
End-to-End DNN Training with Block Floating Point Arithmetic

Mario Drumond¹ Tao Lin¹ Martin Jaggi¹ Babak Falsafi¹

Abstract

DNNs are ubiquitous datacenter workloads, requiring orders of magnitude more computing power from servers than traditional workloads. As such, datacenter operators are forced to adopt domain-specific accelerators that employ half-precision floating-point (FP) numeric representations to improve arithmetic density. Unfortunately, even these representations are not dense enough, and are, therefore, sub-optimal for DNNs. We propose a hybrid approach that employs dense block floating-point (BFP) arithmetic on dot product computations and FP arithmetic elsewhere. While using BFP improves the performance of dot product operations, that compose most of DNN computations, allowing values to freely float between dot product operations leads to a better choice of tensor exponents when converting values to back BFP. We show that models trained with hybrid BFP-FP arithmetic either match or outperform their FP32 counterparts, leading to more compact models and denser arithmetic in computing platforms.

1. Introduction

Today’s ubiquitous online services are often driven by DNNs to provide custom-tailored user content. Delivering faster inference and more accurate training, however, is often limited by the arithmetic density of the underlying hardware platform. Most users resort to graphics processing units (GPUs) as the platform of choice for training neural networks because GPUs offer high arithmetic density per silicon area through full precision floating-point (FP32) units. However, even traditional GPUs have proved not to have dense enough arithmetic for DNNs, and accelerators are moving to narrow arithmetic to improve logic density. For instance, NVIDIA’s Volta (nvi, 2018) and Google’s TPU2 architectures employ half-precision floating point (FP16) arithmetic.

Unfortunately, optimizing floating-point – even narrow FP16 logic – has been a daunting task for device designers. Sequential implementation of floating-point logic is quite slow and parallelizing the logic is prohibitively resource intensive compromising density. A promising solution to this

problem is to utilize fixed-point arithmetic, which promises great gains in both speed and density; unfortunately, performing training with fixed-point networks has been unsuccessful to this point due to the lack of dynamic range inherent in the fixed-point representation.

Signal processing platforms have historically resorted to block floating-point (BFP), whose representation is shown in Figure 1, as a way to optimize for both performance and density. The use of BFP has allowed signal processors to convert common algorithms (e.g., FFT) to dense and parallel integer arithmetic hardware. We observe that BFPs are also likely to be effective in neural networks, increasing the arithmetic density of accelerators and improving the dynamic range of fixed-point-like arithmetic taking the first step towards effective training in dense arithmetic. Naive application of BFP to DNN training, however, is not straight forward. Tensor values often drift during training requiring a new choice of exponent – or quantization points.

In this paper, we make the observation that in DNNs, the majority of the arithmetic operations executed are performed as part of dot product calculations, and therefore, limiting dense fixed-point-like arithmetic to only replacing the dot products still allows us to accelerate the majority of the network. As such, the rest of the operations can be implemented in traditional floating-point logic with little performance degradation. We propose a hybrid BFP-FP framework where values float freely between dot product computations in BFP, resulting in better choice of exponents, and perform the rest of the training in traditional floating-point arithmetic. Hybrid BFP-FP training also makes the underlying hardware more friendly to users, who can use complex arithmetic, undisturbed by limitations imposed by BFP implementations.

The separation between dot products and other operations already exists in commodity hardware in NVIDIA Volta’s FP16 Tensor Cores (nvi, 2018) and in Google’s inference-only, fixed-point based accelerator, Tensor Processing Unit (Jouppi et al., 2017) architecture. We just take one step further and use different numeric representations for these different operations. Hybrid BFP-FP representations enable a new class of efficient accelerators transparently implementing dense arithmetic for DNN while maintaining usability.

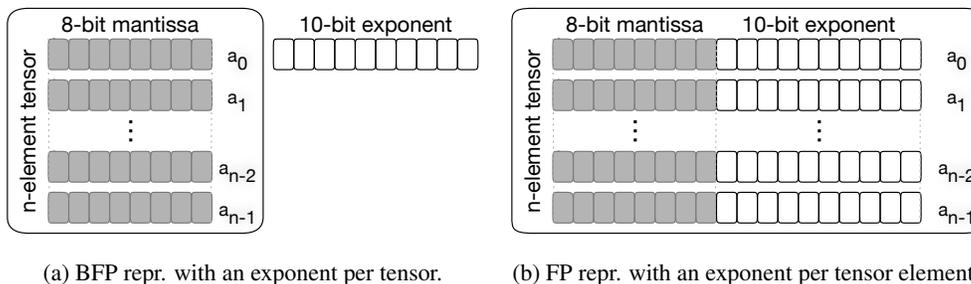


Figure 1. A n -element tensor in BFP and FP representations. BFP tensors save space and simplify computations by sharing exponents across tensors.

This paper’s contributions are: (1) a hybrid BFP-FP DNN training framework to optimize the quantization points while maximizing fixed-point arithmetic in dot products, and (2) an exploration of the design space showing that DNNs trained on BFP with 12- and 8-bit mantissas not only match the quality of DNNs trained with FP32 but sometimes surpass them.

