HALP, it computes the bias that it is going to use for the dynamic bias adjustment, using the same formula as the linear regression case. Next, it does all the high-precision pre-computation it will need to do bit centering later for each of the objective function components. This involves: (1) computing the offsets for each intermediate value in the computation; (2) computing any values that depend on those offsets and are needed for the later computation of the deltas; (3) converting those values to the low-precision format; and (4) storing them somewhere they can be accessed later. Also, at the end of the outer loop, we add a check to validate that the loss is decreasing at each epoch: if it ever increases, we reset the outer iterate to the value at the previous iteration (this results in a decrease in the loss since its value at the current iteratio must be greater than its value at the previous one).

In each inner loop iteration, HALP samples an example i at random, just as standard SVRG does. Next, it loads the pre-computed values it needs to compute ∇f_i using bit centering: these are exactly those values that were computed and stored earlier in the outer loop. Using these values, it computes the delta for $\nabla f_i(\tilde{w} + \delta_{w,k,t})$ using bit centering. Finally, it updates the delta of the model using this gradient it computed using bit centering.

As long as the operations we use to compute the gradient samples satisfy the fundamental axiom of bit-centered arithmetic, the numerical error of the inner loop will become small as $\tilde{g}_k$ becomes small, just as it did for the linear regression case. Just like in that simpler case, here we can prove HALP converges at a linear rate to arbitrarily accurate solutions. But first, we will need to state and justify some assumptions.

Assumption 8 *Assume that there exists global constants a_{lo}, c_{lo}, and $c_{\mathrm{lo\text{-}bias}}$ such that the bit-centered computation of the gradient samples, for any $\tilde{w}_k$ and any $\delta_{w,k,t}$ that might appear*

Algorithm 7 HALP: Bit-centered, dynamic-bias-adjusted SVRG for general loss functions

given: N low-precision training examples x_i and labels y_i, number of epochs K, epoch length T, low-precision step size α, initial low-precision iterate $\tilde{w}_1$, and bias factor $\zeta > 0$.

for $k = 1$ **to** K **do**

> ▷ *high-precision precomputation on host*
>
> $\tilde{g}_k \leftarrow \underset{\text{H}}{\nabla} f(\mathsf{fl}_{\text{hi}}(\tilde{w}_k)) = \left(\underset{i=1}{\overset{N}{\underset{\text{H}}{\sum}}} \underset{\text{H}}{\nabla} f_i(\tilde{w}_k) \right) \underset{\text{H}}{\oslash} N$
>
> **compute bias:** $B \leftarrow \left\lfloor \underset{\text{H}}{\log_2} \left(\zeta \cdot \|\tilde{g}_k\|_{\text{H}} \right) \right\rfloor$
>
> $\tilde{h}_k \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(\tilde{g}_k)$
>
> $\tilde{s}_k \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(2 \underset{\text{H}}{\odot} \|\tilde{g}_k\|_{\text{H}} \underset{\text{H}}{\oslash} \mu)$
>
> **for** $i = 1$ **to** N **do**
>
> > ▷ **precompute in high-precision and store in low-precision** all numbers necessary to compute $\nabla f_i(w)$ using bit centering with offset of $o_w = \tilde{w}_k$
> >
> > $\sigma_{k,i} \leftarrow \mathrm{cache}_{\underset{\text{C}}{\nabla f_i}}(\tilde{w}_k)$
>
> **end for**

$\delta_{w,k,0} \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(0)$

> ▷ *low-precision computation on accelerator*
>
> **for** $t = 1$ **to** T **do**
>
> > **sample** i uniformly from $\{1, \ldots, N\}$
> >
> > **load** $\sigma_{k,i}$
> >
> > **compute using bit centering:** $\delta_{v,k,t-1} \leftarrow \delta_{\underset{\text{C}}{\nabla f_i}}(\delta_{w,k,t}, \sigma_{k,i})$
> >
> > **update model:** $\delta_{w,k,t} \leftarrow \delta_{w,k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\delta_{v,k,t} \underset{\text{L}}{\oplus} \tilde{h}_k \right)$
> >
> > **if** $\|\delta_{w,k,t}\|_{\text{L}} > \tilde{s}_k$ **then**
> >
> > > $\delta_{w,k,t} \leftarrow \mathsf{fl}_{\text{lo-bias}(B)}(0)$
> >
> > **end if**
>
> **end for**

option I: set $\tilde{w}_{k+1} \leftarrow \tilde{w}_k \underset{\text{H}}{\oplus} \mathsf{fl}_{\text{hi-from}(B)}(\delta_{w,k,T})$

option II: sample t uniformly from $\{0, \ldots, T-1\}$,

  then set $\tilde{w}_{k+1} \leftarrow \tilde{w}_k \underset{\text{H}}{\oplus} \mathsf{fl}_{\text{hi-from}(B)}(\delta_{w,k,t})$

if $\underset{\text{H}}{f}(\tilde{w}_{k+1}) \leq \underset{\text{H}}{f}(\tilde{w}_k)$ **then**

  roll back epoch: $\tilde{w}_{k+1} \leftarrow \tilde{w}_k$

end if

end for

return $\tilde{w}_{K+1}$

---

in the course of the algorithm, has a delta with error bounded by

$$\left\| \underset{C}{\nabla f_i}((\tilde{w}_k, \delta_{w,k,t})) - \nabla f_i(\tilde{w}_k + \delta_{w,k,t}) \right\|$$

$$= a_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \varepsilon_{\text{lo}} + c_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \eta_{\text{lo}} + c_{\text{lo-bias}} \cdot \eta_{\text{lo-bias}(B)} \cdot$$

This assumption can be justified directly from the fundamental axiom of bit-centered arithmetic. To see why, first notice that $\tilde{w}_k$ and $w_{k,t}$ that we could encounter during computation must live in some compact set. This holds because the loss of the outer iterates is non-increasing, so for every k it will hold that

$$\tilde{w}_k \in \{w | f(w) \le f(\tilde{w}_1)\}$$

which is compact because f is strongly convex. As a consequence, we also know that $w_{k,t}$ must live in a compact set as well, since its distance from $\tilde{w}_k$ (i.e. the magnitude of $\delta_{w,k,t}$) is always bounded from above by $\tilde{s}_k$. Since the $\tilde{w}_k$ and $w_{k,t}$ all must live in a compact set, by the extreme value theorem there must exist upper bounds for the continuous bounding functions A_{lo}, C_{lo}, and $C_{\text{lo-bias}}$ from the fundamental axiom, bounds which hold across all examples i and all possible values of $\tilde{w}_k$ and $w_{k,t}$; since these upper bounds exist, then a_{lo}, c_{lo}, and $c_{\text{lo-bias}}$ that satisfy this assumption must exist as well. (Note that we are ignoring the high-precision part of the bounds from A_{hi} and C_{hi} because we are assuming from Assumption 2 that there is no high-precision error.)

Assumption 9 *Assume that there exists global constants r_{delta} and $r_{\text{lo-bias}}$ such that every number z that could possibly be computed in the inner loop of Algorithm 7 is bounded in magnitude by*

$$|z| \le r_{\text{delta}} \cdot \|\delta_{w,k,t}\| + r_{\text{lo-bias}} \cdot \eta_{\text{lo-bias}(B)}$$

We can justify this assumption via a similar appeal to the fundamental axion of bit-centered arithmetic. Since $\tilde{w}_k$ and $w_{k,t}$ must live in a compact set, by the extreme value theorem there must exist upper bounds for the bounding functions R_{delta} and $R_{\text{lo-bias}}$ from the fundamental axiom, which hold across all numbers that could be computed in the inner loop and all possible values of $\tilde{w}_k$ and $w_{k,t}$. Using these upper bounds, we can derive r_{delta} and $r_{\text{lo-bias}}$ that satisfy Assumption 9.

Using these assumptions, we can prove the following theorem, which guarantees linear-rate convergence of HALP to solutions of arbitrarily low error.

Theorem 10 *Suppose that we run HALP (Algorithm 7) under Assumptions 1 (strong convexity and Lipschitz continuity), 2 (zero-error high-precision arithmetic), 8, and 9, using option II for the epoch update. Suppose that our underflow threshold $\eta_{\text{machine-lo}}$ is small enough that*

$$\left(c_{\text{lo}} \cdot \frac{2}{\mu} + c_{\text{lo-bias}} \cdot \frac{\zeta}{2} + (4L+1)\sqrt{d} \cdot \zeta \right) \cdot \eta_{\text{machine-lo}} \le \varepsilon_{\text{machine-lo}},$$

and that our overflow threshold is large enough that

$$r_{\text{delta}} \cdot \frac{4}{\zeta\mu} + r_{\text{lo-bias}} \cdot \eta_{\text{machine-lo}} \le M_{\text{machine-lo}}.$$

37

For any constant $0 < \gamma < 1$, if we set our step size and epoch lengths to be

$$\alpha = \frac{\gamma}{4L(1+\gamma)}, \qquad T \geq \frac{8\kappa(1+\gamma)}{\gamma^2},$$

then the iterates of HALP will satisfy

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\gamma + \frac{48\kappa(8L + a_{\text{lo}})}{\gamma L} \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right],$$

and overflow cannot possibly occur.

This theorem shows that as long as the underflow threshold and overflow threshold satisfy certain inequalities (which requires some number of exponent bits to do so), and as long as the machine epsilon is sufficiently small, HALP can converge down to arbitrarily accurate solutions, just like exact-arithmetic SVRG can. Note that the dependence on the condition number that we observed earlier for linear-regression HALP is still present here. In order for the algorithm to converge at a linear rate, we will need

$$\frac{48\kappa(8L + a_{\text{lo}})}{\gamma L} \cdot \varepsilon_{\text{machine-lo}} \ll 1,$$

which happens when

$$\varepsilon_{\text{machine-lo}} = O\left(\min\left(\frac{1}{\kappa}, \frac{L}{\kappa \cdot a_{\text{lo}}}\right)\right).$$

So we certainly need the machine epsilon to be small relative to κ^{-1}, which suggests that we need more bits as the conditioning becomes worse. Specifically, since the machine epsilon is typically about two raised to the negative of the number of mantissa bits, we will need roughly

$$b = O(\log(\kappa))$$

bits of precision to enable the linear convergence rate in Theorem 10. Note that we may also need additional bits to handle the numerical imprecision that comes from computing the gradient samples, which is represented by the a_{lo} term.

7. Experiments

The goal of our experiments is to confirm that HALP can achieve high-accuracy solutions with the majority of the computation being done in low precision computation while some high precision computation is performed infrequently to recenter the solution. In this section we not only show that HALP can lead to high-accuracy solutions on a convex problem in Section 7.1 (multi-class logistic regression), but also that it can lead to high-accuracy solutions on non-convex neural network applications in Section 7.2 (deep convolutional neural networks and a recurrent neural network).

Low-Precision Format Due to recent hardware trends (Jouppi et al., 2017; Micikevicius, 2017), we focus on 16-bit training as this is the popular floating point precision available in recent (and projected future) hardware generations (Corporation, 2018). Unfortunately there is not a consensus on the data type for 16-bit floating point numbers, so we built two simulators to validate our claims: one in C++ that uses the `bfloat16` floating point format and one in PyTorch that uses `float16` floating point format.
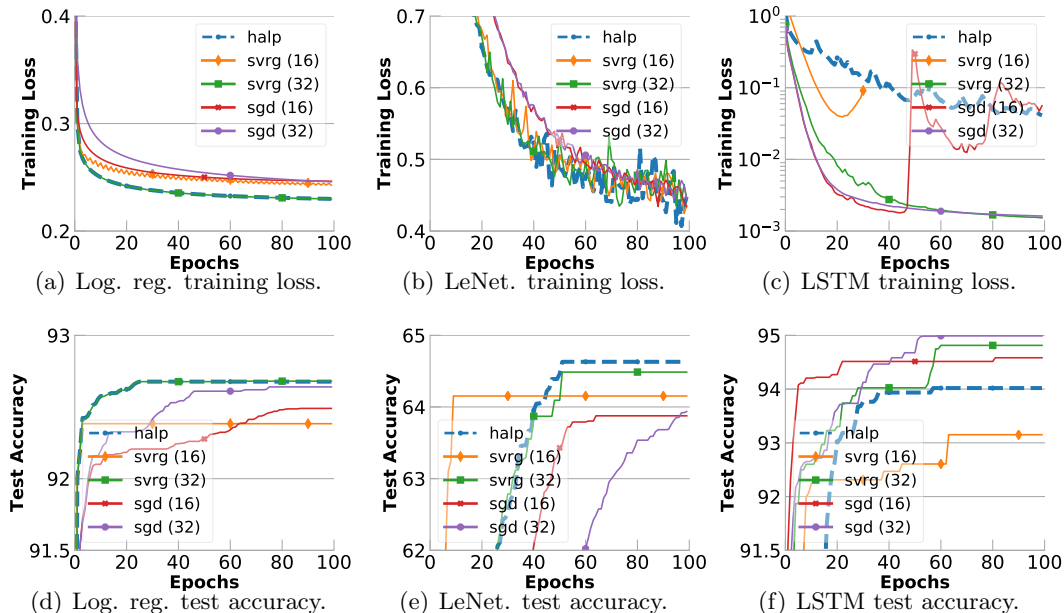
Figure 7: Training loss and test accuracy for logistic regression on MNIST, LeNet on CIFAR10 and LSTM on Penn Treebank.

Setup To test HALP, we evaluate it against the performance of four baseline training algorithms (32-bit SGD, 16-bit SGD, 32-bit SVRG, and 16-bit SVRG). Therefore, low-precision (LP) is synonymous with 16-bit throughout the entirety of this section. For hyperparameter selection, we performed grid search using constant learning rates throughout training. For more details on the hyperparameters we swept, see Appendix A. We present the training loss (averaged over each epoch) and the test accuracy as our evaluation metrics. To ensure a fair comparison, all training algorithms on each task are trained under the same search grid. Within the grid, for each training algorithm, we pick the configuration that generates the best evaluation metric; the configurations for reporting training loss and test accuracy are picked independently. All configurations are run for 100 epochs and with a recentering, or computation of the full gradient for all SVRG-based algorithms, every epoch. To assure statistically meaningful results, our reported metrics are averaged from three runs using different random seeds. For more experimental details, see Appendix A.

7.1 Multi-Class Logistic Regression Results

In Figures 7(a) and 7(d) we show that HALP, using `float16` numbers, strictly outperforms (in terms of test accuracy and training loss) LP-SGD and LP-SVRG on logistic regression on the MNIST (LeCun, 1998) digit classification task. HALP outperforms both LP-SGD and LP-SVRG here due to its lower-magnitude quantization noise from bit centering. HALP outperforms full-precision SGD here due to its lower-magnitude gradient noise. All experiments in this section were run with a minibatch size of 100.

7.2 Deep Learning Results

We also evaluate HALP on two non-convex applications for which the theory provides no guarantees. Surprisingly, even on these applications we show that HALP can be useful as training algorithm. We present HALP results on both a convolutional neural network (CNN) and a recurrent neural network (LSTM).

CNN discussion. In Figure 7(b) and Figure 7(e), we show that HALP can outperform low-precision algorithms (LP-SGD and LP-SVRG) and match full-precision algorithms (SGD and SVRG) in terms of both training loss and test accuracy. To test HALP's performance on a CNN, we train a 5-layer LeNet (LeCun et al., 1998) model on the CIFAR10 image classification dataset (Krizhevsky et al., 2014). Specifically, we use the reference model configuration from Kuang (2018): 2 convolutional layers with 6 and 16 5x5 filters, followed by 2 linear layers with 120 and 84 hidden units. We train the model using minibatch size 128 and with a ℓ_2 regularization of 0.0005.

LSTM discussion. In Figures 7(c) and 7(f) we show that HALP can outperform LP-SVRG while approaching the performance of full-precision SVRG on a part of speech tagging application. Unfortunately, SVRG-based algorithms suffer a degradation when compared to all SGD-based algorithms on this application. As a result, HALP similarly performs worse than both LP-SGD and SGD while matching the performance of full-precision SVRG. One possible way to mitigate this effect would be to run the recentering step more often than once every epoch. The part-of-speech tagging application we run is over the popular Penn Treebank dataset Marcus et al. (1993). The model that we use is a simple architecture consisting of an embedding layer, LSTM layer, and classifier at the end of the network. The embedding layer is not pretrained, and is trained entirely during the duration of our run. We trained with a batch size of 1 in this section. Note that the LP-SVRG loss line shown in Figure 7(c) ends early because the loss values become infinity despite that this hyperparameter setting achieved (earlier on) the lowest loss out of all tested configurations.

8. Conclusion

In this paper we presented HALP, a new SGD variant that is able to theoretically converge at a linear rate while primarily using low-precision computation. HALP leverages SVRG to reduce noise from gradient variance, and introduces *bit centering* to reduce noise from quantization. We proved that HALP converges arbitrarily close to the global optimum on strongly convex problems and validated our results extensively on convex (linear and logistic regression) and non-convex (CNN and LSTM) applications.

model	hyperparam	grid
Log. reg.	α	$\{0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1.0\}$
	μ	$\{0.0, 0.9\}$
	ℓ_2	$\{$1e-5, 5e-5, 1e-4, 5e-4, 1e-3, 5e-3, 1e-2, 5e-2, 1e-1$\}$
LeNet	α	$\{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5\}$
	μ	$\{0.0, 0.9\}$
	ℓ_2	$\{$5e-4$\}$
LSTM	α	$\{0.001, 0.01, 0.05, 0.01, 1.0, 2.0, 4.0\}$
	μ	$\{0.0, 0.9\}$
	ℓ_2	$\{0,$ 1e-5, 1e-4, 1e-3$\}$

Table 3: The learning rate α, momentum μ, regularization strength ℓ_2 values for the grid search in Section 7. Note we did not additionally search over regularizer strength for LeNet; instead, we use the recommended value by Kuang (2018).

model	hyperparam	SGD (32)	SVRG (32)	LP-SGD (16)	LP-SVRG (16)	HALP (16)
Log. reg.	α	0.01	0.05	0.01	0.05	0.05
	μ	0.9	0.9	0.9	0.9	0.9
	ℓ_2	1e-5	1e-5	1e-5	1e-5	1e-5
LeNet	α	0.0005	0.001	0.001	0.005	0.001
	μ	0.0	0.9	0.9	0.9	0.9
	ℓ_2	5e-4	5e-4	5e-4	5e-4	5e-4
LSTM	α	2.0	2.0	2.0	0.1	0.05
	μ	0	0.0	0	0.9	0.9
	ℓ_2	1e-5	1e-5	1e-4	1e-4	1e-3

Table 4: The learning rate α, momentum μ, regularization strength ℓ_2 achieving the highest test accuracy in Section 7.

Appendix A. Extended Evaluation

A.1 Detailed Experiment Setup for Section 7

In Section 7, we perform grid search on learning rate and momentum. From the grid search, we pick the configuration achieving the lowest training loss (averaged within epoch) and highest test accuracy respectively for each training algorithm. In Table 3, we present the learning rate and momentum values we used for the grid search. In addition, we demonstrate the configurations achieving lowest training loss in Table 5 and the ones achieving highest test accuracy in Table 4.

Appendix B. Details of Results

In this appendix, we present some details that, for brevity, were not included in the main body of the manuscript.

model	hyperparam.	SGD (32)	SVRG (32)	LP-SGD (16)	LP-SVRG (16)	HALP (16)
Log. reg.	α	0.1	0.05	0.5	0.5	0.05
	μ	0.0	0.9	0.0	0.0	0.9
	ℓ_2	1e-5	1e-5	1e-5	1e-5	1e-5
LeNet	α	0.05	0.005	0.05	0.005	0.005
	ℓ_2	5e-4	5e-4	5e-4	5e-4	5e-4
	μ	0.0	0.9	0.0	0.9	0.9
LSTM	α	0.1	0.05	1.0	1e-3	0.1
	μ	0.9	0.9	0	0.9	0.9
	ℓ_2	0	0	1e-5	0	1e-4

Table 5: The learning rate α and momentum μ, regularization strength ℓ_2 achieving the lowest training loss in Section 7.

Algorithm 8 SVRG: Stochastic Variance-Reduced Gradient
---
given: N loss gradients ∇f_i, number of epochs K, epoch length T, step size α, and initial iterate $\tilde{w}_1$.
for $k = 1$ **to** K **do**
$\quad \tilde{g}_k \leftarrow \nabla f(\tilde{w}_k) = \frac{1}{N} \sum_{i=1}^{N} \nabla f_i(\tilde{w}_k)$
$\quad w_{k,0} \leftarrow \tilde{w}_k$
$\quad$ **for** $t = 1$ **to** T **do**
$\quad\quad$ **sample** i uniformly from $\{1, \ldots, N\}$
$\quad\quad w_{k,t} \leftarrow w_{k,t-1} - \alpha \left(\nabla f_i(w_{k,t-1}) - \nabla f_i(\tilde{w}_k) + \tilde{g}_k \right)$
$\quad$ **end for**
$\quad$ **option I:** set $\tilde{w}_{k+1} \leftarrow w_{k,T}$
$\quad$ **option II: sample** t uniformly from $\{0, \ldots, T-1\}$, then set $\tilde{w}_{k+1} \leftarrow w_{k,t}$
end for
return $\tilde{w}_{K+1}$

B.1 SVRG

For completeness, we present the basic SVRG algorithm for minimizing an objective, Algorithm 8. Compared with standard stochastic gradient descent, SVRG is able to converge at a linear rate because it periodically uses full gradients $\tilde{g}_k$ to reduce the variance of its stochastic gradient estimators. Note that the two outer-loop update options come from the paper that originally proposed SVRG, Johnson and Zhang (2013). In this and subsequent work Harikandeh et al. (2015), it has been standard to use option II for the theoretical analysis (as it simplifies the derivation) while using option I for all empirical experiments. We will continue to do this for all the SVRG variants we introduce here.

B.2 Notation

When we use low-precision computation with a biased exponent, we are still effectively computing with the same low-precision format: it will have the same number of bits and use the same base computational operations, and only the actual number represented will be different. To be consistent with our previous notation, we let $\mathsf{fl}_{\text{lo-bias}(B)}(x)$ denote the result

of converting x into a biased low-precision format with extra exponent bias B. Note that while this is a different format in terms of what the bits represent, many of the computations in the biased low-precision format can be done with the same hardware and instructions that can be used to compute in the original low-precision format. For example, it can be readily seen that an add or subtract operation in the biased format (for two numbers with the same extra bias B) involves exactly the same binary computation as in the original format. Furthermore, the product of two biased floating point numbers, one with extra bias B_1 and the other with extra bias B_2, can be done with the same binary computation as in the original format, but results in a number with extra bias $B_1 + B_2$. As a special case of this, the product of a biased floating point number (with extra bias B) and an ordinary floating point number can be computed with an ordinary multiply and results in a number with extra bias B. Similar results hold for division. Importantly, the bias values that result from all these computations are fixed *a priori* once the bias of the inputs is set. As a result, we can precompute all bias values in the outer loop and use them at runtime only when necessary. The only computations for which the bias value must be explicitly taken into account at runtime are: (1) conversions to and from the high-precision format; (2) conversions between biased and standard low-precision formats; and (3) To help with clarity, we will indicate these conversions explicitly in our algorithm statements with $\mathsf{fl}_{\text{hi-from}(B)}(\cdot)$, $\mathsf{fl}_{\text{lo-from}(B)}(\cdot)$, and $\mathsf{fl}_{\text{lo-bias}(B_1)\text{-from}(B_2)}(\cdot)$, respectively.

B.3 Details of Bit-Centered Operations

In this section, we provide additional details of the bit-centered operations described in the main body of the manuscript.

Subtraction   The analysis is the same as for addition. We define

$$x \underset{\text{C}}{\ominus} y \overset{\text{def}}{=} \left(\left(o_x \underset{\text{H}}{\ominus} o_y \right), \left(\delta_x \underset{\text{L}}{\ominus} \delta_y \right) \right).$$

Explicitly, this decomposes into

$$o_{\underset{\text{C}}{\ominus}}(o_x, o_y) \overset{\text{def}}{=} o_x \underset{\text{H}}{\ominus} o_y$$

$$\text{cache}_{\underset{\text{C}}{\ominus}}(o_x, o_y) \overset{\text{def}}{=} \emptyset$$

$$\delta_{\underset{\text{C}}{\ominus}}(\delta_x, \delta_y, \text{cache}) \overset{\text{def}}{=} \delta_x \underset{\text{L}}{\ominus} \delta_y.$$

This will have error

$$\left| x \underset{\text{C}}{\ominus} y - (x - y) \right| \le |x - y| \cdot \varepsilon_{\text{hi}} + |\delta_x - \delta_y| \cdot (\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + \eta_{\text{lo-bias}(B)} + \eta_{\text{hi}}.$$

The magnitude of the delta that results from this operation will be bounded by

$$\left| \delta_{x \underset{\text{C}}{\ominus} y} \right| = \left| \delta_x \underset{\text{L}}{\ominus} \delta_y \right| \le |\delta_x - \delta_y| \cdot \varepsilon_{\text{lo}} + \eta_{\text{lo-bias}(B)}.$$

Division Suppose that we want to do division using bit centering, to compute $x \underset{\text{C}}{\oslash} y$ where $x = o_x + \delta_x$ and $y = o_y + \delta_y$. The natural way to do this is

$$x \underset{\text{C}}{\oslash} y \stackrel{\text{def}}{=} \left((o_x \underset{\text{H}}{\oslash} o_y), \left(\delta_x \underset{\text{L}}{\ominus} (\mathsf{fl}_{\text{lo}}(o_x \underset{\text{H}}{\oslash} o_y) \underset{\text{L}}{\odot} \delta_y) \right) \underset{\text{L}}{\oslash} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{\text{L}}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right) \right).$$

Explicitly, this decomposes into

$$o_{\underset{\text{C}}{\oslash}}(o_x, o_y) \stackrel{\text{def}}{=} o_x \underset{\text{H}}{\oslash} o_y$$

$$\text{cache}_{\underset{\text{C}}{\oslash}}(o_x, o_y) \stackrel{\text{def}}{=} (\mathsf{fl}_{\text{lo}}(o_x \underset{\text{H}}{\oslash} o_y), \mathsf{fl}_{\text{lo}}(o_y))$$

$$\delta_{\underset{\text{C}}{\odot}}(\delta_x, \delta_y, (\sigma_{x/y}, \sigma_y)) \stackrel{\text{def}}{=} \left(\delta_x \underset{\text{L}}{\ominus} (\sigma_{x/y} \underset{\text{L}}{\odot} \delta_y) \right) \underset{\text{L}}{\oslash} \left(\sigma_y \underset{\text{L}}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right).$$

From this, we will have

$$x \underset{\text{C}}{\oslash} y = (o_x \underset{\text{H}}{\oslash} o_y) + \left(\delta_x \underset{\text{L}}{\ominus} (\mathsf{fl}_{\text{lo}}(o_x \underset{\text{H}}{\oslash} o_y) \underset{\text{L}}{\odot} \delta_y) \right) \underset{\text{L}}{\oslash} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{\text{L}}{\oplus} \delta_y \right)$$

$$= (o_x/o_y) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}}$$

$$+ \Big(\big(\delta_x - ((((o_x/o_y) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}}) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}}) \cdot \delta_y)$$

$$\cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)} \big) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)} \Big)$$

$$/ \, ((o_y \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}} + \delta_y \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}}) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo}}) \cdot (1 + \varepsilon_{\text{lo}}) + \eta_{\text{lo-bias}(B)}$$

$$= (o_x/o_y) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{hi}}$$

$$+ \Big(\delta_x \cdot (1 + 2\varepsilon_{\text{lo}}) - (o_x/o_y) \cdot \delta_y \cdot (1 + 4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + (\eta_{\text{lo}} + \eta_{\text{hi}}) \cdot \delta_y + 2\eta_{\text{lo-bias}(B)} \Big)$$

$$/ \, (o_y \cdot (1 + 2\varepsilon_{\text{lo}}) + \delta_y \cdot (1 + 2\varepsilon_{\text{lo}}) + 3\eta_{\text{lo}}) + \eta_{\text{lo-bias}(B)}$$

$$= (o_x/o_y) \cdot (1 + \varepsilon_{\text{hi}}) + \eta_{\text{lo-bias}(B)} + \eta_{\text{hi}}$$

$$+ \frac{\delta_x \cdot (1 + 2\varepsilon_{\text{lo}}) - (o_x/o_y) \cdot \delta_y \cdot (1 + 4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + (\eta_{\text{lo}} + \eta_{\text{hi}}) \cdot \delta_y + 2\eta_{\text{lo-bias}(B)}}{o_y \cdot (1 + 2\varepsilon_{\text{lo}}) + \delta_y \cdot (1 + 2\varepsilon_{\text{lo}}) + 3\eta_{\text{lo}}}.$$

Next, we need some way to deal with the factors of ε_{lo} in the denominator. We can do this by noticing that, for small ϵ,

$$\frac{1}{x + \epsilon} = \frac{1}{x} \left(1 - \frac{\epsilon}{x} \right) + O(\epsilon^2).$$

44

This means that we can reduce our expression above to

$$x \oslash_{\mathrm{C}} y = (o_x/o_y) \cdot (1 + \varepsilon_{\mathrm{hi}}) + \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{hi}}$$
$$+ \frac{\delta_x \cdot (1 + 2\varepsilon_{\mathrm{lo}}) - (o_x/o_y) \cdot \delta_y \cdot (1 + 4\varepsilon_{\mathrm{lo}} + \varepsilon_{\mathrm{hi}}) + (\eta_{\mathrm{lo}} + \eta_{\mathrm{hi}}) \cdot \delta_y + 2\eta_{\mathrm{lo\text{-}bias}(B)}}{o_y + \delta_y}$$
$$\cdot \left(1 + \frac{o_y \cdot 2\varepsilon_{\mathrm{lo}} + \delta_y \cdot 2\varepsilon_{\mathrm{lo}} + 3\eta_{\mathrm{lo}}}{o_y + \delta_y}\right)$$
$$= (o_x/o_y) \cdot (1 + \varepsilon_{\mathrm{hi}}) + \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{hi}}$$
$$+ \frac{\delta_x - (o_x/o_y) \cdot \delta_y}{o_y + \delta_y} \cdot \left(1 + \frac{o_y \cdot 2\varepsilon_{\mathrm{lo}} + \delta_y \cdot 2\varepsilon_{\mathrm{lo}} + 3\eta_{\mathrm{lo}}}{o_y + \delta_y}\right)$$
$$+ \frac{\delta_x \cdot 2\varepsilon_{\mathrm{lo}} - (o_x/o_y) \cdot \delta_y \cdot (4\varepsilon_{\mathrm{lo}} + \varepsilon_{\mathrm{hi}}) + (\eta_{\mathrm{lo}} + \eta_{\mathrm{hi}}) \cdot \delta_y + 2\eta_{\mathrm{lo\text{-}bias}(B)}}{o_y + \delta_y}.$$

From here, notice that

$$\frac{\delta_x - (o_x/o_y) \cdot \delta_y}{o_y + \delta_y} = \frac{\delta_x \cdot o_y - o_x \cdot \delta_y}{(o_y + \delta_y) \cdot o_y} = \frac{(o_x + \delta_x) \cdot o_y - o_x \cdot (o_y + \delta_y)}{(o_y + \delta_y) \cdot o_y} = \frac{o_x + \delta_x}{o_y + \delta_y} - \frac{o_x}{o_y},$$

and so

$$x \oslash_{\mathrm{C}} y = (o_x/o_y) \cdot (1 + \varepsilon_{\mathrm{hi}}) + \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{hi}}$$
$$+ \left(\frac{x}{y} - \frac{o_x}{o_y}\right) \cdot \left(1 + \frac{o_y \cdot 2\varepsilon_{\mathrm{lo}} + \delta_y \cdot 2\varepsilon_{\mathrm{lo}} + 3\eta_{\mathrm{lo}}}{o_y + \delta_y}\right)$$
$$+ \frac{\delta_x \cdot 2\varepsilon_{\mathrm{lo}} - (o_x/o_y) \cdot \delta_y \cdot (4\varepsilon_{\mathrm{lo}} + \varepsilon_{\mathrm{hi}}) + (\eta_{\mathrm{lo}} + \eta_{\mathrm{hi}}) \cdot \delta_y + 2\eta_{\mathrm{lo\text{-}bias}(B)}}{o_y + \delta_y}$$
$$= \frac{x}{y} + \frac{o_x}{o_y} \cdot \varepsilon_{\mathrm{hi}} + \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{hi}}$$
$$+ \left(\frac{\delta_x - (o_x/o_y) \cdot \delta_y}{o_y + \delta_y}\right) \cdot \left(\frac{o_y \cdot 2\varepsilon_{\mathrm{lo}} + \delta_y \cdot 2\varepsilon_{\mathrm{lo}} + 3\eta_{\mathrm{lo}}}{o_y + \delta_y}\right)$$
$$+ \frac{\delta_x \cdot 2\varepsilon_{\mathrm{lo}} - (o_x/o_y) \cdot \delta_y \cdot (4\varepsilon_{\mathrm{lo}} + \varepsilon_{\mathrm{hi}}) + (\eta_{\mathrm{lo}} + \eta_{\mathrm{hi}}) \cdot \delta_y + 2\eta_{\mathrm{lo\text{-}bias}(B)}}{o_y + \delta_y}$$
$$= \frac{x}{y} + \frac{o_x}{o_y} \cdot \varepsilon_{\mathrm{hi}} + \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{hi}}$$
$$+ \left(\frac{\delta_x - (o_x/o_y) \cdot \delta_y}{y}\right) \cdot \left(\frac{y \cdot 2\varepsilon_{\mathrm{lo}} + \delta_y \cdot 4\varepsilon_{\mathrm{lo}} + 3\eta_{\mathrm{lo}}}{y}\right)$$
$$+ \frac{\delta_x \cdot 2\varepsilon_{\mathrm{lo}} - (o_x/o_y) \cdot \delta_y \cdot (4\varepsilon_{\mathrm{lo}} + \varepsilon_{\mathrm{hi}}) + (\eta_{\mathrm{lo}} + \eta_{\mathrm{hi}}) \cdot \delta_y + 2\eta_{\mathrm{lo\text{-}bias}(B)}}{y}$$
$$= \frac{x}{y} + \frac{o_x}{o_y} \cdot \varepsilon_{\mathrm{hi}} + \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{hi}}$$
$$+ \left(\frac{\delta_x - (o_x/o_y) \cdot \delta_y}{y^2}\right) \cdot (\delta_y \cdot 4\varepsilon_{\mathrm{lo}} + 3\eta_{\mathrm{lo}})$$
$$+ \frac{\delta_x \cdot 4\varepsilon_{\mathrm{lo}} + (o_x/o_y) \cdot \delta_y \cdot (6\varepsilon_{\mathrm{lo}} + \varepsilon_{\mathrm{hi}}) + (\eta_{\mathrm{lo}} + \eta_{\mathrm{hi}}) \cdot \delta_y + 2\eta_{\mathrm{lo\text{-}bias}(B)}}{y}.$$