2. Related Work

Inference with reduced precision. Quantization (Qua, 2017) is a widely used technique for DNN inference. BFP (Song et al., 2017) has also been proposed for inference. These proposals take DNNs trained using full precision floating-point and quantize their weights in order to use cheap fixed-point logic during inference. These DNNs often have to be retrained with quantized weights to recover precision. Quantized inference takes advantage of the fact that, at inference time, weights are static, and so are their exponents. Unfortunately, it is not clear how to derive gradients of DNNs with quantized weights. Quantized-inference techniques also cannot be used for training. We introduce a technique to train DNNs with performance that matches quantized inference.

Binarized and Ternary Neural Networks. Binarized (Courbariaux et al., 2016) and Ternary (Li & Liu, 2016; Zhu et al., 2016) neural networks are another way to compress models. Although these networks require hardware that is orders of magnitude simpler for inference, they are trained in a similar way to traditional neural networks, with both activations and parameters represented with floating-point. Therefore, these approaches are orthogonal to BFP-based training, because BFP is meant as a replacement for floating-point during training.

Training with binary operations in forward and backward passes. (Courbariaux et al., 2015; Rastegari et al., 2016) use binary operations for forward and backward passes but not for calculating and accumulating weight gradients. Their approach redesigns SGD and is not transparent

to users, requiring redesign of networks with numeric representation in mind.

Training with low bit-width gradients. QSGD (Alistarh et al., 2017) is a compression framework for SGD gradients that seeks to reduce communication bandwidth requirements for distributed SGD implementations. BFP can also be seen as a way to reduce the communication requirements of SGD, as it reduces the number of bits used to represent each number by removing exponents of individual values.

Training with end-to-end low precision. ZipML (Zhang et al., 2017), DoReFa (Zhou et al., 2016), and Flexpoint (Köster et al., 2017) introduce methods to train models using end-to-end low precision. They use fixed-point arithmetic to represent weights and activations during forward and backward passes. DoReFa (Zhou et al., 2016) requires techniques to control the activations’ magnitudes, and is unable to quantize the first and last layers of networks. We use BFP for its wide dynamic range of representable values, obviating the need for controlling the magnitude of the values of both activations and gradients.

ZipML (Zhang et al., 2017) takes a more theoretical approach to find the optimal quantization points for each dataset, performing both computations and communication using fixed-point arithmetic. We use BFP instead, effectively computing quantization points by choosing tensor exponents, and doing so at a finer granularity, at conversion time for activations and update time for weights.

Flexpoint (Köster et al., 2017) uses a BFP number representation. However, their training method adds complexity to SGD. During training, they calculate quantization points using the Autoflex algorithm, and perform all operations in fixed-point logic. We calculate quantization points on-the-fly every time tensors are converted to BFP.

We observe – and verify through an FPGA prototype – that, as long as dot product calculation’s intermediate values remain in fixed-point-like representations, conversions are infrequent enough that the hardware area dedicated to con-

versions accounts for a small fraction of the total accelerator area.

Training with half precision floating-point. Half precision floating-point (FP16) (Dally, 2015) is quickly becoming the state-of-the-art for neural networks training, with both Google’s TPU2 (goo, 2017) and NVIDIA’s Volta (nvi, 2018) GPUs adopting half-precision floating-point as their arithmetic representation. However, FP16 suffers from limited range, and it often requires weights and gradients to be scaled in order to converge. Also, FP16 incurs larger area and power requirements in hardware. BFP solves this problem by sharing exponents across matrices, enabling the usage of exponents with large bit-widths with little communication overheads, preserving large dynamic range and obviating gradient scaling techniques.

3. Specialized Arithmetic for DNNs

Due to the massive computational requirements for DNNs when employed in datacenter scale online services, operators such as Google started adopting specialized numeric representations for DNNs. So far, accelerators have employed fixed-point for inference (Jouppi et al., 2017), and narrow floating-point representations, such as FP16 (goo, 2017; nvi, 2018), for training. From a hardware design perspective, the use of reduced precision arithmetic allows silicon designers to improve logic density and energy-efficiency, while minimizing the number of bits used to represent models relaxes demands on both memory capacity and bandwidth. From the user’s perspective, arithmetic representations must be usable, not resulting in accuracy deterioration for models, nor requiring novel algorithmic techniques to recover model performance.

FP32 representations are usable but inefficient. They represent numbers with a 24-bit wide mantissa and a 8-bit wide exponent. In terms of precision, the 24-bit mantissa used in FP32 is overkill for DNNs. Figure 2a shows the training loss and Table 1 shows the test error of a ResNet-20 model trained on CIFAR-10 with truncated mantissas. The model converges even with a 4-bit mantissa, achieving best performance with 16-bit mantissas, and failing to converge only with 1-bit mantissas. In contrast, while the 8-bit exponent provides an appropriate dynamic range as shown in Figure 2b, training already suffers with a 6-bit exponent, and completely fails to converge with a 2-bit exponent. Silicon implementations of the the mantissa-exponent encoding normalize output mantissa of every operation. Normalization is implemented by a shifter in silicon, an expensive hardware structure in terms of area and power.