Supposing that $\varepsilon_{\text{machine-hi}} \leq \varepsilon_{\text{machine-lo}}$ and similarly $\eta_{\text{machine-hi}} \leq \eta_{\text{machine-lo}}$, it follows that the error will be bounded by

$$\left| x \underset{\text{C}}{\oslash} y - x/y \right| \leq \left| \frac{o_x}{o_y} \right| \cdot \varepsilon_{\text{hi}} + \eta_{\text{lo-bias}(B)} + \eta_{\text{hi}}$$
$$+ \left(\frac{|\delta_x| + |o_x/o_y| \cdot |\delta_y|}{y^2} \right) \cdot (|y| \cdot 7\varepsilon_{\text{lo}} + |\delta_y| \cdot 4\varepsilon_{\text{lo}} + 3\eta_{\text{lo}})$$
$$+ \frac{|\delta_y| \cdot 2\eta_{\text{lo}} + 2\eta_{\text{lo-bias}(B)}}{|y|}.$$

Note that, as with our operations above, this expression becomes small (in the low-precision part of the error) as δ_x and δ_y become small. The magnitude of the delta of the result of this operator will be bounded by

$$\left| \delta_{x \underset{\text{C}}{\oslash} y} \right| = \left| \left(\delta_x \underset{\text{L}}{\ominus} (\mathsf{fl}_{\text{lo}}(o_x \underset{\text{H}}{\oslash} o_y) \underset{\text{L}}{\odot} \delta_y) \right) \underset{\text{L}}{\oslash} \left(\mathsf{fl}_{\text{lo}}(o_y) \underset{\text{L}}{\oplus} \mathsf{fl}_{\text{lo}}(\delta_y) \right) \right|$$
$$= \left| \frac{\delta_x - (o_x/o_y) \cdot \delta_y}{y} \cdot \left(1 + \frac{o_y \cdot 2\varepsilon_{\text{lo}} + \delta_y \cdot 2\varepsilon_{\text{lo}} + 3\eta_{\text{lo}}}{y} \right) \right.$$
$$\left. + \frac{\delta_x \cdot 2\varepsilon_{\text{lo}} - (o_x/o_y) \cdot \delta_y \cdot (4\varepsilon_{\text{lo}} + \varepsilon_{\text{hi}}) + (\eta_{\text{lo}} + \eta_{\text{hi}}) \cdot \delta_y + 2\eta_{\text{lo-bias}(B)}}{y} + \eta_{\text{lo-bias}(B)} \right|.$$

Notice that this expression becomes small as the input deltas become small, just as was the case for the previous operators we analyzed.

Appendix C. Proofs

In this appendix, we prove the main theorems presented in the paper. Before we prove the main theorems presented in the paper, we will prove the following lemmas, which will be useful later.

For completeness, we start by re-stating the proof of following lemma, which was presented as equation (8) in Johnson and Zhang (2013).

Lemma 11 *Under the standard condition of Lipschitz continuity, if i is sampled uniformly at random from $\{1, \ldots, N\}$, then for any w,*

$$\mathbf{E}\left[\|\nabla f_i(w) - \nabla f_i(w^*)\|^2 \right] \leq 2L\left(f(w) - f(w^*) \right).$$

Proof of Lemma 11 For any i, define

$$g_i(w) = f_i(w) - f_i(w^*) - (w - w^*)^T \nabla f_i(w^*).$$

Clearly, if i is sampled randomly as in the lemma statement, $\mathbf{E}[g_i(w)] = f(w)$. But also, w^* must be the minimizer of g_i, so for any w

$$g_i(w^*) \leq \min_\eta g_i(w - \eta \nabla g_i(w))$$
$$\leq \min_\eta \left(g_i(w) - \eta \|\nabla g_i(w)\|^2 + \frac{\eta^2 L}{2} \|\nabla g_i(w)\|^2 \right)$$
$$= g_i(w) - \frac{1}{2L} \|\nabla g_i(w)\|^2.$$

where the second inequality follows from the Lipschitz continuity property. Re-writing this in terms of f_i and averaging over all the i now proves the lemma statement. ∎

Lemma 12 *Suppose that we compute the inner loop of SVRG with some computation error. That is, the update step is*

$$w_t = w_{t-1} - \alpha \left(\nabla f_t(w_{t-1}) - \nabla f_t(\tilde{w}) + \nabla f(\tilde{w}) \right) + u_t$$

where f_t is the random training example chosen at timestep t, and where u_t is the error term and satisfies $\|u_t\| \leq \Delta$. Note that u_t is not necessarily independent of ∇f_t. Assume that the step size is small enough that $2\mu - \alpha L^2 > 0$. Then, given some fixed w_{t-1} and $\tilde{w}$ (which are not random variables and are not part of the expected values taken in this Lemma statement),

$$\mathbf{E}\left[\|w_t - w^*\|^2\right] \leq \|w_{t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}) - f(w^*))$$

$$+ 2\Delta \left((1 + \alpha^2 L^2) \|w_{t-1} - w^*\| + \sqrt{2}\alpha L \|\tilde{w} - w^*\| \right) + \Delta^2.$$

Proof Define

$$\nabla h_t(w_{t-1}) = \nabla f_t(w_{t-1}) - \nabla f_t(\tilde{w}) + \nabla f(\tilde{w}).$$

Using this, we can rewrite the update step as

$$w_t - w^* = w_{t-1} - w^* - \alpha \nabla h_t(w_{t-1}) + u_t.$$

Let's start by just looking at the right side of this without the u_t term, which is the same term that we get in the proof of ordinary full-precision SVRG. Taking the expected value, and applying the fact that $\mathbf{E}[h_t] = f$,

$$\mathbf{E}\left[\|w_{t-1} - w^* - \alpha \nabla h_t(w_{t-1})\|^2\right]$$

$$= \|w_{t-1} - w^*\|^2 - 2\alpha \mathbf{E}\left[(w_{t-1} - w^*)^T \nabla h_t(w_{t-1})\right] + \alpha^2 \mathbf{E}\left[\|\nabla h_t(w_{t-1})\|^2\right]$$

$$= \|w_{t-1} - w^*\|^2 - 2\alpha(w_{t-1} - w^*)^T \nabla f(w_{t-1}) + \alpha^2 \mathbf{E}\left[\|\nabla h_t(w_{t-1})\|^2\right]$$

$$\leq \|w_{t-1} - w^*\|^2 - 2\alpha(f(w_{t-1}) - f(w^*)) + \alpha^2 \mathbf{E}\left[\|\nabla h_t(w_{t-1})\|^2\right]$$

This second-order term can be further bounded by

$$\mathbf{E}\left[\|\nabla h_t(w_{t-1})\|^2\right]$$

$$= \mathbf{E}\left[\|\nabla f_t(w_{t-1}) - \nabla f_t(\tilde{w}) + \nabla f(\tilde{w})\|^2\right]$$

$$= \mathbf{E}\left[\|\nabla f_t(w_{t-1}) - \nabla f_t(w^*) - (\nabla f_t(\tilde{w}) - \nabla f_t(w^*) - \nabla f(\tilde{w}))\|^2\right]$$

$$\leq \mathbf{E}\left[2 \|\nabla f_t(w_{t-1}) - \nabla f_t(w^*)\|^2 + 2 \|\nabla f_t(\tilde{w}) - \nabla f_t(w^*) - \nabla f(\tilde{w})\|^2\right]$$

$$= \mathbf{E}\left[2 \|\nabla f_t(w_{t-1}) - \nabla f_t(w^*)\|^2\right]$$

$$\qquad + \mathbf{E}\left[2 \|\nabla f_t(\tilde{w}) - \nabla f_t(w^*) - \mathbf{E}_{j \sim \mathrm{Unif}(1,...,N)} [\nabla f_j(\tilde{w}) - \nabla f_j(w^*)]\|^2\right]$$

$$\leq \mathbf{E}\left[2 \|\nabla f_t(w_{t-1}) - \nabla f_t(w^*)\|^2 + 2 \|\nabla f_t(\tilde{w}) - \nabla f_t(w^*)\|^2\right]$$

where the first inequality holds because $\|x + y\|^2 \leq 2\|x\|^2 + 2\|y\|^2$ and the second holds because the variance is always upper bounded by the second moment. We can now apply Lemma 11 to this last expression, which produces

$$\mathbf{E}\left[\|\nabla h_t(w_{t-1})\|^2\right] \leq 4L(f(w_{t-1}) - f(w^*)) + 4L(f(\tilde{w}) - f(w^*)).$$

Substituting this back into our above bound, we get

$$\begin{aligned}
\mathbf{E}\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|^2\right] &\leq \|w_{t-1} - w^*\|^2 - 2\alpha(f(w_{t-1}) - f(w^*)) \\
&\quad + \alpha^2\left(4L(f(w_{t-1}) - f(w^*)) + 4L(f(\tilde{w}) - f(w^*))\right) \\
&= \|w_{t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{t-1}) - f(w^*)) \\
&\quad + 4\alpha^2 L(f(\tilde{w}) - f(w^*)).
\end{aligned}$$

Another way to bound this same expression is the following. Using convexity of the function f,

$$\begin{aligned}
\mathbf{E}&\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|^2\right] \\
&= \|w_{t-1} - w^*\|^2 - 2\alpha\mathbf{E}\left[(w_{t-1} - w^*)^T\nabla f(w_{t-1})\right] + \alpha^2\mathbf{E}\left[\|\nabla h_t(w_{t-1})\|^2\right] \\
&\leq \mathbf{E}\left[\|w_{t-1} - w^*\|^2\right] + \alpha^2\mathbf{E}\left[\|\nabla h_t(w_{t-1})\|^2\right].
\end{aligned}$$

Further bounding the second-order term using part of our analysis above and Lipschitz continuity of f_t,

$$\begin{aligned}
\mathbf{E}&\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|^2\right] \\
&\leq \|w_{t-1} - w^*\|^2 + \alpha^2\mathbf{E}\left[2\|\nabla f_t(w_{t-1}) - \nabla f_t(w^*)\|^2 + 2\|\nabla f_t(\tilde{w}) - \nabla f_t(w^*)\|^2\right] \\
&\leq \|w_{t-1} - w^*\|^2 + 2\alpha^2 L^2\left(\|w_{t-1} - w^*\|^2 + \|\tilde{w} - w^*\|^2\right) \\
&\leq (1 + 2\alpha^2 L^2)\|w_{t-1} - w^*\|^2 + 2\alpha^2 L^2\|\tilde{w} - w^*\|^2.
\end{aligned}$$

It follows that, by Jensen's inequality,

$$\begin{aligned}
\mathbf{E}\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|\right] &\leq \mathbf{E}\left[\sqrt{\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|^2}\right] \\
&\leq \sqrt{\mathbf{E}\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|^2\right]} \\
&\leq \sqrt{(1 + 2\alpha^2 L^2)\|w_{t-1} - w^*\|^2 + 2\alpha^2 L^2\|\tilde{w} - w^*\|^2} \\
&\leq \sqrt{(1 + 2\alpha^2 L^2)\|w_{t-1} - w^*\|^2} + \sqrt{2\alpha^2 L^2\|\tilde{w} - w^*\|^2} \\
&\leq (1 + \alpha^2 L^2)\|w_{t-1} - w^*\| + \sqrt{2}\alpha L\|\tilde{w} - w^*\|.
\end{aligned}$$

48

Returning to our original expression,

$$
\begin{aligned}
\mathbf{E}\left[\|w_t - w^*\|^2\right] = {} & \mathbf{E}\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|^2\right] \\
& + 2\mathbf{E}\left[u_t^T\left(w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\right)\right] + \mathbf{E}\left[\|u_t\|^2\right] \\
\leq {} & \mathbf{E}\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|^2\right] \\
& + 2\Delta\mathbf{E}\left[\|w_{t-1} - w^* - \alpha\nabla h_t(w_{t-1})\|\right] + \Delta^2 \\
\leq {} & \|w_{t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}) - f(w^*)) \\
& + 2\Delta\left((1 + \alpha^2 L^2)\|w_{t-1} - w^*\| + \sqrt{2}\alpha L\|\tilde{w} - w^*\|\right) + \Delta^2.
\end{aligned}
$$

This is what we wanted to show. $\blacksquare$

Now we are ready to prove Theorem 3. Our proof of this theorem follows the structure of the proof of the original SVRG convergence result in Johnson and Zhang (2013).
Proof of Theorem 3 We start by looking at the inner loop update step for LP-SVRG

$$
w_{k,t} = w_{k,t-1} \ominus_{\mathrm{L}} \alpha \odot_{\mathrm{L}} \left(\left(x_i^T \odot_{\mathrm{L}} w_{k,t-1}\right) \odot_{\mathrm{L}} \mathsf{fl}_{\mathrm{lo}}(x_i) \ominus_{\mathrm{L}} \left(x_i^T \odot_{\mathrm{L}} \tilde{w}_k\right) \odot_{\mathrm{L}} \mathsf{fl}_{\mathrm{lo}}(x_i) \oplus_{\mathrm{L}} \tilde{h}_k\right).
$$

Define

$$
v_{k,t-1} = w_{k,t-1} - \alpha\left(x_i^T w_{k,t-1} \cdot x_i - x_i^T \tilde{w}_k \cdot x_i + \tilde{h}_k\right),
$$

where all the arithmetic in the above expression is exact. This $v_{k,t-1}$ is what the original SVRG algorithm would step to, if it were using exact arithmetic. First, we want to find a bound on the error between $v_{k,t-1}$ and $w_{k,t-1}$: this is the error induced by using low-precision arithmetic. We can start by noticing that, from our standard bound on the error of a dot product, we will have that

$$
x_i^T \odot_{\mathrm{L}} w_{k,t-1} = x_i^T w_{k,t-1} + |x_i|^T |w_{k,t-1}| \cdot d \cdot \varepsilon_{\mathrm{lo}} = x_i^T w_{k,t-1} + \|x_i\| \cdot \|w_{k,t-1}\| \cdot d \cdot \varepsilon_{\mathrm{lo}}.
$$

Note that in the expressions above, as usual $\varepsilon_{\mathrm{lo}}$ refers to some number that is bounded from above by the machine epsilon; multiple instances of $\varepsilon_{\mathrm{lo}}$ may denote different values. Similarly, we will have

$$
x_i \odot_{\mathrm{L}} \tilde{w}_k = x_i^T \tilde{w}_k + \|x_i\| \cdot \|\tilde{w}_k\| \cdot d \cdot \varepsilon_{\mathrm{lo}}.
$$

So, the jth entry of the gradient sample will be

$$
\begin{aligned}
\left(x_i^T \odot_{\mathrm{L}} w_{k,t-1}\right) \odot_{\mathrm{L}} x_{i,j} &= \left(\left(x_i^T \odot_{\mathrm{L}} w_{k,t-1}\right) \cdot x_{i,j}\right) \cdot (1 + \varepsilon_{\mathrm{lo}}) \\
&= \left((x_i^T w_{k,t-1} + \|x_i\| \cdot \|w_{k,t-1}\| \cdot d \cdot \varepsilon_{\mathrm{lo}}) \cdot x_{i,j}\right) \cdot (1 + \varepsilon_{\mathrm{lo}}) \\
&= x_i^T w_{k,t-1} \cdot x_{i,j} + x_i^T w_{k,t-1} \cdot x_{i,j} \cdot \varepsilon_{\mathrm{lo}} + \|x_i\| \cdot \|w_{k,t-1}\| \cdot x_{i,j} \cdot d \cdot \varepsilon_{\mathrm{lo}}, \\
&= x_i^T w_{k,t-1} \cdot x_{i,j} + \|x_i\| \cdot \|w_{k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

where as usual we ignore factors of $O(\varepsilon_{\mathrm{lo}}^2)$. Similarly, it will hold that

$$\left(x_i^T \underset{\mathrm{L}}{\odot} \tilde{w}_k\right) \underset{\mathrm{L}}{\odot} x_{i,j} = x_i^T \tilde{w}_k \cdot x_{i,j} + \|x_i\| \cdot \|\tilde{w}_k\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}}.$$