Using FP16 mitigates the area issues of FP32, employing narrow 11-bit mantissas and 5-bit exponents. However, FP16 is still expensive compared to fixed-point logic. For

Table 1. Test error of ResNet-20 on CIFAR-10 with narrow FP representations

Mantissa bit-widths				
2	4	8	16	24
9.77%	8.22%	8.05%	7.97%	8.42%
Exponent bit-widths				
2	6	8		
-%	14.67%	8.42%		

instance, although the area of an FP16 multiplier is $4.7\times$ smaller than that of a FP32 multiplier in 45nm manufacturing process node (Dally, 2015), it is $13\times$ larger than its 8-bit fixed-point counterpart. FP16 is also notoriously difficult to use, as the 5-bit exponent results in narrow dynamic range that is not sufficient to represent gradients throughout the training process. As such, from a usability perspective, the numeric representation must have wide dynamic range. Dynamic range is important during the training process, as the loss value decreases, and the gradient values also decrease.

Given these requirements, we identify block floating-point (BFP) as the ideal numeric representation for DNNs. BFP represents numbers with a mantissa and exponent, like floating-point, but exponents are shared across entire tensors, as shown in Figure 1, resulting in dot products that can be computed entirely in fixed-point logic. Since over 99% of the arithmetic operations executed by DNN training and inference are dot product computations, we are able to fold almost all the DNNs’ computations into fixed-point logic.

4. DNN Training using BFP Arithmetic

4.1. Using BFP in DNNs computation

Equation (1) computes the real value a_i of an element i of a BFP tensor a with mantissa m_i^a and exponent e_a .

$$a_i = m_i^a \times 2^{e_a} \quad (1)$$

Equation (2) calculates the dot product between BFP tensors a and b , each with N elements.

$$\begin{aligned} a \cdot b &= \sum_{i=1}^N \left((m_i^a \times 2^{e_a}) \times (m_i^b \times 2^{e_b}) \right) \\ &= 2^{e_a+e_b} \times (m^a \cdot m^b) \end{aligned} \quad (2)$$

The dot product $m^a \cdot m^b$ is computed entirely with fixed-point arithmetic, without the alignment of intermediate values, since all elements m_i^a and m_i^b are fixed-point. Thus, BFP dot products can only be computed with fixed-point

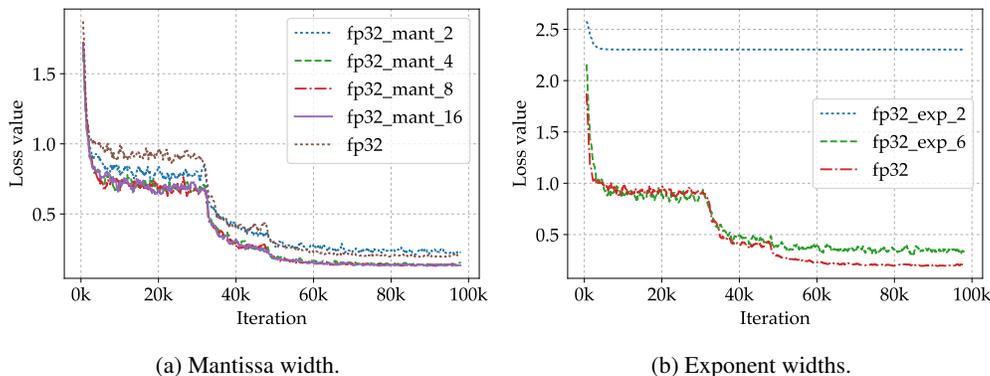


Figure 2. ResNet-20 training loss on CIFAR-10 with various floating-point configurations.

arithmetic if the entire sub-tensors that take part in dot products share the same exponent.

In a DNN’s fully-connected layers, this requirement translates to one exponent per activation tensor and weight matrix column in the forward pass, and one exponent per activation gradient tensor and weight matrix row in the backward pass. Since storing the weight matrix in two views (with both per-row and per-column exponent) is not possible, we use a single exponent for the entire weight matrix. The requirements are similar for convolutions: one exponent per activation input and kernel matrix. When computing weight gradients, the dot products are computed across batches, and therefore, entire batches of activations and gradients must share exponents to take advantage of fixed-point dot products,

4.2. Hybrid BFP-FP DNN training

BFP should be used for the most demanding, dot product based, computations, with other operations being performed in floating-point-like representations. This configuration enables the bulk of the DNN operations to be performed in efficient fixed-point logic, and facilitates the use of various activation functions or techniques like batch normalization without the restrictions imposed by BFP.