Subtracting these expressions, we get

$$
\begin{aligned}
\left(x_i^T \underset{\mathrm{L}}{\odot} w_{k,t-1}\right) &\underset{\mathrm{L}}{\odot} x_{i,j} \underset{\mathrm{L}}{\ominus} \left(x_i^T \underset{\mathrm{L}}{\odot} \tilde{w}_k\right) \underset{\mathrm{L}}{\odot} x_{i,j} \\
&= \left(\left(x_i^T \underset{\mathrm{L}}{\odot} w_{k,t-1}\right) \underset{\mathrm{L}}{\odot} x_{i,j} - \left(x_i^T \underset{\mathrm{L}}{\odot} \tilde{w}_k\right) \underset{\mathrm{L}}{\odot} x_{i,j}\right) \cdot (1 + \varepsilon_{\mathrm{lo}}) \\
&= \left(x_i^T w_{k,t-1} \cdot x_{i,j} - x_i^T \tilde{w}_k \cdot x_{i,j}\right. \\
&\qquad \left. + \|x_i\| \cdot \|w_{k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}} + \|x_i\| \cdot \|\tilde{w}_k\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}}\right) \cdot (1 + \varepsilon_{\mathrm{lo}}) \\
&= x_i^T w_{k,t-1} \cdot x_{i,j} - x_i^T \tilde{w}_k \cdot x_{i,j} \\
&\qquad + \left(x_i^T w_{k,t-1} \cdot x_{i,j} - x_i^T \tilde{w}_k \cdot x_{i,j}\right) \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + \|x_i\| \cdot \|w_{k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}} + \|x_i\| \cdot \|\tilde{w}_k\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}} \\
&= x_i^T w_{k,t-1} \cdot x_{i,j} - x_i^T \tilde{w}_k \cdot x_{i,j} \\
&\qquad + \|x_i\| \cdot \|w_{k,t-1}\| \cdot |x_{i,j}| \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}} + \|x_i\| \cdot \|\tilde{w}_k\| \cdot |x_{i,j}| \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}} \\
&= x_i^T w_{k,t-1} \cdot x_{i,j} - x_i^T \tilde{w}_k \cdot x_{i,j} \\
&\qquad + \|x_i\| \cdot (\|w_{k,t-1}\| + \|\tilde{w}_k\|) \cdot |x_{i,j}| \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

If follows that by bounding the total error across all j,

$$
\begin{aligned}
\left\| \left(x_i^T \underset{\mathrm{L}}{\odot} w_{k,t-1}\right) \underset{\mathrm{L}}{\odot} x_i \underset{\mathrm{L}}{\ominus} \left(x_i^T \underset{\mathrm{L}}{\odot} \tilde{w}_k\right) \underset{\mathrm{L}}{\odot} x_i - \left(x_i^T w_{k,t-1} \cdot x_i - x_i^T \tilde{w}_k \cdot x_i\right) \right\| \\
= \left\| \|x_i\| \cdot (\|w_{k,t-1}\| + \|\tilde{w}_k\|) \cdot |x_i| \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}} \right\| \\
= \|x_i\|^2 \cdot (\|w_{k,t-1}\| + \|\tilde{w}_k\|) \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

Next, since we know by assumption that the Lipschitz constant of ∇f_i is L, for any w and v,

$$\left|x_i^T w - x_i^T v\right| \cdot \|x_i\| \le L \|w - v\|.$$

It is straightforward to show that this holds if and only if

$$\|x_i\|^2 \le L.$$

Therefore,

$$
\begin{aligned}
\left\| \left(x_i^T \underset{\mathrm{L}}{\odot} w_{k,t-1}\right) \underset{\mathrm{L}}{\odot} x_i \underset{\mathrm{L}}{\ominus} \left(x_i^T \underset{\mathrm{L}}{\odot} \tilde{w}_k\right) \underset{\mathrm{L}}{\odot} x_i - \left(x_i^T w_{k,t-1} \cdot x_i - x_i^T \tilde{w}_k \cdot x_i\right) \right\| \\
= L \cdot (\|w_{k,t-1}\| + \|\tilde{w}_k\|) \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

Now, to simplify this, let

$$\phi = x_i^T w_{k,t-1} \cdot x_i - x_i^T \tilde{w}_k \cdot x_i$$

and

$$\hat{\phi} = \left(x_i^T \underset{\text{L}}{\odot} w_{k,t-1} \right) \underset{\text{L}}{\odot} x_i \underset{\text{L}}{\ominus} \left(x_i^T \underset{\text{L}}{\odot} \tilde{w}_k \right) \underset{\text{L}}{\odot} x_i.$$

Then we can restate our above bound as

$$\left\| \hat{\phi} - \phi \right\| = L \cdot \left(\|w_{k,t-1}\| + \|\tilde{w}_k\| \right) \cdot (d+2) \cdot \varepsilon_{\text{lo}}.$$

Our inner loop update step for LP-SVRG is

$$w_{k,t-1} = w_{k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\hat{\phi} \underset{\text{L}}{\oplus} \tilde{h}_k \right).$$

We know from the fundamental axiom of floating point arithmetic that for the jth component of this expression,

$$
\begin{aligned}
w_{k,t-1,j} &\underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\hat{\phi}_j \underset{\text{L}}{\oplus} \tilde{h}_{k,j} \right) \\
&= \left(w_{k,t-1,j} - \left(\alpha \cdot \left(\hat{\phi}_j + \tilde{g}_{k,j} \cdot (1+\varepsilon_{\text{lo}}) \right) \cdot (1+\varepsilon_{\text{lo}}) \right) \cdot (1+\varepsilon_{\text{lo}}) \right) \cdot (1+\varepsilon_{\text{lo}}) \\
&= w_{k,t-1} - \alpha \left(\hat{\phi}_j + \tilde{g}_{k,j} \right) + \left(w_{k,t-1,j} + 3\alpha \cdot \hat{\phi}_j + 4\alpha \cdot \tilde{g}_{k,j} \right) \varepsilon_{\text{lo}}, \\
&= w_{k,t-1} - \alpha \left(\phi_j + \tilde{g}_{k,j} \right) - \alpha \left(\hat{\phi}_j - \phi_j \right) + \left(w_{k,t-1,j} + 3\alpha \cdot \phi_j + 4\alpha \cdot \tilde{g}_{k,j} \right) \varepsilon_{\text{lo}}.
\end{aligned}
$$

where again we ignore factors in $O(\varepsilon_{\text{lo}}^2)$. From this, we can bound the total norm of the error

$$
\begin{aligned}
&\left\| \left(w_{k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\hat{\phi} \underset{\text{L}}{\oplus} \tilde{h}_k \right) \right) - \left(w_{k,t-1} - \alpha \left(\phi_j + \tilde{g}_{k,j} \right) \right) \right\| \\
&\qquad \leq \alpha \left\| \hat{\phi}_j - \phi_j \right\| + \left(\|w_{k,t-1}\| + 3\alpha \cdot \|\phi\| + 4\alpha \cdot \|\tilde{g}_k\| \right) \varepsilon_{\text{lo}} \\
&\qquad = \alpha L \cdot \left(\|w_{k,t-1}\| + \|\tilde{w}_k\| \right) \cdot (d+2) \cdot \varepsilon_{\text{lo}} + \left(\|w_{k,t-1}\| + 3\alpha \cdot \|\phi\| + 4\alpha \cdot \|\tilde{g}_k\| \right) \varepsilon_{\text{lo}}.
\end{aligned}
$$

Finally, we can bound

$$\|\phi\| = \left| (x_i^T w_{k,t-1}) - (x_i^T \tilde{w}_k) \right| \cdot \|x_i\| \leq L \|w_{k,t-1} - \tilde{w}_k\|,$$

and so we get

$$
\begin{aligned}
&\left\| \left(w_{k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\hat{\phi} \underset{\text{L}}{\oplus} \tilde{h}_k \right) \right) - \left(w_{k,t-1} - \alpha \left(\phi_j + \tilde{g}_{k,j} \right) \right) \right\| \\
&\qquad = \alpha L \cdot \left(\|w_{k,t-1}\| + \|\tilde{w}_k\| \right) \cdot (d+2) \cdot \varepsilon_{\text{lo}} + \left(\|w_{k,t-1}\| + 3\alpha \cdot (L \|w_{k,t-1} - \tilde{w}_k\|) + 4\alpha \cdot \|\tilde{g}_k\| \right) \varepsilon_{\text{lo}} \\
&\qquad = \alpha L \cdot \left(\|w_{k,t-1}\| + \|\tilde{w}_k\| \right) \cdot (d+5) \cdot \varepsilon_{\text{lo}} + \left(\|w_{k,t-1}\| + 4\alpha \cdot \|\tilde{g}_k\| \right) \varepsilon_{\text{lo}}.
\end{aligned}
$$

Finally, we use the fact that

$$\|\tilde{g}_k\| = \|\nabla f(\tilde{w}_k) - \nabla f(w^*)\| \leq L \cdot \|\tilde{w}_k - w^*\|$$

to get

$$
\left\| \left(w_{k,t-1} \underset{\mathrm{L}}{\ominus} \alpha \underset{\mathrm{L}}{\odot} \left(\hat{\phi} \underset{\mathrm{L}}{\oplus} \tilde{h}_k \right) \right) - (w_{k,t-1} - \alpha \left(\phi_j + \tilde{g}_{k,j} \right)) \right\|
$$

$$
\begin{aligned}
&= \alpha L \cdot (\|w_{k,t-1}\| + \|\tilde{w}_k\|) \cdot (d+5) \cdot \varepsilon_{\mathrm{lo}} + (\|w_{k,t-1}\| + 4\alpha L \cdot \|\tilde{w}_k - w^*\|) \varepsilon_{\mathrm{lo}} \\
&= \alpha L \cdot (\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\| + 2\|w^*\|) \cdot (d+5) \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + (\|w_{k,t-1} - w^*\| + \|w^*\| + 4\alpha L \cdot \|\tilde{w}_k - w^*\|) \varepsilon_{\mathrm{lo}} \\
&\leq \Big((\alpha L(d+5) + 1) \cdot \|w_{k,t-1} - w^*\| + \alpha L(d+9) \cdot \|\tilde{w}_k - w^*\| \\
&\quad + (2\alpha L(d+5) + 1) \cdot \|w^*\| \Big) \cdot \varepsilon_{\mathrm{machine\text{-}lo}} + O(\varepsilon_{\mathrm{machine\text{-}lo}}{}^2).
\end{aligned}
$$

Now we've proved a bound on the distance between the actual step that is done in LP-SVRG and the step that would have been taken with exact number arithmetic. Next, we simplify this result a bit. Suppose (we will make this rigorous through our choice of α later) that $4\alpha L \leq 1$ and that $d > 7$. Then,

$$
\left\| \left(w_{k,t-1} \underset{\mathrm{L}}{\ominus} \alpha \underset{\mathrm{L}}{\odot} \left(\hat{\phi} \underset{\mathrm{L}}{\oplus} \tilde{h}_k \right) \right) - (w_{k,t-1} - \alpha \left(\phi_j + \tilde{g}_{k,j} \right)) \right\|
$$

$$
\begin{aligned}
&\leq \frac{1}{4} \Big((4\alpha L(d+5) + 4) \cdot \|w_{k,t-1} - w^*\| + 4\alpha L(d+9) \cdot \|\tilde{w}_k - w^*\| \\
&\quad + (8\alpha L(d+5) + 4) \cdot \|w^*\| \Big) \cdot \varepsilon_{\mathrm{machine\text{-}lo}} + O(\varepsilon_{\mathrm{machine\text{-}lo}}{}^2) \\
&\leq \frac{1}{4} \Big((d+9) \cdot \|w_{k,t-1} - w^*\| + (d+9) \cdot \|\tilde{w}_k - w^*\| \\
&\quad + (2d+14) \cdot \|w^*\| \Big) \cdot \varepsilon_{\mathrm{machine\text{-}lo}} + O(\varepsilon_{\mathrm{machine\text{-}lo}}{}^2) \\
&\leq \frac{1}{4} \Big((d+3d) \cdot \|w_{k,t-1} - w^*\| + (d+3d) \cdot \|\tilde{w}_k - w^*\| \\
&\quad + (2d+2d) \cdot \|w^*\| \Big) \cdot \varepsilon_{\mathrm{machine\text{-}lo}} + O(\varepsilon_{\mathrm{machine\text{-}lo}}{}^2) \\
&\leq (\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\| + \|w^*\|) \cdot d \cdot \varepsilon_{\mathrm{machine\text{-}lo}} + O(\varepsilon_{\mathrm{machine\text{-}lo}}{}^2).
\end{aligned}
$$

Next, we show how this can be used to prove a bound on the convergence of LP-SVRG. First, by the result of Lemma 12, we will have

$$
\begin{aligned}
\mathbf{E}\left[\|w_{k,t} - w^*\|^2 \right] &\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) \\
&\quad + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 2\Delta \left((1 + \alpha^2 L^2) \|w_{k,t-1} - w^*\| + \sqrt{2}\alpha L \|\tilde{w}_k - w^*\| \right) + \Delta^2,
\end{aligned}
$$

where

$$
\Delta \leq (\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\| + \|w^*\|) \cdot d \cdot \varepsilon_{\mathrm{machine\text{-}lo}} + O(\varepsilon_{\mathrm{machine\text{-}lo}}{}^2).
$$

Subject to our assumption above that $4\alpha L \leq 1$, and noticing that $\Delta = O(\varepsilon_{\text{machine-lo}})$, we can ignore the $O(\Delta^2)$ term and simplify this to

$$
\begin{aligned}
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] &\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) \\
&\quad + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 3\Delta\left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\| + \|w^*\|\right) + O(\varepsilon_{\text{machine-lo}}{}^2) \\
&\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) \\
&\quad + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 3\left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\| + \|w^*\|\right)^2 \cdot d \cdot \varepsilon_{\text{machine-lo}} \\
&\quad + O(\varepsilon_{\text{machine-lo}}{}^2) \\
&\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) \\
&\quad + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 9\left(\|w_{k,t-1} - w^*\|^2 + \|\tilde{w}_k - w^*\|^2 + \|w^*\|^2\right) \cdot d \cdot \varepsilon_{\text{machine-lo}} \\
&\quad + O(\varepsilon_{\text{machine-lo}}{}^2).
\end{aligned}
$$

As a consequence of the strong convexity property,

$$
\frac{\mu}{2}\|w - w^*\|^2 \leq f(w) - f(w^*),
$$

so

$$
\begin{aligned}
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] &\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) \\
&\quad + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 9\left(\frac{2}{\mu}(f(w_{k,t-1}) - f(w^*)) + \frac{2}{\mu}(f(\tilde{w}_k) - f(w^*)) + \|w^*\|^2\right) \cdot d \cdot \varepsilon_{\text{machine-lo}} \\
&\quad + O(\varepsilon_{\text{machine-lo}}{}^2) \\
&\leq \|w_{k,t-1} - w^*\|^2 - \left(2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right)(f(w_{k,t-1}) - f(w^*)) \\
&\quad + \left(4\alpha^2 L + 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right)(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 9\|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2).
\end{aligned}
$$

Note that the above expected values were all taken conditioned on the current value of $w_{k,t-1}$ and $\tilde{w}_k$. Now, if we take the full expected value and telescope-sum this over the whole epoch, we get

$$
\begin{aligned}
\mathbf{E}\left[\|w_{k,T} - w^*\|^2\right] &\leq \mathbf{E}\left[\|w_{k,0} - w^*\|^2\right] \\
&\quad - \left(2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right)\sum_{t=1}^{T}\mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right] \\
&\quad + \left(4\alpha^2 L + 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right] \\
&\quad + 9T\|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2).
\end{aligned}
$$

Using the fact that $w_{k,0} = \tilde{w}_k$ and the norm of anything is non-negative,

$$0 \leq \mathbf{E}\left[\|\tilde{w}_k - w^*\|^2\right] - \left(2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \sum_{t=1}^{T} \mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right]$$
$$+ \left(4\alpha^2 L + 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ 9T\|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2).$$

As a consequence of the strong convexity property,

$$\frac{\mu}{2}\|\tilde{w}_k - w^*\|^2 \leq f(\tilde{w}_k) - f(w^*),$$

so,

$$0 \leq \frac{2}{\mu}\mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right] - \left(2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \sum_{t=1}^{T} \mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right]$$
$$+ \left(4\alpha^2 L + 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ 9T\|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2)$$
$$\leq - \left(2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \sum_{t=1}^{T} \mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right]$$
$$+ \left(\frac{2}{\mu T} + 4\alpha^2 L + 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ 9T\|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2).$$

If we use option II to assign the next outer iterate, then

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] = \frac{1}{T}\sum_{t=1}^{T} \mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right],$$

and so

$$0 \leq - \left(2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right]$$
$$+ \left(\frac{2}{\mu T} + 4\alpha^2 L + 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ 9T\|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2).$$

Now dividing to isolate terms on the left side,

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \frac{\frac{2}{\mu T} + 4\alpha^2 L + 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}}{2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}}\mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ \frac{9\|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}}}{2\alpha(1 - 2\alpha L) - 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}} + O(\varepsilon_{\text{machine-lo}}{}^2).$$