This configuration also leads to better choices of exponents when values are converted to BFP. Whenever an operation causes a change in the value distribution within a tensor, it is beneficial to perform this operation in a representation that allows individual values to freely float so that, when the tensor is converted back to BFP, a more appropriate exponent is chosen. For instance, with the hybrid approach, tensors resulting from long chains of dot products followed by activation functions always have their exponent adjusted after activations when they are converted to dot products. In contrast, entirely in BFP would result in tensors with exponents that do not reflect their value distribution after long chains of operations, incurring loss of precision, and

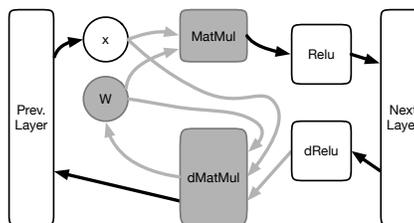


Figure 3. Traditional neural network layer dataflow using block floating point. The white boxes and black arrows indicate computations and values flowing in FP representation, and the grey boxes and arrows indicate operations and intermediate values in BFP.

requiring exponent adjusting techniques.

Figure 3 illustrates the dataflow of the forward and backward passes of a fully connected layer. Weights are stored in BFP format throughout the training process, to take advantage of the compressed nature of BFP representations. This reduces memory bandwidth during both forward and backward passes, as well as the amount of communication during parameter updates.

4.3. BFP design space

Floating-point tensors converted to BFP lose precision when tensors have a wide range of values. The BFP implementation can minimize the loss of precision during conversions by choosing an appropriate exponent for the tensor and rounding numbers with bias free policy.

We evaluated three exponent policies: *max*, *min* and *avg*. The *max* policy uses the exponent of the largest value in the tensor, rounding values that are too small. This is the policy used by traditional BFP implementations. The *min* policy guarantees that the smallest value in the tensor is represented by a minimum number of bits, and may lead to the clipping of large tensor elements. The *avg* policy guarantees that

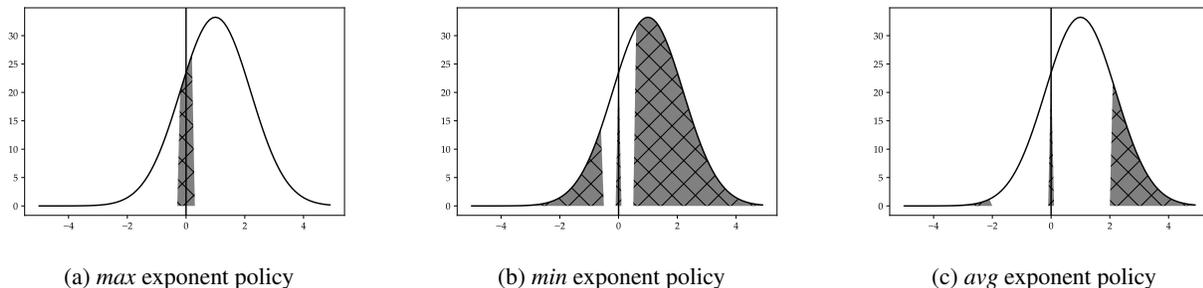


Figure 4. Matrix values rounded/saturated out by various exponent policies. The grey region of the curve indicates values that are lost when a tensor exponent is set. Values lost around 0 are rounded while other lost values are saturated.

the average value of the tensor can be represented by a minimum number of bits, and is a compromise between the two aforementioned policies. Figure 4 illustrates the ranges of values lost by each policy.

We evaluated two rounding policies: round-to-nearest and stochastic rounding (Gupta et al., 2015). Round-to-nearest (*determ*) deterministically rounds numbers to the nearest value, while stochastic rounding (*stoc*) stochastically rounds numbers with probability depending on the remainder of the number. We will show that rounding policies play a larger role when operating with narrow mantissas.

4.4. FPGA BFP prototype

To illustrate the area trade-offs of hybrid BFP-FP accelerators, we synthesized a proof-of-concept accelerator, shown in figure 5. We implemented the basic operations needed for neural network training (i.e., matrix multiplication, transpose, convolutions and data movement operations) using a dataflow similar to (Chen et al., 2016).

The matrix multiplication unit employs 75×75 systolic array of multiply-accumulate (MAC) units that feed a 75-wide activation/loss unit. The matrix multiplication unit operates on BFP values and the other units operate on custom floating point representation that features a 10-bit exponent and a 8-bit mantissa. In steady state, the matrix multiplication unit computes 75 dot products taking 75-wide tensors as inputs per cycle.

The FP-to-BFP units convert tensors by detecting the maximum exponent of the input FP tensors and normalizing the mantissas accordingly, while the BFP-to-FP unit normalizes the mantissas according to the single given exponent. The activation/loss and the conversion units are capable of processing a single 75-wide tensor per cycle. Weights are kept in BFP throughout the entire training process and during inference.

We synthesized the accelerator in a Stratix V 5SGSD5 FPGA at a clock rate of 200MHz. We achieve a maxi-

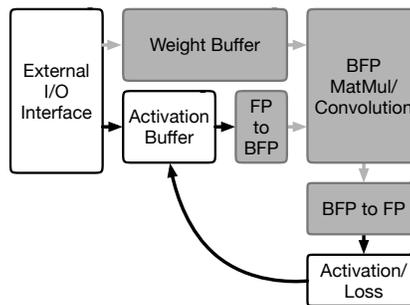


Figure 5. Hybrid BFP-FP accelerator with BFP

imum throughput of 1 TOP/s when using 8-bit wide MACs in the matrix-multiplier with FP activations, the FP-to-BFP and the BFP-to-FP conversion units occupying less than 10% of the FPGA resources. This is an $8.5 \times$ throughput improvement over a variant of the accelerator that employs FP16 MAC units synthesized on the same FPGA.

5. Methodology

5.1. Implementation

We train DNNs with the hybrid approach, using BFP in the compute-intensive operations (matrix multiplications, convolutions) and FP32 in the other operations. We modified TensorFlow’s (Abadi et al., 2016) matrix multiplications and convolution operations to reproduce the behaviour of BFP matrix multipliers in both the forward and backward passes.

We used TensorFlow’s *defun* function to create a new op that processes the inputs and outputs of both the forward and backward passes of another tensorflow op, to simulate the usage of BFP. In the forward pass, shown in Figure 6a, we convert both inputs (x and w) to BFP, giving the x tensor one exponent per training input and the w tensor one

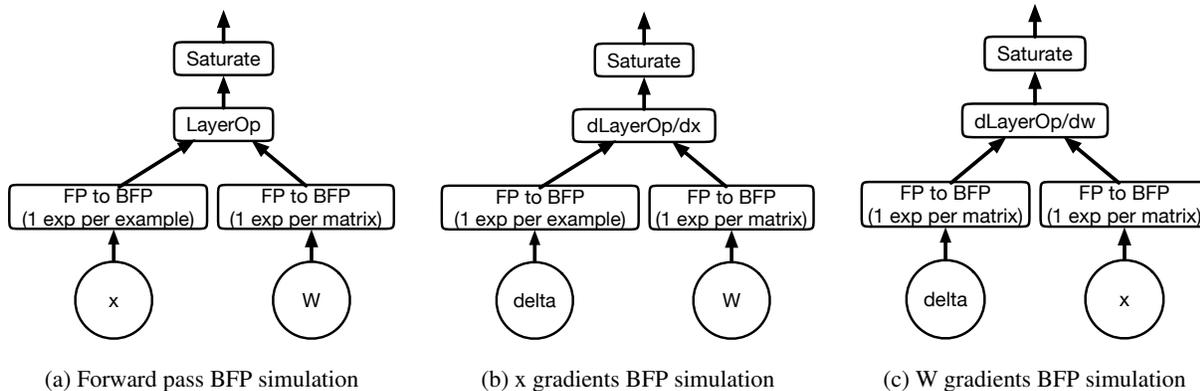


Figure 6. BFP simulation in TensorFlow. Both *layerOp* and their derivatives are native FP32 operations executed in a traditional GPU.

exponent per matrix. Then we execute the target operation with native floating-point arithmetic, and saturate the outputs of the original op, to simulate the saturation that occurs in fixed-point matrix multipliers. In the backward pass, we perform the same pre-/post-processing of the inputs/outputs of the x derivative (Figure 6b), but handle the w derivative differently (Figure 6c) since it performs a reduction across entire batches. Thus, to emulate the behavior of an accelerator with native BFP, we convert inputs to BFP tensors that share exponents across the entire batch. Finally, we re-align weights and their gradients during updates to simulate the update of weights stored in BFP.

Using *defun* enables us to evaluate the impact of the hybrid approach on training quality without building and integrating a full-blown distributed BFP accelerator into a machine learning framework. It also enables us to take advantage of the highly optimized GPU kernels already available for all the different varieties of convolution and fully-connected layers.

5.2. Evaluation Setup

Datasets. We experiment with a set of popular image classification tasks.

- CIFAR-10 and CIFAR-100 (Krizhevsky, 2009). Each consists of a training set of size 50K and a test set of size 10K. Instances are 32×32 color images representing 10 or 100 classes. We adopt a standard data augmentation scheme (He et al., 2016; Huang et al., 2016), by randomly cropping and flipping. For pre-processing, we normalize the data using the channel means and standard deviations.

Note that we use a model trained on CIFAR-10 to explore the design space of block-floating point implementations, and report the overall performance of BFP on the more challenging CIFAR-100.

- The SVHN (Netzer et al., 2011) dataset consists of color images of house numbers collected by Google Street View. The data format is the image of size 32×32 centered around a single character. It consists of 73K images in the training set and 26K images in the test set. We do not use data augmentation and only divide the pixel values by 255 for data pre-processing.

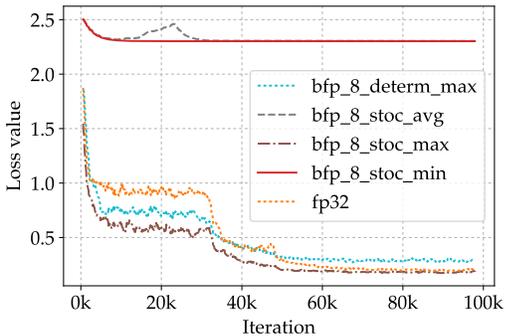
Evaluation Metric. To evaluate the impact of BFP, we tune the models using only FP32, and then train the same models from scratch with the same hyper-parameters in BFP. We report training loss and best top-1 error.