Note that in general, we can simplify a fraction with

$$\frac{a+\epsilon}{b-\epsilon} = \frac{a/b}{1-\epsilon/b} + \frac{\epsilon}{b} + O(\epsilon^2)$$

$$= (a/b)(1+\epsilon/b) + \frac{\epsilon}{b} + O(\epsilon^2)$$

$$= \frac{a}{b} + \frac{a\epsilon}{b^2} + \frac{\epsilon}{b} + O(\epsilon^2)$$

$$= \frac{a}{b} + \frac{a+b}{b^2} \cdot \epsilon + O(\epsilon^2).$$

So,

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \le \left(\frac{\frac{2}{\mu T} + 4\alpha^2 L}{2\alpha(1-2\alpha L)} + \frac{\frac{2}{\mu T} + 4\alpha^2 L + 2\alpha(1-2\alpha L)}{(2\alpha(1-2\alpha L))^2} \cdot 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right)$$
$$\cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ \frac{9\left\|w^*\right\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}}}{2\alpha(1-2\alpha L)} + O(\varepsilon_{\text{machine-lo}}^2).$$

Next, if we set the step size and epoch length in terms of $\gamma \in (0,1)$ as

$$\alpha = \frac{\gamma}{4L(1+\gamma)}$$

and

$$T = \frac{8\kappa(1+\gamma)}{\gamma^2}$$

as it is in the original SVRG analysis (note that this validates our above assumption that $4\alpha L < 1$), we get that

$$2\alpha(1-2\alpha L) = 2 \cdot \frac{\gamma}{4L(1+\gamma)} \cdot \frac{2+\gamma}{2(1+\gamma)} = \frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}$$

and

$$\frac{2}{\mu T} + 4\alpha^2 L = \frac{\gamma^2}{4L(1+\gamma)} + \frac{\gamma^2}{4L(1+\gamma)^2} = \frac{\gamma^2(2+\gamma)}{4L(1+\gamma)^2}$$

and so

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \le \left(\frac{\frac{\gamma^2(2+\gamma)}{4L(1+\gamma)^2}}{\frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}} + \frac{\frac{\gamma^2(2+\gamma)}{4L(1+\gamma)^2} + \frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}}{\left(\frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}\right)^2} \cdot 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right)$$
$$\cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ \frac{9\left\|w^*\right\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}}}{\frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}} + O(\varepsilon_{\text{machine-lo}}^2)$$
$$\le \left(\gamma + \frac{4L(1+\gamma)^3}{\gamma(2+\gamma)} \cdot 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$+ \frac{36L(1+\gamma)^2 \cdot \left\|w^*\right\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}}}{\gamma(2+\gamma)} + O(\varepsilon_{\text{machine-lo}}^2).$$

Using the fact that $0 < \gamma < 1$, we can further bound this with

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\gamma + \frac{4L(1+1)^3}{\gamma(2+0)} \cdot 18\mu^{-1} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$

$$+ \frac{36L(1+1)^2 \cdot \|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}}}{\gamma(2+0)} + O(\varepsilon_{\text{machine-lo}}{}^2)$$

$$\leq \left(\gamma + \frac{288\kappa}{\gamma} \cdot d \cdot \varepsilon_{\text{machine-lo}}\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$

$$+ \frac{72L \cdot \|w^*\|^2 \cdot d \cdot \varepsilon_{\text{machine-lo}}}{\gamma} + O(\varepsilon_{\text{machine-lo}}{}^2).$$

This is what we wanted to prove. ∎

Proof of Theorem 4 We start by looking at the inner loop update step for BC-SVRG.

$$\delta_{w,k,t} = \delta_{w,k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_i \underset{\text{L}}{\oplus} \tilde{h}_k\right).$$

From the same analysis of the dot product used in the proof of Theorem 3, it will hold that

$$\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_{i,j} = x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\text{lo}}.$$

When the $\tilde{h}_k$ term is added, we get

$$\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_{i,j} \underset{\text{L}}{\oplus} \tilde{h}_{k,j}$$

$$= \left(\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_{i,j} + \tilde{h}_{k,j}\right) \cdot (1 + \varepsilon_{\text{lo}})$$

$$= \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\text{lo}} + \tilde{g}_{k,j} \cdot (1 + \varepsilon_{\text{lo}})\right) \cdot (1 + \varepsilon_{\text{lo}})$$

$$= x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} + x_i^T \delta_{w,k,t-1} \cdot x_{i,j} \cdot \varepsilon_{\text{lo}}$$

$$\quad + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\text{lo}} + 2\tilde{g}_{k,j} \cdot \varepsilon_{\text{lo}}$$

$$= x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+2) \cdot \varepsilon_{\text{lo}} + 2\tilde{g}_{k,j} \cdot \varepsilon_{\text{lo}}$$

where as usual we ignore factors of $O(\varepsilon_{\text{lo}}{}^2)$ by rolling them up into the other ε_{lo} terms. Now multiplying this with α,

$$\alpha \underset{\text{L}}{\odot} \left(\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_{i,j} \underset{\text{L}}{\oplus} \tilde{h}_{k,j}\right)$$

$$= \alpha \cdot \left(\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_{i,j} \underset{\text{L}}{\oplus} \tilde{h}_{k,j}\right) \cdot (1 + \varepsilon_{\text{lo}})$$

$$= \alpha \cdot \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+2) \cdot \varepsilon_{\text{lo}} + 2\tilde{g}_{k,j} \cdot \varepsilon_{\text{lo}}\right) \cdot (1 + \varepsilon_{\text{lo}})$$

$$= \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j}\right) + \alpha \cdot \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+3) \cdot \varepsilon_{\text{lo}} + 3\alpha \cdot \tilde{g}_{k,j} \cdot \varepsilon_{\text{lo}}.$$

Finally, subtracting this from $\delta_{w,k,t-1}$,

$$
\begin{aligned}
\delta_{w,k,t-1,j} \ominus_{\mathrm{L}} \alpha & \odot_{\mathrm{L}} \left(\left(x_i^T \odot_{\mathrm{L}} \delta_{w,k,t-1} \right) \odot_{\mathrm{L}} x_{i,j} \oplus_{\mathrm{L}} \tilde{h}_{k,j} \right) \\
&= \left(\delta_{w,k,t-1,j} - \alpha \odot_{\mathrm{L}} \left(\left(x_i^T \odot_{\mathrm{L}} \delta_{w,k,t-1} \right) \odot_{\mathrm{L}} x_{i,j} \oplus_{\mathrm{L}} \tilde{h}_{k,j} \right) \right) \cdot (1 + \varepsilon_{\mathrm{lo}}) \\
&= \Big(\delta_{w,k,t-1,j} - \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} \right) \\
&\qquad + \alpha \cdot \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+3) \cdot \varepsilon_{\mathrm{lo}} + 3\alpha \cdot \tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}} \Big) \cdot (1 + \varepsilon_{\mathrm{lo}}) \\
&= \delta_{w,k,t-1,j} - \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} \right) \\
&\qquad + \delta_{w,k,t-1,j} \cdot \varepsilon_{\mathrm{lo}} + \alpha \cdot \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+4) \cdot \varepsilon_{\mathrm{lo}} + 4\alpha \cdot \tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

Taking the norm, it follows that the total error from low-precision computation in the inner loop update step is

$$
\begin{aligned}
\Delta &= \left\| \delta_{w,k,t-1} \ominus_{\mathrm{L}} \alpha \odot_{\mathrm{L}} \left(\left(x_i^T \odot_{\mathrm{L}} \delta_{w,k,t-1} \right) \odot_{\mathrm{L}} x_i \oplus_{\mathrm{L}} \tilde{h}_k \right) - \left(\delta_{w,k,t-1} - \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_i + \tilde{g}_k \right) \right) \right\| \\
&= \|\delta_{w,k,t-1}\| \cdot \varepsilon_{\mathrm{lo}} + \alpha \cdot \|x_i\|^2 \cdot \|\delta_{w,k,t-1}\| \cdot (d+4) \cdot \varepsilon_{\mathrm{lo}} + 4\alpha \cdot \|\tilde{g}_{k,j}\| \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

Next, leveraging the fact that by Lipschitz continuity, $\|x_i\|^2 \le L$ and $\|\tilde{g}_k\| \le L \cdot \|\tilde{w}_k - w^*\|$,

$$
\Delta = \left((1 + \alpha L(d+4)) \|\delta_{w,k,t-1}\| + 4\alpha L \cdot \|\tilde{w}_k - w^*\| \right) \cdot \varepsilon_{\mathrm{lo}}.
$$

Now, if we define

$$
w_{k,t} = \tilde{w}_k + \delta_{w,k,t},
$$

that is, $w_{k,t}$ is the total number represented using bit centering (the sum of the offset and the delta), then we can rewrite this as

$$
\begin{aligned}
\Delta &= \left((1 + \alpha L(d+4)) \|w_{k,t-1} - \tilde{w}_k\| + 4\alpha L \cdot \|\tilde{w}_k - w^*\| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&= \left((1 + \alpha L(d+4)) \|w_{k,t-1} - w^*\| + (1 + \alpha L(d+4)) \|\tilde{w}_k - w^*\| + 4\alpha L \cdot \|\tilde{w}_k - w^*\| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&= \left((1 + \alpha L(d+4)) \|w_{k,t-1} - w^*\| + (1 + \alpha L(d+8)) \|\tilde{w}_k - w^*\| \right) \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

Next, we simplify this a bit. Suppose (we will make this rigorous through our choice of α later) that $4\alpha L \le 1$ and that $d \ge 4$. Then,

$$
\begin{aligned}
\Delta &= \frac{1}{4} \left((4 + 4\alpha L(d+4)) \|w_{k,t-1} - w^*\| + (4 + 4\alpha L(d+8)) \|\tilde{w}_k - w^*\| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&= \frac{1}{4} \left((d+8) \|w_{k,t-1} - w^*\| + (d+12) \|\tilde{w}_k - w^*\| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&= \frac{1}{4} \left((d+3d) \|w_{k,t-1} - w^*\| + (d+3d) \|\tilde{w}_k - w^*\| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&= \left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\| \right) \cdot d \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

Finally, we bound this using the bound on $\varepsilon_{\mathrm{lo}}$,

$$
\Delta = \left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\| \right) \cdot d \cdot \left(\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2) \right).
$$

57

Note that the big-O in the term of $O(\varepsilon_{\text{machine-lo}}{}^2)$ does not hide any terms that depend on the iterates or statistical problem parameters (L and μ); this can be verified by looking at the previous derivation. Now, we apply Lemma 12, which gives us

$$
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] \leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*))
$$
$$
+ 2\Delta\left((1 + \alpha^2 L^2)\|w_{k,t-1} - w^*\| + \sqrt{2}\alpha L\|\tilde{w}_k - w^*\|\right) + \Delta^2,
$$

where Δ is as defined above. Subject to our earlier assumption that $4\alpha L \leq 1$, we can simplify this to

$$
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] \leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*))
$$
$$
+ 3\Delta\left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\|\right) + \Delta^2.
$$

Next, we substitute our above bound on Δ, which gives us

$$
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] \leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*))
$$
$$
+ 3\left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\|\right)^2 \cdot d \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2))
$$
$$
+ \left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\|\right)^2 \cdot d^2 \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2))^2
$$
$$
\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*))
$$
$$
+ 3\left(\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\|\right)^2 \cdot d \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2))
$$
$$
\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*))
$$
$$
+ 6\left(\|w_{k,t-1} - w^*\|^2 + \|\tilde{w}_k - w^*\|^2\right) d \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2)).
$$

For simplicity, for the next stretch of the proof we define

$$
\epsilon = \varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}{}^2).
$$

As a consequence of the strong convexity property,

$$
\frac{\mu}{2}\|w - w^*\|^2 \leq f(w) - f(w^*).
$$

So, we can bound this further with

$$
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] \leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*))
$$
$$
+ \frac{12}{\mu}\left((f(w_{k,t-1}) - f(w^*)) + (f(\tilde{w}_k) - f(w^*))\right) d \cdot \epsilon
$$
$$
\leq \|w_{k,t-1} - w^*\|^2 - \left(2\alpha(1 - 2\alpha L) - 12\mu^{-1}d \cdot \epsilon\right)(f(w_{k,t-1}) - f(w^*))
$$
$$
+ \left(4\alpha^2 L + 12\mu^{-1}d \cdot \epsilon\right)(f(\tilde{w}_k) - f(w^*)).
$$

The expectations above were all taken with respect to $\delta_{w,k,t-1}$ and $\tilde{w}_k$. Now taking the full expectation and summing over an epoch, we get

$$
\begin{aligned}
\mathbf{E}\left[\|w_{k,T} - w^*\|^2\right] &\leq \mathbf{E}\left[\|w_{k,0} - w^*\|^2\right] \\
&\quad - \sum_{t=1}^{T}\left(2\alpha(1 - 2\alpha L) - 12\mu^{-1}d \cdot \epsilon\right)\mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right] \\
&\quad + \sum_{t=1}^{T}\left(4\alpha^2 L + 12\mu^{-1}d \cdot \epsilon\right)\mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right] \\
&\leq \mathbf{E}\left[\|w_{k,0} - w^*\|^2\right] \\
&\quad - \left(2\alpha(1 - 2\alpha L) - 12\mu^{-1}d \cdot \epsilon\right)\sum_{t=1}^{T}\mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right] \\
&\quad + \left(4\alpha^2 L + 12\mu^{-1}d \cdot \epsilon\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].
\end{aligned}
$$

As usual, now if we use Option II to assign the next outer iterate, we will have

$$
T \cdot \mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] = \sum_{t=1}^{T}\mathbf{E}\left[f(w_{k,t-1}) - f(w^*)\right].
$$

Applying this, the fact that the expected value of anything is always nonnegative, and the fact that $w_{k,0} = \tilde{w}_k$ gives us

$$
\begin{aligned}
0 &\leq \mathbf{E}\left[\|\tilde{w}_k - w^*\|^2\right] - \left(2\alpha(1 - 2\alpha L) - 12\mu^{-1}d \cdot \epsilon\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \\
&\quad + \left(4\alpha^2 L + 12\mu^{-1}d \cdot \epsilon\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].
\end{aligned}
$$

Applying strong convexity again produces

$$
\begin{aligned}
0 &\leq \frac{2}{\mu}\mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right] - \left(2\alpha(1 - 2\alpha L) - 12\mu^{-1}d \cdot \epsilon\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \\
&\quad + \left(4\alpha^2 L + 12\mu^{-1}d \cdot \epsilon\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right] \\
&= -\left(2\alpha(1 - 2\alpha L) - 12\mu^{-1}d \cdot \epsilon\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \\
&\quad + \left(\frac{2}{\mu T} + 4\alpha^2 L + 12\mu^{-1}d \cdot \epsilon\right) \cdot T \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].
\end{aligned}
$$

Dividing to isolate the $\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right]$ term,

$$
\begin{aligned}
\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] &\leq \frac{\frac{2}{\mu T} + 4\alpha^2 L + 12\mu^{-1}d \cdot \epsilon}{2\alpha(1 - 2\alpha L) - 12\mu^{-1}d \cdot \epsilon} \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right] \\
&= \left(\frac{\frac{2}{\mu T} + 4\alpha^2 L}{2\alpha(1 - 2\alpha L)} + \frac{\frac{2}{\mu T} + 4\alpha^2 L + 2\alpha(1 - 2\alpha L)}{(2\alpha(1 - 2\alpha L))^2} \cdot 12\mu^{-1}d \cdot (\epsilon + O(\epsilon^2))\right) \\
&\quad \cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right],
\end{aligned}
$$

where in the last line we are doing a first-order expansion in terms of ϵ. Now, we return back to our notation using $\varepsilon_{\text{machine-lo}}$,

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right]$$
$$\leq \left(\frac{\frac{2}{\mu T} + 4\alpha^2 L}{2\alpha(1 - 2\alpha L)} + \frac{\frac{2}{\mu T} + 4\alpha^2 L + 2\alpha(1 - 2\alpha L)}{(2\alpha(1 - 2\alpha L))^2} \cdot 12\mu^{-1}d \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right)$$
$$\cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right],$$

Next, if (as in the analysis of LP-SVRG) we set the step size and epoch length in terms of $\gamma \in (0, 1)$ as

$$\alpha = \frac{\gamma}{4L(1 + \gamma)}$$

and

$$T = \frac{8\kappa(1 + \gamma)}{\gamma^2}$$

as it is in the original SVRG analysis (note that this validates our above assumption that $4\alpha L < 1$), we get that

$$2\alpha(1 - 2\alpha L) = 2 \cdot \frac{\gamma}{4L(1 + \gamma)} \cdot \frac{2 + \gamma}{2(1 + \gamma)} = \frac{\gamma(2 + \gamma)}{4L(1 + \gamma)^2}$$

and

$$\frac{2}{\mu T} + 4\alpha^2 L = \frac{\gamma^2}{4L(1 + \gamma)} + \frac{\gamma^2}{4L(1 + \gamma)^2} = \frac{\gamma^2(2 + \gamma)}{4L(1 + \gamma)^2}.$$