Training. We train CIFAR-10/CIFAR-100 with ResNet (He et al., 2016) and WideResNet (Zagoruyko & Komodakis, 2016), and SHVN (Netzer et al., 2011) with ResNet, using various configurations of BFP.

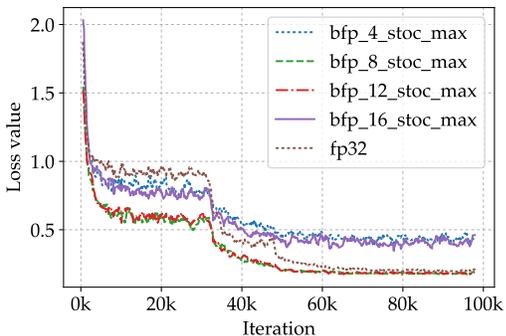
Our models are trained by momentum SGD with a mini-batch size of 128. We use a weight decay of $1e - 4$ and momentum of 0.9 for our datasets. We trained models on CIFAR-10 and CIFAR-100 for 250 epochs starting with a learning rate of 0.1, and dividing it by 10 at 32K, 48K and 64K iterations (He et al., 2016). We trained the SVHN models for 160 epochs, starting from an initial learning rate of 0.01, and dividing it by 10 at epochs 80 and 120.

6. Evaluation

We now evaluate DNN training with the hybrid approach, that is referred to as BFP for simplicity, comparing it to FP32-based training. We start with a BFP design space exploration, where we train a ResNet-20 model on CIFAR-10, to explore the different choices of exponent range and rounding policy, as well as various mantissa bit-widths. Then we compare BFP- with FP32-based training for more challenging models on CIFAR-100 and SVHN. Our evaluation intends to show that BFP can be used as a drop-in replace-



(a) Rounding and exponent policies.



(b) Bit-width

Figure 7. BFP design space exploration: training loss of ResNet-20 on CIFAR-10 with various BFP configurations.

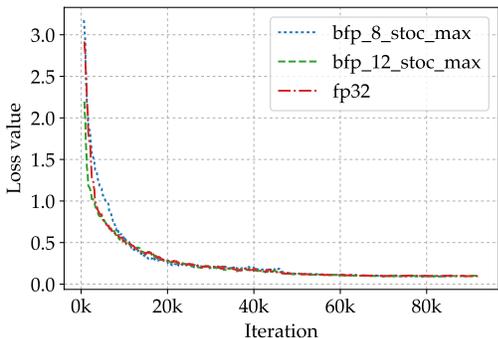


Figure 8. ResNet-50 training loss on SVHN

ment for FP32.

6.1. Exploring the BFP design space

Figure 7 shows the training loss of BFP with various rounding and exponent policies and mantissa bit-widths. We trained models on all the possible configurations, but we only show the best performing one when varying a parameter in the design space.

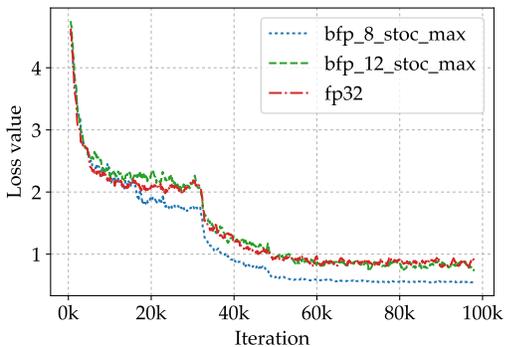
Exponent policy. The exponent policy that works the best is *max*, as shown in Figure 7a. The *min* policy saturates most values due to very small outliers in tensors, resulting in exponents that are too small, and preventing convergence. This problem is mitigated by the *avg* policy that chooses more reasonable exponents, but still incurs too many saturations. Models trained with *avg* also do not converge. The *max* policy works the best because it does not incur saturation at all, and the models do not seem to suffer from the loss of low valued activations and gradients, even with narrow mantissas.

Rounding policy. Models using *stoc* outperform their *determ* counterparts consistently, especially with narrow mantissas. For instance, when using 4-bit mantissas (not shown in the Figure 7a), the *determ* policy leads to divergence.

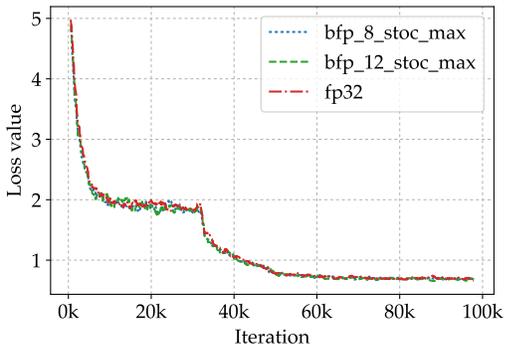
Mantissa bit-width. Figure 7b shows BFP performance with various mantissa widths. Both 12- and 8-bit-mantissa BFP outperform FP32 while 4- and 16-bit-mantissa performs worse than FP32. This result, also observed in other models and datasets, indicates a sweet spot in the mantissa bit-width design space. We believe that 12- and 8-bit mantissas regularize the weight matrices, compensating for the loss of precision incurred by BFP. This result also appears in FP32 representations, where using 16-bit mantissas outperforms the baseline with 24-bit wide mantissas, as shown in Table 1. Although 4-bit-mantissa BFP is outperformed by FP32, it still converges, uncovering a quality-performance trade-off: users that can tolerate models with lower quality can achieve better energy-efficiency during training and inference. Overall, BFP converges with narrow 8- and 12-bit wide mantissas resulting in a 4× smaller model than the FP32 baseline. The overhead of carrying an exponent per tensor is negligible, since tensors often carry 100’s of elements.