So,

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\frac{\frac{\gamma^2(2+\gamma)}{4L(1+\gamma)^2}}{\frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}} + \frac{\frac{\gamma^2(2+\gamma)}{4L(1+\gamma)^2} + \frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}}{\left(\frac{\gamma(2+\gamma)}{4L(1+\gamma)^2}\right)^2} \cdot 12\mu^{-1}d \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right)$$
$$\cdot \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$\leq \left(\gamma + \frac{4L(1+\gamma)^3}{\gamma(2+\gamma)} \cdot 12\mu^{-1} \cdot d \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].$$

Using the fact that $0 < \gamma < 1$, we can further bound this with

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\gamma + \frac{4L(1+1)^3}{\gamma(2+0)} \cdot 12\mu^{-1} \cdot d \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right]$$
$$\leq \left(\gamma + \frac{192L \cdot d}{\gamma\mu} \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].$$

This is what we wanted to prove. ∎

Lemma 13 *In a situation where a dot product could underflow (but can not overflow), the result of a n-dimensional dot product done in floating point arithmetic satisfies*

$$x^T \odot y = x^T y + |x|^T |y| \cdot n \cdot \epsilon + 2n \cdot \eta.$$

Proof We will prove this by induction. First, note that, by the fundamental axiom of floating point arithmetic with underflow,

$$x_i \odot y_i = x_i \cdot y_i \cdot (1 + \epsilon) + \eta.$$

It follows directly that the lemma is true for $n = 1$. Now, supposing (by way of induction) the lemma is true for $n - 1$, let s denote the dot product of all but the last entries of the two vectors (x_n and y_n). Then,

$$x^T \odot y = s \oplus x_n \odot y_n.$$

Applying the fundamental axiom again, we get

$$\begin{aligned}
x^T \odot y &= (s + (x_n \odot y_n)) \cdot (1 + \epsilon) + \eta \\
&= \left((x^T y - x_n y_n) + \left(|x|^T |y| - |x_n y_n| \right) \cdot (n - 1) \cdot \epsilon + 2(n - 1) \cdot \eta \right. \\
&\qquad \left. + x_n \cdot y_n \cdot (1 + \epsilon) + \eta \right) \cdot (1 + \epsilon) + \eta \\
&= \left(x^T y + |x|^T |y| \cdot (n - 1) \cdot \epsilon + (2n - 1) \cdot \eta \right) \cdot (1 + \epsilon) + \eta \\
&= x^T y + |x|^T |y| \cdot n \cdot \epsilon + 2n \cdot \eta.
\end{aligned}$$

This proves the lemma. ∎

Proof of Theorem 5 We start by looking at the inner loop update step for HALP, which is the same as the one for BC-SVRG.

$$\delta_{w,k,t} = \delta_{w,k,t-1} \underset{L}{\ominus} \alpha \underset{L}{\odot} \left(\left(x_i^T \underset{L}{\odot} \delta_{w,k,t-1} \right) \underset{L}{\odot} x_i \underset{L}{\oplus} \tilde{h}_k \right).$$

Unlike in the analysis of BC-SVRG, here we are also considering overflow and underflow, so the proof is a bit more involved. By Lemma 13 and Cauchy-Schwarz,

$$x_i^T \underset{L}{\odot} \delta_{w,k,t-1} = x_i^T \delta_{w,k,t-1} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot d \cdot \varepsilon_{\mathrm{lo}} + 2 \cdot d \cdot \eta_{\mathrm{lo\text{-}bias}(B)}.$$

Next, multiplying by the jth entry of x_i,

$$\begin{aligned}
& \left(x_i^T \underset{L}{\odot} \delta_{w,k,t-1} \right) \underset{L}{\odot} x_{i,j} \\
&= \left(x_i^T \underset{L}{\odot} \delta_{w,k,t-1} \right) \cdot x_{i,j} \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \left(x_i^T \delta_{w,k,t-1} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot d \cdot \varepsilon_{\mathrm{lo}} + 2 \cdot d \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \right) \cdot x_{i,j} \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot x_{i,j} \cdot d \cdot \varepsilon_{\mathrm{lo}} + 2 x_{i,j} \cdot d \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \\
&\qquad + x_i^T \delta_{w,k,t-1} \cdot x_{i,j} \cdot \varepsilon_{\mathrm{lo}} + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot x_{i,j} \cdot (d + 1) \cdot \varepsilon_{\mathrm{lo}} + (1 + 2 x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)}.
\end{aligned}$$

When the $\tilde{h}_k$ term is added, we get

$$
\begin{aligned}
&\left(x_i^T \underset{\mathrm{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\mathrm{L}}{\odot} x_{i,j} \underset{\mathrm{L}}{\oplus} \tilde{h}_{k,j} \\
&= \left(\left(x_i^T \underset{\mathrm{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\mathrm{L}}{\odot} x_{i,j} + \tilde{h}_{k,j} \right) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \Big(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}} + (1 + 2x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \\
&\qquad + \tilde{g}_{k,j} \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \Big) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} + x_i^T \delta_{w,k,t-1} \cdot x_{i,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+1) \cdot \varepsilon_{\mathrm{lo}} + 2\tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + (1 + 2x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{lo\text{-}bias}(B)} + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}} + 2\tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + (3 + 2x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)}
\end{aligned}
$$

where as usual we ignore factors of $O({\varepsilon_{\mathrm{lo}}}^2)$ and similar by rolling them up into the other $\varepsilon_{\mathrm{lo}}$ and $\eta_{\mathrm{lo\text{-}bias}(B)}$ terms. Now multiplying this with α,

$$
\begin{aligned}
&\alpha \underset{\mathrm{L}}{\odot} \left(\left(x_i^T \underset{\mathrm{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\mathrm{L}}{\odot} x_{i,j} \underset{\mathrm{L}}{\oplus} \tilde{h}_{k,j} \right) \\
&= \alpha \cdot \left(\left(x_i^T \underset{\mathrm{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\mathrm{L}}{\odot} x_{i,j} \underset{\mathrm{L}}{\oplus} \tilde{h}_{k,j} \right) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \alpha \cdot \Big(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} + \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+2) \cdot \varepsilon_{\mathrm{lo}} + 2\tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + (3 + 2x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \Big) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} \right) + \alpha \cdot \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+3) \cdot \varepsilon_{\mathrm{lo}} + 3\alpha \cdot \tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + (1 + 3\alpha + 2\alpha x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)}
\end{aligned}
$$

Finally, subtracting this from $\delta_{w,k,t-1}$,

$$
\begin{aligned}
&\delta_{w,k,t-1,j} \underset{\mathrm{L}}{\ominus} \alpha \underset{\mathrm{L}}{\odot} \left(\left(x_i^T \underset{\mathrm{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\mathrm{L}}{\odot} x_{i,j} \underset{\mathrm{L}}{\oplus} \tilde{h}_{k,j} \right) \\
&= \left(\delta_{w,k,t-1,j} - \alpha \underset{\mathrm{L}}{\odot} \left(\left(x_i^T \underset{\mathrm{L}}{\odot} \delta_{w,k,t-1} \right) \underset{\mathrm{L}}{\odot} x_{i,j} \underset{\mathrm{L}}{\oplus} \tilde{h}_{k,j} \right) \right) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \Big(\delta_{w,k,t-1,j} - \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} \right) \\
&\qquad + \alpha \cdot \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+3) \cdot \varepsilon_{\mathrm{lo}} + 3\alpha \cdot \tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + (1 + 3\alpha + 2\alpha x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \Big) \cdot (1 + \varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \delta_{w,k,t-1,j} - \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_{i,j} + \tilde{g}_{k,j} \right) \\
&\qquad + \delta_{w,k,t-1,j} \cdot \varepsilon_{\mathrm{lo}} + \alpha \cdot \|x_i\| \cdot \|\delta_{w,k,t-1}\| \cdot |x_{i,j}| \cdot (d+4) \cdot \varepsilon_{\mathrm{lo}} + 4\alpha \cdot \tilde{g}_{k,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\qquad + (2 + 3\alpha + 2\alpha x_{i,j} \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \cdot
\end{aligned}
$$

62

Taking the norm, it follows that the total error from low-precision computation in the inner loop update step is

$$
\begin{aligned}
\Delta &= \left\| \delta_{w,k,t-1} \ominus_{\mathrm{L}} \alpha \odot_{\mathrm{L}} \left(\left(x_i^T \odot_{\mathrm{L}} \delta_{w,k,t-1} \right) \odot_{\mathrm{L}} x_i \oplus_{\mathrm{L}} \tilde{h}_k \right) - \left(\delta_{w,k,t-1} - \alpha \left(x_i^T \delta_{w,k,t-1} \cdot x_i + \tilde{g}_k \right) \right) \right\| \\
&= \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + \alpha \cdot \| x_i \|^2 \cdot \| \delta_{w,k,t-1} \| \cdot (d+4) \cdot \varepsilon_{\mathrm{lo}} + 4\alpha \cdot \| \tilde{g}_k \| \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + (2\sqrt{d} + 3\alpha\sqrt{d} + 2\alpha \| x_i \| \cdot d) \cdot \eta_{\mathrm{lo\text{-}bias}(B)}.
\end{aligned}
$$

Now, we know from the construction of the HALP algorithm that

$$
\begin{aligned}
\eta_{\mathrm{machine\text{-}lo\text{-}bias}(B)} &= 2^B \cdot \eta_{\mathrm{machine\text{-}lo}} \\
&= 2^{\lfloor \log_2(\zeta \cdot \| \tilde{g}_k \|) \rfloor} \cdot \eta_{\mathrm{machine\text{-}lo}} \\
&\leq 2^{\log_2(\zeta \cdot \| \tilde{g}_k \|) - 1} \cdot \eta_{\mathrm{machine\text{-}lo}} \\
&= \frac{1}{2} \cdot \zeta \cdot \| \tilde{g}_k \| \cdot \eta_{\mathrm{machine\text{-}lo}}.
\end{aligned}
$$

Substituting this in above produces

$$
\begin{aligned}
\Delta &= \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + \alpha \cdot \| x_i \|^2 \cdot \| \delta_{w,k,t-1} \| \cdot (d+4) \cdot \varepsilon_{\mathrm{lo}} + 4\alpha \cdot \| \tilde{g}_k \| \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{2}(2\sqrt{d} + 3\alpha\sqrt{d} + 2\alpha \| x_i \| \cdot d) \cdot \zeta \cdot \| \tilde{g}_k \| \cdot \eta_{\mathrm{lo}}.
\end{aligned}
$$

Next, leveraging the fact that by Lipschitz continuity, $\| x_i \|^2 \leq L$ and $\| \tilde{g}_k \| \leq L \cdot \| \tilde{w}_k - w^* \|$,

$$
\begin{aligned}
\Delta &= \left((1 + \alpha L(d+4)) \| \delta_{w,k,t-1} \| + 4\alpha L \cdot \| \tilde{w}_k - w^* \| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{2}(2\sqrt{d} + 3\alpha\sqrt{d} + 2\alpha\sqrt{L} \cdot d) \cdot \zeta \cdot L \cdot \| \tilde{w}_k - w^* \| \cdot \eta_{\mathrm{lo}}.
\end{aligned}
$$

Now, if we define

$$
w_{k,t} = \tilde{w}_k + \delta_{w,k,t},
$$

that is, $w_{k,t}$ is the total number represented using bit centering (the sum of the offset and the delta), then we can rewrite this as

$$
\begin{aligned}
\Delta &= \left((1 + \alpha L(d+4)) \| w_{k,t-1} - \tilde{w}_k \| + 4\alpha L \cdot \| \tilde{w}_k - w^* \| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{2}(2\sqrt{d} + 3\alpha\sqrt{d} + 2\alpha\sqrt{L} \cdot d) \cdot \zeta \cdot L \cdot \| \tilde{w}_k - w^* \| \cdot \eta_{\mathrm{lo}} \\
&= \left((1 + \alpha L(d+4)) \| w_{k,t-1} - w^* \| + (1 + \alpha L(d+4)) \| \tilde{w}_k - w^* \| + 4\alpha L \cdot \| \tilde{w}_k - w^* \| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{2}(2\sqrt{d} + 3\alpha\sqrt{d} + 2\alpha\sqrt{L} \cdot d) \cdot \zeta \cdot L \cdot \| \tilde{w}_k - w^* \| \cdot \eta_{\mathrm{lo}} \\
&= \left((1 + \alpha L(d+4)) \| w_{k,t-1} - w^* \| + (1 + \alpha L(d+8)) \| \tilde{w}_k - w^* \| \right) \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{2}(2\sqrt{d} + 3\alpha\sqrt{d} + 2\alpha\sqrt{L} \cdot d) \cdot \zeta \cdot L \cdot \| \tilde{w}_k - w^* \| \cdot \eta_{\mathrm{lo}}.
\end{aligned}
$$

Next, we simplify this a bit. Suppose (we will make this rigorous through our choice of α later) that $4\alpha L \leq 1$ and that $d \geq 16$. Then,

$$
\begin{aligned}
\Delta &= \frac{1}{4}\left((4 + 4\alpha L(d+4))\left\|w_{k,t-1} - w^*\right\| + (4 + 4\alpha L(d+8))\left\|\tilde{w}_k - w^*\right\|\right) \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{8}(8L\sqrt{d} + 12\alpha L\sqrt{d} + 8\alpha L\sqrt{L} \cdot d) \cdot \zeta \cdot \left\|\tilde{w}_k - w^*\right\| \cdot \eta_{\mathrm{lo}} \\
&= \frac{1}{4}\left((d+8)\left\|w_{k,t-1} - w^*\right\| + (d+12)\left\|\tilde{w}_k - w^*\right\|\right) \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{8}(8L\sqrt{d} + 3\sqrt{d} + 2\sqrt{L} \cdot d) \cdot \zeta \cdot \left\|\tilde{w}_k - w^*\right\| \cdot \eta_{\mathrm{lo}}.
\end{aligned}
$$

We can further bound this last term with

$$
\begin{aligned}
\frac{1}{8}(8L\sqrt{d} + 3\sqrt{d} + 2\sqrt{L} \cdot d) &\leq \frac{d}{8}\left(\frac{8L}{\sqrt{d}} + \frac{3}{\sqrt{d}} + \sqrt{L}\right) \\
&\leq \frac{d}{8}\left(\frac{8L}{4} + \frac{3}{4} + \frac{L+1}{2}\right) \\
&= \frac{d}{8}\left(\frac{5L}{2} + \frac{5}{4}\right) \\
&\leq \frac{d}{2}(L+1).
\end{aligned}
$$

Substituting this back in above, and leveraging our assumption that

$$(L+1)\zeta \cdot \eta_{\mathrm{machine\text{-}lo}} \leq \varepsilon_{\mathrm{machine\text{-}lo}},$$

we get

$$
\begin{aligned}
\Delta &= \frac{1}{4}\left((d+8)\left\|w_{k,t-1} - w^*\right\| + (d+12)\left\|\tilde{w}_k - w^*\right\|\right) \cdot \varepsilon_{\mathrm{lo}} + \frac{d}{2}(L+1) \cdot \zeta \cdot \left\|\tilde{w}_k - w^*\right\| \cdot \eta_{\mathrm{lo}} \\
&= \frac{1}{4}\left((d+8)\left\|w_{k,t-1} - w^*\right\| + (d+12)\left\|\tilde{w}_k - w^*\right\|\right) \cdot \varepsilon_{\mathrm{lo}} + \frac{d}{2} \cdot \left\|\tilde{w}_k - w^*\right\| \cdot \varepsilon_{\mathrm{lo}} \\
&= \frac{1}{4}\left((d+8)\left\|w_{k,t-1} - w^*\right\| + (3d+12)\left\|\tilde{w}_k - w^*\right\|\right) \cdot \varepsilon_{\mathrm{lo}} \\
&\leq \frac{1}{4}\left((d+3d)\left\|w_{k,t-1} - w^*\right\| + (3d+d)\left\|\tilde{w}_k - w^*\right\|\right) \cdot \varepsilon_{\mathrm{lo}} \\
&= \left(\left\|w_{k,t-1} - w^*\right\| + \left\|\tilde{w}_k - w^*\right\|\right) \cdot d \cdot \varepsilon_{\mathrm{lo}}.
\end{aligned}
$$

Finally, we bound this using the bound on $\varepsilon_{\mathrm{lo}}$,

$$\Delta = \left(\left\|w_{k,t-1} - w^*\right\| + \left\|\tilde{w}_k - w^*\right\|\right) \cdot d \cdot (\varepsilon_{\mathrm{machine\text{-}lo}} + O(\varepsilon_{\mathrm{machine\text{-}lo}}{}^2)).$$

Now, when we perform the update in HALP, let $u_{k,t}$ denote the value of the iterate before possibly re-setting the value of δ to zero in the if statement. This is just the value for the model that we get after the inner loop update step of SVRG, so we can apply Lemma 12, which gives us

$$
\begin{aligned}
\mathbf{E}\left[\left\|u_{k,t} - w^*\right\|^2\right] &\leq \left\|w_{k,t-1} - w^*\right\|^2 - 2\alpha(1-2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 2\Delta\left((1+\alpha^2 L^2)\left\|w_{k,t-1} - w^*\right\| + \sqrt{2}\alpha L\left\|\tilde{w}_k - w^*\right\|\right) + \Delta^2,
\end{aligned}
$$

where Δ is as defined above. Let $w_{k,t}$ denote the value after the if statement. From our analysis in Section 4.3, we know that the resetting of δ, if it happens, only brings the value of $w_{k,t}$ closer to the optimum. As a consequence,

$$\|w_{k,t} - w^*\|^2 \leq \|u_{k,t} - w^*\|^2 \,;$$

usually these terms will be equal since on almost all iterations the body of that if block is not run and δ is not reset to zero. Plugging this into our formula above, we get

$$\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] \leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*))$$
$$+ 2\Delta\left((1 + \alpha^2 L^2)\|w_{k,t-1} - w^*\| + \sqrt{2}\alpha L\|\tilde{w}_k - w^*\|\right) + \Delta^2.$$

This is the exact same expression that we got in the proof of Theorem 4. As a consequence, if we apply the exact same logic as we did in that proof, we can reach the same conclusion, which is

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\gamma + \frac{192\kappa d}{\gamma} \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right)\mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].$$

This is the first part of what we wanted to prove.