6.2. BFP vs. FP32

Figure 8 shows the training loss of ResNet models trained on SVHN, and Table 3 shows the test error for these models. Both BFP and FP32 behave similarly during training, resulting in similar test errors. Figures 9 show the training loss of ResNet and WideResNet models trained on CIFAR-100, and Table 3 shows the test error for these models. For all



(a) ResNet-50



(b) WideResNet-52

Figure 9. CIFAR-100 training loss with BFP.

models, 12- and 8-bit mantissa BFP either outperforms or matches FP32, with 8-bit mantissa representations being the best in all models and datasets. BFP is robust to various models and different datasets, and it is indeed a viable, more efficient alternative to FP32 for general purpose deep learning. Using 8-bit mantissa BFP results in models that are 4 \times smaller than baseline FP32 with most computations performed on fixed-point arithmetic.

7. Conclusion

DNNs have become ubiquitous in datacenter settings, forcing operators to adopt specialized hardware to execute and train them. However, DNN training depends on floating-point number representations for convergence, severely limiting the efficiency of accelerators. In this paper, we propose the use of a hybrid BFP-FP number representation for the dot product computations in DNN training. We show that the hybrid approach leads to efficient hardware, with the bulk of the silicon real-estate spent on efficient fixed-point logic. Finally, we evaluate the hybrid approach, and show that, for all models evaluated, BFP training either matches or outperforms their counterparts trained with FP32. BFP results in more compact models, with 4 \times and 2 \times smaller mod-

Table 2. BFP design space exploration test error

Exponent Policies			
Config.	<i>min</i>	<i>avg</i>	<i>max</i>
BFP8_STOC	–	–	8.25%

Range policies

Config.	<i>determ</i>	<i>stoc</i>
BFP8_MAX	9.50%	8.25%

Bit-widths

Config.	4	8	12	16
BFP8_STOC_MAX	13.39%	8.25%	8.41%	8.43%

Table 3. ResNet-50 and WideResNet-28-10 test error.

Config.	ResNet-50		Wide-ResNet-28-10
	CIFAR-100	SVHN	CIFAR-100
FP32	33.02%	4.56%	27.69%
BFP8	29.60%	4.34%	27.69%
BFP12	32.40%	4.50%	27.93%

els when compared to FP32 and FP16, respectively. BFP also leads to faster accelerators, with 8-bit BFP achieving 8.5 \times higher throughput when compared to FP16. Higher throughput leads to faster and more energy-efficient DNN training/inference, while model compression leads to lower bandwidth requirements for off-chip memory, lower capacity requirements for on-chip memory and lower communication bandwidth requirements for distributed training.

References

- Cloud tpu, 2017. URL <https://cloud.google.com/tpu>. Accessed: 2018-01-31.
- How to quantize neural networks with tensorflow, 2017. URL <https://www.tensorflow.org/performance/quantization>. Accessed: 2018-01-31.
- Artificial intelligence architecture, 2018. URL <https://www.nvidia.com/en-us/data-center/volta-gpu-architecture>. Accessed: 2018-01-31.
- Abadi, Martín, Barham, Paul, Chen, Jianmin, Chen, Zhifeng, Davis, Andy, Dean, Jeffrey, Devin, Matthieu, Ghemawat, Sanjay, Irving, Geoffrey, Isard, Michael, Kudlur, Manjunath, Levenberg, Josh, Monga, Rajat, Moore, Sherry, Murray, Derek Gordon, Steiner, Benoit, Tucker, Paul A., Vasudevan, Vijay, Warden, Pete, Wicke, Martin, Yu, Yuan,

- and Zheng, Xiaoqiang. TensorFlow: A System for Large-Scale Machine Learning. pp. 265–283, 2016.
- Alistarh, Dan, Grubic, Demjan, Li, Jerry, Tomioka, Ryota, and Vojnovic, Milan. QSGD: Communication-Efficient SGD via Gradient Quantization and Encoding. pp. 1707–1718, 2017.
- Chen, Yu-Hsin, Emer, Joel S., and Sze, Vivienne. Eyeriss: A Spatial Architecture for Energy-Efficient Dataflow for Convolutional Neural Networks. pp. 367–379, 2016.
- Courbariaux, Matthieu, Bengio, Yoshua, and David, Jean-Pierre. BinaryConnect: Training Deep Neural Networks with binary weights during propagations. pp. 3123–3131, 2015.
- Courbariaux, Matthieu, Hubara, Itay, Soudry, Daniel, El-Yaniv, Ran, and Bengio, Yoshua. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.
- Dally, William. High performance hardware for machine learning, 2015. URL <https://media.nips.cc/Conferences/2015/tutorialslides/Dally-NIPS-Tutorial-2015.pdf>. Accessed: 2018-01-31.
- Gupta, Suyog, Agrawal, Ankur, Gopalakrishnan, Kailash, and Narayanan, Pritish. Deep Learning with Limited Numerical Precision. pp. 1737–1746, 2015.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep Residual Learning for Image Recognition. pp. 770–778, 2016.
- Huang, Gao, Sun, Yu, Liu, Zhuang, Sedra, Daniel, and Weinberger, Kilian Q. Deep networks with stochastic depth. In *European Conference on Computer Vision*, pp. 646–661. Springer, 2016.
- Jouppi, Norman P., Young, Cliff, Patil, Nishant, Patterson, David, Agrawal, Gaurav, Bajwa, Raminder, Bates, Sarah, Bhatia, Suresh, Boden, Nan, Borchers, Al, Boyle, Rick, luc Cantin, Pierre, Chao, Clifford, Clark, Chris, Coriell, Jeremy, Daley, Mike, Dau, Matt, Dean, Jeffrey, Gelb, Ben, Ghaemmaghami, Tara Vazir, Gottipati, Rajendra, Gulland, William, Hagmann, Robert, Ho, C. Richard, Hogberg, Doug, Hu, John, Hundt, Robert, Hurt, Dan, Ibarz, Julian, Jaffey, Aaron, Jaworski, Alek, Kaplan, Alexander, Khaitan, Harshit, Killebrew, Daniel, Koch, Andy, Kumar, Naveen, Lacy, Steve, Laudon, James, Law, James, Le, Diemthu, Leary, Chris, Liu, Zhuyuan, Lucke, Kyle, Lundin, Alan, MacKean, Gordon, Maggiore, Adriana, Mahony, Maire, Miller, Kieran, Nagarajan, Rahul, Narayanaswami, Ravi, Ni, Ray, Nix, Kathy, Norrie, Thomas, Omernick, Mark, Penukonda, Narayana, Phelps, Andy, Ross, Jonathan, Ross, Matt, Salek, Amir, Samadiani, Emad, Severn, Chris, Sizikov, Gregory, Snelham, Matthew, Souter, Jed, Steinberg, Dan, Swing, Andy, Tan, Mercedes, Thorson, Gregory, Tian, Bo, Toma, Horia, Tuttle, Erick, Vasudevan, Vijay, Walter, Richard, Wang, Walter, Wilcox, Eric, and Yoon, Doe Hyun. In-Datacenter Performance Analysis of a Tensor Processing Unit. pp. 1–12, 2017. doi: 10.1145/3079856.3080246.
- Köster, Urs, Webb, Tristan, Wang, Xin, Nassar, Marcel, Bansal, Arjun K., Constable, William, Elibol, Oguz, Hall, Stewart, Hornof, Luke, Khosrowshahi, Amir, Kloss, Carey, Pai, Ruby J., and Rao, Naveen. Flexpoint: An Adaptive Numerical Format for Efficient Training of Deep Neural Networks. pp. 1740–1750, 2017.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009.
- Li, Fengfu and Liu, Bin. Ternary Weight Networks. *CoRR*, abs/1605.04711, 2016.
- Netzer, Yuval, Wang, Tao, Coates, Adam, Bissacco, Alessandro, Wu, Bo, and Ng, Andrew Y. Reading digits in natural images with unsupervised feature learning. *Deep Learning and Unsupervised Feature Learning Workshop*, 2011.
- Rastegari, Mohammad, Ordonez, Vicente, Redmon, Joseph, and Farhadi, Ali. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks. In *ECCV*, pp. 525–542. Springer International Publishing, September 2016.
- Song, Zhouhui, Liu, Zhenyu, Wang, Chunlu, and Wang, Dongsheng. Computation Error Analysis of Block Floating Point Arithmetic Oriented Convolution Neural Network Accelerator Design. *CoRR*, abs/1709.07776, 2017.
- Zagoruyko, Sergey and Komodakis, Nikos. Wide Residual Networks. 2016.
- Zhang, Hantian, Li, Jerry, Kara, Kaan, Alistarh, Dan, Liu, Ji, and Zhang, Ce. ZipML: Training Linear Models with End-to-End Low Precision, and a Little Bit of Deep Learning. pp. 4035–4043, 2017.
- Zhou, Shuchang, Ni, Zekun, Zhou, Xinyu, Wen, He, Wu, Yuxin, and Zou, Yuheng. DoReFa-Net: Training Low Bitwidth Convolutional Neural Networks with Low Bitwidth Gradients. *CoRR*, abs/1606.06160, 2016.
- Zhu, Chenzhuo, Han, Song, Mao, Huizi, and Dally, William J. Trained Ternary Quantization. *CoRR*, abs/1612.01064, 2016.