The second part is to show that no overflow occurs. To do this, we consider the largest possible value that could result in the inner loop of HALP, ignoring floating point error for now. Because of our if-guard on the magnitude of the iterates, all the numbers used to store $\delta_{w,k,t-1}$ are bounded in magnitude by

$$2\mu^{-1} \cdot \|\tilde{g}_k\|\,.$$

Within the computation of the dot product

$$x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}$$

all the numbers computed will be bounded in magnitude by

$$|x_i|^T \cdot |\delta_{w,k,t-1}| \leq \|x_i\| \cdot \|\delta_{w,k,t-1}\| \leq 2\sqrt{L}\mu^{-1} \cdot \|\tilde{g}_k\|\,.$$

Similarly, the result of computing the scalar-vector product

$$\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_i$$

will have numbers all bounded in magnitude by

$$\|x_i\|^2 \cdot \|\delta_{w,k,t-1}\| \leq 2L\mu^{-1} \cdot \|\tilde{g}_k\|\,.$$

The result of adding this to $\tilde{h}_k$

$$\left(x_i^T \underset{\text{L}}{\odot} \delta_{w,k,t-1}\right) \underset{\text{L}}{\odot} x_i \underset{\text{L}}{\oplus} \tilde{h}_k$$

will have numbers all bounded in magnitude by

$$2L\mu^{-1} \cdot \|\tilde{g}_k\| + \|\tilde{g}_k\| = (2\kappa + 1) \cdot \|\tilde{g}_k\|.$$

Finally, the result of multiplying this by α will be bounded by, given that $4\alpha L \leq 1$,

$$2\alpha L\mu^{-1} \cdot \|\tilde{g}_k\| + \|\tilde{g}_k\| = \frac{2\kappa + 1}{4L} \cdot \|\tilde{g}_k\|.$$

If there is overflow in the addition with δ, it will not matter, because we will reset δ to zero anyway. So all our numbers are bounded in magnitude by

$$\max\left(2\mu^{-1}, 2\sqrt{L}\mu^{-1}, 2\kappa, 2\kappa + 1, \frac{2\kappa + 1}{4L}\right) \cdot \|\tilde{g}_k\|.$$

We can see immediately that $2\sqrt{L}\mu^{-1}$ and 2κ are redundant in this max. Additionally, since $\mu^{-1} = \kappa/L$, we can further simplify this upper bound to

$$\max\left(2\kappa + 1, \frac{2\kappa + 1}{L}\right) \cdot \|\tilde{g}_k\| = \max\left(1, L^{-1}\right) \cdot (2\kappa + 1) \cdot \|\tilde{g}_k\|.$$

So we will be guaranteed to not overflow if

$$\max\left(1, L^{-1}\right) \cdot (2\kappa + 1) \cdot \|\tilde{g}_k\| \leq M_{\text{machine-lo-bias}(B)}.$$

We know from the construction of the HALP algorithm that

$$\begin{aligned}
M_{\text{machine-lo-bias}(B)} &= 2^B \cdot M_{\text{machine-lo}} \\
&= 2^{\lfloor \log_2(\zeta \cdot \|\tilde{g}_k\|) \rfloor} \cdot M_{\text{machine-lo}} \\
&\leq 2^{\log_2(\zeta \cdot \|\tilde{g}_k\|) - 1} \cdot M_{\text{machine-lo}} \\
&= \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot M_{\text{machine-lo}}.
\end{aligned}$$

Therefore we will be guaranteed to not overflow if

$$\frac{4\kappa + 2}{\zeta} \cdot \max\left(1, L^{-1}\right) \leq M_{\text{machine-lo}}.$$

Now, all these computations are actually being performed in floating-point, not exact arithmetic. So, they will see some relative error as a result. When we take this relative error into account, we get

$$\frac{4\kappa + 2}{\zeta} \cdot \max\left(1, L^{-1}\right) \cdot (1 + O(\varepsilon_{\text{machine-lo}})) \leq M_{\text{machine-lo}},$$

which is the desired result. This proves the theorem. $\blacksquare$

Proof of Theorem 10 We start by analyzing Assumption 8. This assumption says that

$$\left|\nabla_{\text{C}} f_i((\tilde{w}_k, \delta_{w,k,t})) - \nabla f_i(\tilde{w}_k + \delta_{w,k,t})\right|$$
$$= a_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \varepsilon_{\text{lo}} + c_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \eta_{\text{lo}} + c_{\text{lo-bias}} \cdot \eta_{\text{lo-bias}(B)}.$$

That is, the error in computing the delta is bounded by this expression. In terms of the computation in the HALP algorithm, we can write this as

$$\|\delta_{v,k,t} - (\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k))\| = a_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \varepsilon_{\text{lo}} + c_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \eta_{\text{lo}} + c_{\text{lo-bias}} \cdot \eta_{\text{lo-bias}(B)}.$$

Just as was the case for the analysis of HALP for linear regression, we have that the machine epsilon for the biased low-precision format is bounded by

$$\begin{aligned}
\eta_{\text{machine-lo-bias}(B)} &= 2^B \cdot \eta_{\text{machine-lo}} \\
&= 2^{\lfloor \log_2(\zeta \cdot \|\tilde{g}_k\|) \rfloor} \cdot \eta_{\text{machine-lo}} \\
&\leq 2^{\log_2(\zeta \cdot \|\tilde{g}_k\|) - 1} \cdot \eta_{\text{machine-lo}} \\
&= \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot \eta_{\text{machine-lo}},
\end{aligned}$$

so

$$\|\delta_{v,k,t} - (\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k))\| = a_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \varepsilon_{\text{lo}} + c_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \eta_{\text{lo}} + c_{\text{lo-bias}} \cdot \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot \eta_{\text{lo}}.$$

Since we know that

$$\|\delta_{w,k,t}\| \leq \tilde{s}_k = \frac{2}{\mu} \cdot \|\tilde{g}_k\|,$$

it follows that we can bound the error with

$$\|\delta_{v,k,t} - (\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k))\| = a_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \varepsilon_{\text{lo}} + \left(c_{\text{lo}} \cdot \frac{2}{\mu} + c_{\text{lo-bias}} \cdot \frac{\zeta}{2} \right) \cdot \|\tilde{g}_k\| \cdot \eta_{\text{lo}}.$$

To simplify this, define

$$c = c_{\text{lo}} \cdot \frac{2}{\mu} + c_{\text{lo-bias}} \cdot \frac{\zeta}{2}$$

and

$$u = \delta_{v,k,t} - (\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k))$$

Then,

$$\delta_{v,k,t} = \nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) + u$$

and

$$\|u\| \leq a_{\text{lo}} \cdot \|\delta_{w,k,t}\| \cdot \varepsilon_{\text{lo}} + c \cdot \|\tilde{g}_k\| \cdot \eta_{\text{lo}}$$

Next, we look at the inner update step of HALP. This update step is

$$\delta_{w,k,t} = \delta_{w,k,t-1} \underset{\text{L}}{\ominus} \alpha \underset{\text{L}}{\odot} \left(\delta_{v,k,t} \underset{\text{L}}{\oplus} \tilde{h}_k \right).$$

67

Analyzing its error in the jth coordinate, we get

$$
\begin{aligned}
\delta_{w,k,t,j} &= \delta_{w,k,t-1,j} \ominus_{\mathrm{L}} \alpha \odot_{\mathrm{L}} \left(\delta_{v,k,t,j} \oplus_{\mathrm{L}} \tilde{h}_{k,j} \right) \\
&= \Bigg(\delta_{w,k,t-1,j} - \Big(\alpha \cdot \Big(\Big(\big(\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) + u \big)_j \\
&\quad + \tilde{g}_{k,j} \cdot (1+\varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \Big) \cdot (1+\varepsilon_{\mathrm{lo}}) \\
&\qquad + \eta_{\mathrm{lo\text{-}bias}(B)} \Big) \cdot (1+\varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \Big) \Bigg) \cdot (1+\varepsilon_{\mathrm{lo}}) + \eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \delta_{w,k,t-1,j} \cdot (1+\varepsilon_{\mathrm{lo}}) \\
&\quad - \Big(\alpha \cdot \Big(\big(\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) \big)_j \cdot (1+3\varepsilon_{\mathrm{lo}}) + u_j \Big) \\
&\qquad + \alpha \cdot \tilde{g}_{k,j} \cdot (1+4\varepsilon_{\mathrm{lo}}) + 2\alpha \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \Big) + 2\eta_{\mathrm{lo\text{-}bias}(B)} \\
&= \delta_{w,k,t-1,j} - \alpha \cdot \big(\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) + \tilde{g}_k \big)_j + \delta_{w,k,t-1,j} \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \Big(\alpha \cdot \Big(\big(\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) \big)_j \cdot 3\varepsilon_{\mathrm{lo}} + u_j \Big) \\
&\qquad + \alpha \cdot \tilde{g}_{k,j} \cdot 4\varepsilon_{\mathrm{lo}} + 2\alpha \cdot \eta_{\mathrm{lo\text{-}bias}(B)} \Big) + 2\eta_{\mathrm{lo\text{-}bias}(B)}.
\end{aligned}
$$

So, the error of computing the inner update step will be

$$
\begin{aligned}
\Delta &= \big\| \delta_{w,k,t} - \big(\delta_{w,k,t-1} - \alpha \cdot \big(\nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) + \tilde{g}_k \big) \big) \big\| \\
&\leq \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + \alpha \cdot \| \nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) \| \cdot 3\varepsilon_{\mathrm{lo}} + \alpha \cdot \| u \| \\
&\quad + \alpha \cdot \| \tilde{g}_k \| \cdot 4\varepsilon_{\mathrm{lo}} + 2\alpha \cdot \sqrt{d}\eta_{\mathrm{lo\text{-}bias}(B)} + 2\sqrt{d}\eta_{\mathrm{lo\text{-}bias}(B)} \\
&\leq \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + \alpha \cdot \| \nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) \| \cdot 3\varepsilon_{\mathrm{lo}} + \alpha \cdot \big(a_{\mathrm{lo}} \cdot \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + c \cdot \| \tilde{g}_k \| \cdot \eta_{\mathrm{lo}} \big) \\
&\quad + \alpha \cdot \| \tilde{g}_k \| \cdot 4\varepsilon_{\mathrm{lo}} + 2\alpha \cdot \sqrt{d}\eta_{\mathrm{lo\text{-}bias}(B)} + 2\sqrt{d}\eta_{\mathrm{lo\text{-}bias}(B)}.
\end{aligned}
$$

Since ∇f_i is Lipschitz continuous,

$$
\| \nabla f_i(w_{k,t}) - \nabla f_i(\tilde{w}_k) \| \leq L \cdot \| w_{k,t} - \tilde{w}_k \| = L \cdot \| \delta_{w,k,t} \|.
$$

So, applying this and our earlier bound on $\eta_{\mathrm{lo\text{-}bias}(B)}$,

$$
\begin{aligned}
\Delta &= \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + \alpha L \cdot \| \delta_{w,k,t-1} \| \cdot 3\varepsilon_{\mathrm{lo}} + \alpha \cdot \big(a_{\mathrm{lo}} \cdot \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + c \cdot \| \tilde{g}_k \| \cdot \eta_{\mathrm{lo}} \big) \\
&\quad + \alpha \cdot \| \tilde{g}_k \| \cdot 4\varepsilon_{\mathrm{lo}} + \alpha \cdot \sqrt{d} \cdot \zeta \cdot \| \tilde{g}_k \| \cdot \eta_{\mathrm{lo}} + \sqrt{d} \cdot \zeta \cdot \| \tilde{g}_k \| \cdot \eta_{\mathrm{lo}} \\
&= (1 + 3\alpha L + \alpha a_{\mathrm{lo}}) \cdot \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + 4\alpha \cdot \| \tilde{g}_k \| \cdot \varepsilon_{\mathrm{lo}} + \big(\alpha c + (1+\alpha)\sqrt{d} \cdot \zeta \big) \| \tilde{g}_k \| \cdot \eta_{\mathrm{lo}}.
\end{aligned}
$$

Next, leveraging the fact that by Lipschitz contiuity, $\| \tilde{g}_k \| \leq L \cdot \| \tilde{w}_k - w^* \|$,

$$
\begin{aligned}
\Delta &\leq (1 + 3\alpha L + \alpha a_{\mathrm{lo}}) \cdot \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + 4\alpha L \cdot \| \tilde{w}_k - w^* \| \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \big(\alpha L c + (L+\alpha L)\sqrt{d} \cdot \zeta \big) \| \tilde{w}_k - w^* \| \cdot \eta_{\mathrm{lo}}.
\end{aligned}
$$

We can simplify this by using the assumption that $4\alpha L \leq 1$, which we will validate by our choice of α later. Using this, we get

$$
\begin{aligned}
\Delta &\leq \left(2 + \frac{a_{\mathrm{lo}}}{4L} \right) \cdot \| \delta_{w,k,t-1} \| \cdot \varepsilon_{\mathrm{lo}} + \| \tilde{w}_k - w^* \| \cdot \varepsilon_{\mathrm{lo}} \\
&\quad + \frac{1}{4}\big(c + (4L+1)\sqrt{d} \cdot \zeta \big) \| \tilde{w}_k - w^* \| \cdot \eta_{\mathrm{lo}}.
\end{aligned}
$$

Leveraging our assumption that

$$(c + (4L + 1)\sqrt{d} \cdot \zeta) \cdot \eta_{\text{machine-lo}} \leq \varepsilon_{\text{machine-lo}},$$

and recalling that $\delta_{w,k,t-1} = w_{k,t-1} - \tilde{w}$, we get

$$\begin{aligned}
\Delta &\leq \left(2 + \frac{a_{\text{lo}}}{4L}\right) \cdot \|\delta_{w,k,t-1}\| \cdot \varepsilon_{\text{lo}} + 2\|\tilde{w}_k - w^*\| \cdot \varepsilon_{\text{lo}} \\
&\leq \left(2 + \frac{a_{\text{lo}}}{4L}\right) \cdot (\|w_{k,t-1} - \tilde{w}\| + \|\tilde{w}_k - w^*\|) \cdot \varepsilon_{\text{lo}} \\
&\leq \left(2 + \frac{a_{\text{lo}}}{4L}\right) \cdot (\|w_{k,t-1} - \tilde{w}\| + \|\tilde{w}_k - w^*\|) \cdot \varepsilon_{\text{machine-lo}}.
\end{aligned}$$

Finally, we are ready to apply Lemma 12, which gives us

$$\begin{aligned}
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] &\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 2\Delta\left((1 + \alpha^2 L^2)\|w_{k,t-1} - w^*\| + \sqrt{2}\alpha L \|\tilde{w}_k - w^*\|\right) + \Delta^2,
\end{aligned}$$

where Δ is as defined above. Note that this will hold even if we take the if-statement to reset the value of $\delta_{w,k,t}$ to 0, by the same logic used in the analysis of HALP for linear regression. Subject to our assumption that $4\alpha L \leq 1$, we can simplify this to

$$\begin{aligned}
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] &\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 3\Delta (\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\|) + \Delta^2.
\end{aligned}$$

Next, we substitute our above bound on Δ, which produces

$$\begin{aligned}
\mathbf{E}\left[\|w_{k,t} - w^*\|^2\right] &\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 3\left(2 + \frac{a_{\text{lo}}}{4L}\right) \cdot (\|w_{k,t-1} - w^*\| + \|\tilde{w}_k - w^*\|)^2 \cdot \varepsilon_{\text{machine-lo}} \\
&\quad + \left(\left(2 + \frac{a_{\text{lo}}}{4L}\right) \cdot (\|w_{k,t-1} - \tilde{w}\| + \|\tilde{w}_k - w^*\|) \cdot \varepsilon_{\text{machine-lo}}\right)^2 \\
&\leq \|w_{k,t-1} - w^*\|^2 - 2\alpha(1 - 2\alpha L)(f(w_{k,t-1}) - f(w^*)) + 4\alpha^2 L(f(\tilde{w}_k) - f(w^*)) \\
&\quad + 6\left(2 + \frac{a_{\text{lo}}}{4L}\right) \cdot \left(\|w_{k,t-1} - w^*\|^2 + \|\tilde{w}_k - w^*\|^2\right) \cdot \left(\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2)\right).
\end{aligned}$$

Notice that this is the exact same expression we obtain in the proof of convergence of BC-SVRG (Theorem 4), except with

$$2 + \frac{a_{\text{lo}}}{4L}$$

in the place of d in that expression. By the exact same argument we used in that proof, we will get

$$\mathbf{E}\left[f(\tilde{w}_{k+1}) - f(w^*)\right] \leq \left(\gamma + \frac{192\kappa}{\gamma} \cdot \left(2 + \frac{a_{\text{lo}}}{4L}\right) \cdot (\varepsilon_{\text{machine-lo}} + O(\varepsilon_{\text{machine-lo}}^2))\right) \mathbf{E}\left[f(\tilde{w}_k) - f(w^*)\right].$$

Note that this still holds even if we take the if-statement to reset the value of $\tilde{w}_{k+1}$ to be equal to $\tilde{w}_k$, since this operation decreases the value of $f(\tilde{w}_{k+1}) - f(w^*)$. This is what we wanted to prove.

All that remains to be done is to show that there is no overflow. To do this, consider Assumption 9: all the numbers in the inner loop will satisfy

$$|z| \leq r_{\text{delta}} \cdot \|\delta_{w,k,t}\| + r_{\text{lo-bias}} \cdot \eta_{\text{lo-bias}(B)}.$$

By our above bound on $\eta_{\text{lo-bias}(B)}$, we have that

$$|z| \leq r_{\text{delta}} \cdot \|\delta_{w,k,t}\| + r_{\text{lo-bias}} \cdot \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot \eta_{\text{machine-lo}}.$$

Because we do not let $\delta_{w,k,t}$ get too large, it will always hold that

$$\|\delta_{w,k,t}\| \leq \tilde{s} = \frac{2}{\mu} \cdot \|\tilde{g}_k\|,$$

so

$$|z| \leq r_{\text{delta}} \cdot \frac{2}{\mu} \cdot \|\tilde{g}_k\| + r_{\text{lo-bias}} \cdot \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot \eta_{\text{machine-lo}}.$$

The overflow threshold for our biased low-precision representation is

$$M_{\text{machine-lo-bias}(B)} = 2^B \cdot M_{\text{machine-lo}} \leq \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot M_{\text{machine-lo}}.$$

So we will be guaranteed to have no overflow if

$$r_{\text{delta}} \cdot \frac{2}{\mu} \cdot \|\tilde{g}_k\| + r_{\text{lo-bias}} \cdot \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot \eta_{\text{machine-lo}} \leq \frac{1}{2} \cdot \zeta \cdot \|\tilde{g}_k\| \cdot M_{\text{machine-lo}}.$$

This will happen when

$$r_{\text{delta}} \cdot \frac{4}{\zeta \mu} + r_{\text{lo-bias}} \cdot \eta_{\text{machine-lo}} \leq M_{\text{machine-lo}}.$$

This completes the theorem. $\blacksquare$

Proof of Theorem 7 To simplify our notation, let

$$f = f(g_1(x_1, \ldots, x_n), g_2(x_1, \ldots, x_n), \ldots, g_m(x_1, \ldots, x_n)),$$

$$\tilde{f} = o_f + \delta_f = \underset{\text{C}}{f}(\underset{\text{C}}{g_1}(x_1, \ldots, x_n), \underset{\text{C}}{g_2}(x_1, \ldots, x_n), \ldots, \underset{\text{C}}{g_m}(x_1, \ldots, x_n))$$

$$g_i = g_i(x_1, \ldots, x_n),$$

and

$$\tilde{g}_i = o_{g,i} + \delta_{g,i} = \underset{\text{C}}{g_i}(x_1, \ldots, x_n).$$

Then

$$\left| \underset{\text{C}}{h}(x_1, \ldots, x_n) - h(x_1, \ldots, x_n) \right| = \left| \tilde{f} - f \right|.$$

By the fundamental axiom on g_i, we will have

$$|\tilde{g}_i - g_i| \leq A_{\text{hi},g,i} \cdot \varepsilon_{\text{machine-hi}} + C_{\text{hi},g,i} \cdot \eta_{\text{machine-hi}} + A_{\text{lo},g,i} \cdot \left(\sum_{i=1}^{m} |\delta_{x,i}| \right) \cdot \varepsilon_{\text{machine-lo}}$$

$$+ C_{\text{lo-bias},g,i} \cdot \left(\sum_{i=1}^{m} |\delta_{x,i}| \right) \cdot \eta_{\text{machine-lo-bias}(B)} + C_{\text{lo},g,i} \cdot \eta_{\text{machine-lo}}$$

$$\leq A_{\text{hi},g,i} \cdot \varepsilon_{\text{machine-hi}} + C_{\text{hi},g,i} \cdot \eta_{\text{machine-hi}} + A_{\text{lo},g,i} \cdot \left(\sum_{i=1}^{m} |x_i - o_{x,i}| \right) \cdot \varepsilon_{\text{machine-lo}}$$

$$+ C_{\text{lo-bias},g,i} \cdot \left(\sum_{i=1}^{m} |x_i - o_{x,i}| \right) \cdot \eta_{\text{machine-lo-bias}(B)} + C_{\text{lo},g,i} \cdot \eta_{\text{machine-lo}}.$$

Define a new function D_i (a function of the same parameters as $A_{\text{hi},g,i}$), as

$$D_i = A_{\text{hi},g,i} \cdot \varepsilon_{\text{machine-hi}} + C_{\text{hi},g,i} \cdot \eta_{\text{machine-hi}} + A_{\text{lo},g,i} \cdot \left(\sum_{i=1}^{m} |x_i - o_{x,i}| \right) \cdot \varepsilon_{\text{machine-lo}}$$

$$+ C_{\text{lo-bias},g,i} \cdot \left(\sum_{i=1}^{m} |x_i - o_{x,i}| \right) \cdot \eta_{\text{machine-lo-bias}(B)} + C_{\text{lo},g,i} \cdot \eta_{\text{machine-lo}}.$$

It follows that, if for some $\alpha_i \in [-1, 1]$,

$$\tilde{g}_i \leq g_i + \alpha_i D_i.$$

Next, by the fundamental axiom on f, we will have

$$\left| \tilde{f} - f \right| \leq A_{\text{hi},f}(\tilde{g}_1, o_{g,1}, \ldots, \tilde{g}_m, o_{g,m}, \ldots) \cdot \varepsilon_{\text{machine-hi}}$$

$$+ C_{\text{hi},f}(\tilde{g}_1, o_{g,1}, \ldots, \tilde{g}_m, o_{g,m}, \ldots) \cdot \eta_{\text{machine-hi}}$$

$$+ A_{\text{lo},f}(\tilde{g}_1, o_{g,1}, \ldots, \tilde{g}_m, o_{g,m}, \ldots) \cdot \left(\sum_{i=1}^{m} |\delta_{g,i}| \right) \cdot \varepsilon_{\text{machine-lo}}$$

$$+ C_{\text{lo-bias},f}(\tilde{g}_1, o_{g,1}, \ldots, \tilde{g}_m, o_{g,m}, \ldots) \cdot \left(\sum_{i=1}^{m} |\delta_{g,i}| \right) \cdot \eta_{\text{machine-lo-bias}(B)}$$

$$+ C_{\text{lo},f}(\tilde{g}_1, o_{g,1}, \ldots, \tilde{g}_m, o_{g,m}, \ldots) \cdot \eta_{\text{machine-lo}}.$$

From our analysis above, there exists $\alpha_i \in [-1, 1]$ such that

$$\left| \tilde{f} - f \right| \leq A_{\text{hi},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots, g_m + \alpha_m D_m, o_{g,m}, \ldots) \cdot \varepsilon_{\text{machine-hi}}$$

$$+ C_{\text{hi},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots, g_m + \alpha_m D_m, o_{g,m}, \ldots) \cdot \eta_{\text{machine-hi}}$$

$$+ A_{\text{lo},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots, g_m + \alpha_m D_m, o_{g,m}, \ldots) \cdot \left(\sum_{i=1}^{m} |\delta_{g,i}| \right) \cdot \varepsilon_{\text{machine-lo}}$$

$$+ C_{\text{lo-bias},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots, g_m + \alpha_m D_m, o_{g,m}, \ldots) \cdot \left(\sum_{i=1}^{m} |\delta_{g,i}| \right) \cdot \eta_{\text{machine-lo-bias}(B)}$$

$$+ C_{\text{lo},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots, g_m + \alpha_m D_m, o_{g,m}, \ldots) \cdot \eta_{\text{machine-lo}}.$$

It follows that

$$
\begin{aligned}
\left| \tilde{f} - f \right| \leq & \left(\min_{\alpha \in [-1,1]^m} A_{\mathrm{hi},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots) \right) \cdot \varepsilon_{\text{machine-hi}} \\
& + \left(\min_{\alpha \in [-1,1]^m} C_{\mathrm{hi},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots) \right) \cdot \eta_{\text{machine-hi}} \\
& + \left(\min_{\alpha \in [-1,1]^m} A_{\mathrm{lo},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots) \right) \cdot \left(\sum_{i=1}^m |\delta_{g,i}| \right) \cdot \varepsilon_{\text{machine-lo}} \\
& + \left(\min_{\alpha \in [-1,1]^m} C_{\mathrm{lo\text{-}bias},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots) \right) \cdot \left(\sum_{i=1}^m |\delta_{g,i}| \right) \cdot \eta_{\text{machine-lo-bias}(B)} \\
& + \left(\min_{\alpha \in [-1,1]^m} C_{\mathrm{lo},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots) \right) \cdot \eta_{\text{machine-lo}}.
\end{aligned}
$$

If we define $A_{\mathrm{hi},\hat{h}}$ to be

$$
A_{\mathrm{hi},\hat{h}} = \min_{\alpha \in [-1,1]^m} A_{\mathrm{hi},f}(g_1 + \alpha_1 D_1, o_{g,1}, \ldots),
$$

and similarly for the other functions, then we get

$$
\begin{aligned}
\left| \tilde{f} - f \right| \leq & \; A_{\mathrm{hi},\hat{h}} \cdot \varepsilon_{\text{machine-hi}} \\
& + C_{\mathrm{hi},\hat{h}} \cdot \eta_{\text{machine-hi}} \\
& + A_{\mathrm{lo},\hat{h}} \cdot \left(\sum_{i=1}^m |\delta_{g,i}| \right) \cdot \varepsilon_{\text{machine-lo}} \\
& + C_{\mathrm{lo\text{-}bias},\hat{h}} \cdot \left(\sum_{i=1}^m |\delta_{g,i}| \right) \cdot \eta_{\text{machine-lo-bias}(B)} \\
& + C_{\mathrm{lo},\hat{h}} \cdot \eta_{\text{machine-lo}}.
\end{aligned}
$$

Note that $A_{\mathrm{hi},h}$ must be continuous since it is the maximum of a continuous function over a bounded domain, and it will be a function of the same parameters as $A_{\mathrm{hi},g,i}$. Next, from the fundamental axiom on g_i, we will also have

$$
|\delta_{g,i}| \leq R_{\mathrm{delta},g,i} \cdot \left(\sum_{i=1}^m |\delta_{x,i}| \right) + R_{\mathrm{lo\text{-}bias},g,i} \cdot \eta_{\mathrm{lo\text{-}bias}(B)}.
$$

Combining this with our previous expression, it is straightfoward to see that h satisfies the fundamental axiom of bit centered arithmetic. ∎

References

Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. QSGD: Communication-efficient SGD via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*, pages 1707–1718, 2017.

Zeyuan Allen-Zhu and Elad Hazan. Variance reduction for faster non-convex optimization. In *International Conference on Machine Learning*, pages 699–707, 2016.

David Bindel. CS 6210 course notes: Dot products in floating point. `http://www.cs.cornell.edu/~bindel/class/cs6210-f12/notes/lec05.pdf`, 2012.

Doug Burger. Microsoft unveils Project Brainwave for real-time ai. `https://www.microsoft.com/en-us/research/blog/microsoft-unveils-project-brainwave/`, 2017. Accessed: 2018-02-08.

Adrian M Caulfield, Eric S Chung, Andrew Putnam, Hari Angepat, Daniel Firestone, Jeremy Fowers, Michael Haselman, Stephen Heil, Matt Humphrey, Puneet Kaur, et al. Configurable clouds. *IEEE Micro*, 37(3):52–61, 2017.

Intel Corporation. Bfloat16: a hardware numerics definition. `https://software.intel.com/sites/default/files/managed/40/8b/bf16-hardware-numerics-definition-white-paper.pdf`, 2018. Accessed: 2018-12-1.

Matthieu Courbariaux, Jean-Pierre David, and Yoshua Bengio. Training deep neural networks with low precision multiplications. *arXiv preprint arXiv:1412.7024*, 2014.

Christopher De Sa, Ce Zhang, Kunle Olukotun, and Christopher Ré. Taming the wild: A unified analysis of HOGWILD!-style algorithms. In *NIPS*, 2015.

Christopher De Sa, Matthew Feldman, Christopher Ré, and Kunle Olukotun. Understanding and optimizing asynchronous low-precision stochastic gradient descent. *ISCA 2017*, 2017.

Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. *ICML*, 2015.

Reza Harikandeh, Mohamed Osama Ahmed, Alim Virani, Mark Schmidt, Jakub Konečnỳ, and Scott Sallinen. Stop wasting my gradients: Practical SVRG. In *Advances in Neural Information Processing Systems*, pages 2251–2259, 2015.

Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. In *Advances in neural information processing systems*, pages 4107–4115, 2016.

IEEE. IEEE standard for floating-point arithmetic - redline. *IEEE Std 754-2008 (Revision of IEEE Std 754-1985)*, pages 1–82, Aug 2008. doi: 10.1109/IEEESTD.2008.5976968.

Xianyan Jia, Shutao Song, Wei He, Yangzihao Wang, Haidong Rong, Feihu Zhou, Liqiang Xie, Zhenyu Guo, Yuanzhou Yang, Liwei Yu, et al. Highly scalable deep learning training system with mixed-precision: Training imagenet in four minutes. *arXiv preprint arXiv:1807.11205*, 2018.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323, 2013.

Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pages 1–12. ACM, 2017.

Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The CIFAR-10 dataset. `http://www.cs.toronto.edu/kriz/cifar.html`, 2014.

Liu Kuang. Train CIFAR10 with PyTorch. `https://github.com/kuangliu/pytorch-cifar`, 2018.

Yann LeCun. The MNIST database of handwritten digits. `http://yann.lecun.com/exdb/mnist/`, 1998.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Hao Li, Soham De, Zheng Xu, Christoph Studer, Hanan Samet, and Tom Goldstein. Training quantized nets: A deeper understanding. In *Advances in Neural Information Processing Systems*, pages 5813–5823, 2017.

Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. Building a large annotated corpus of English: The Penn Treebank. *COMPUTATIONAL LINGUISTICS*, 19(2):313–330, 1993.

Paulius Micikevicius. Mixed-precision training of deep neural networks. `https://devblogs.nvidia.com/mixed-precision-training-deep-neural-networks/`, 2017. Accessed: 2018-12-1.

Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory F. Diamos, Erich Elsen, David García, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed precision training. *CoRR*, abs/1710.03740, 2017a. URL `http://arxiv.org/abs/1710.03740`.

Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory F. Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed precision training. abs/1710.03740, 2017b.

Nvidia. Deep learning sdk documentation. `https://docs.nvidia.com/deeplearning/sdk/mixed-precision-training/index.html`, 2018. Accessed: 2018-12-1.

Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, pages 525–542. Springer, 2016.

Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alex Smola. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*, pages 314–323, 2016.

Nicolas L Roux, Mark Schmidt, and Francis R Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In *Advances in Neural Information Processing Systems*, pages 2663–2671, 2012.

Antony W Savich and Medhat Moussa. Resource efficient arithmetic effects on RBM neural network solution quality using MNIST. In *2011 International Conference on Reconfigurable Computing and FPGAs*, pages 35–40. IEEE, 2011.

Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs. In *INTERSPEECH*, pages 1058–1062, 2014.

Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(Feb):567–599, 2013.

Nikko Strom. Scalable distributed DNN training using commodity GPU cloud computing. In *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.

Hanlin Tang, Ce Zhang, Shaoduo Gan, Tong Zhang, and Ji Liu. Decentralization meets quantization. *CoRR*, abs/1803.06443, 2018. URL `http://arxiv.org/abs/1803.06443`.

Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. Siam, 1997.

Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, and Kailash Gopalakrishnan. Training deep neural networks with 8-bit floating point numbers. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7685–7694. Curran Associates, Inc., 2018. URL `http://papers.nips.cc/paper/7994-training-deep-neural-networks-with-8-bit-floating-point-numbers.pdf`.

Hantian Zhang, Jerry Li, Kaan Kara, Dan Alistarh, Ji Liu, and Ce Zhang. Zipml: Training linear models with end-to-end low precision, and a little bit of deep learning. In *International Conference on Machine Learning*, pages 4035–4043, 2017.