Interleaver Design for Deep Neural Networks

Sourya Dey, Peter A. Beerel and Keith M. Chugg

Ming Hsieh Department of Electrical Engineering

University of Southern California, Los Angeles, California 90089

{souryade, pabeerel, chugg}@usc.edu

Abstract—We propose a class of interleavers for a novel deep neural network (DNN) architecture that uses algorithmically predetermined, structured sparsity to significantly lower memory and computational requirements, and speed up training. The interleavers guarantee clash-free memory accesses to eliminate idle operational cycles, optimize spread and dispersion to improve network performance, and are designed to ease the complexity of memory address computations in hardware. We present a design algorithm with mathematical proofs for these properties. We also explore interleaver variations and analyze the behavior of neural networks as a function of interleaver metrics.

I. INTRODUCTION

DNNs in machine learning systems are critical drivers of new technologies such as speech processing and autonomous vehicles. Modern DNNs typically have millions of parameters [1], which make them difficult to implement in hardware and slow to train [2]. A suggested solution to these problems is a sparse network, where some form of compression or deletion is employed to reduce the number of parameters [3], [4]. However, an issue with sparse networks is that some neurons may get completely disconnected from neighboring layers and have no effect on the output [5]. A second issue arises when all the neurons in a certain layer which connect to a certain neuron in the next layer are 'close together', such as coming from nearby pixels in an image. This issue is similar to convolutional layers, which are known to be inadequate for classification without the presence of fully connected (FC) classification layers [1], [2]. The term 'layer' will henceforth to classification layer, which is what this work deals with.

We are investigating a class of hardware-optimized DNN architectures which use *pre-defined sparsity*, wherein a connection pattern is algorithmically defined using an *interleaver*, or permutation, for every *junction* between 2 layers prior to training. This has the potential to achieve higher training speed and lower storage complexity compared to approaches which start training the full network and then remove parameters [6], [7]. A related paper [8] has demonstrated that our approach can reduce the memory footprint of FC layers in CNNs by 457x without performance degradation.

This paper is a followup to our previous work [9] and focuses on the design and analysis of interleavers suited to the

requirements of our hardware architecture, which is reviewed in Section II. The key contributions of this paper are:

1) Mathematical formalizations of desirable properties of a class of interleavers usable in DNNs (Section III). The interleavers should implement pseudo-random connection patterns between layers so as to achieve:
   a) Flexible degrees of sparsity in the junctions, while preventing neurons from getting disconnected.
   b) Maximum operational efficiency by avoiding pipeline stalls.
   c) Ease of address computation for on-chip memories.
2) An algorithm to design such interleavers (Section IV-A) and mathematical proofs to show that it satisfies the requirements (Section IV-B).
3) Possible variations in interleaver design (Section IV-C).
4) Relations between network performance and interleaver metrics such as spread and dispersion (Section IV-D), explored through training on different datasets.

II. HARDWARE ARCHITECTURE

A DNN is made up of layers of neurons, and junctions connecting adjacent layers via *weights*, or *edges*. We will use p and n to represent the number of neurons in the preceding (left) and succeeding (right) layers, respectively, of any junction. Every left neuron has a fixed number of edges going from it to the right, and every right neuron has a fixed number of edges coming into it from the left. These numbers are defined as fan-out (*fo*) and fan-in (*fi*), respectively. For a conventional FC junction, $fo = n$ and $fi = p$. Every neuron has associated activation and delta values which are used in the 3 operations – feedforward (FF), backpropagation (BP), and update (UP).

A. Our Architecture

For our sparse architecture, $fo < n$ and $fi < p$, such that $p \times fo = n \times fi = W$, the total number of edges in the junction. They are sequentially indexed on the right side, for example, the 1st right neuron has weights w_0 to w_{fi-1}. Motivated by the fact that the weights feature in all 3 network operations, we designed an *edge-processing* architecture where every junction has a *degree of parallelism (DoP)*, denoted as z, which is the number of weights processed in parallel. (z is chosen such that p/z is an integer). All the weights in each junction are stored in a bank of z memories, each having W/z *cells*, as shown in Fig. 1. This means that 2 weights with indices i and j (i.e. weights w_i and w_j) are in the same memory if $i\%z = j\%z$, where $\%$ is the modulo operator. If
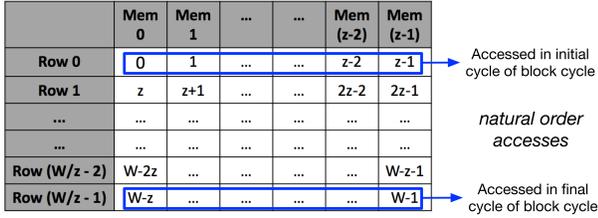
Fig. 1. Weight memory configuration in any junction, showing natural order accesses. The number in each cell is the index of the weight, i.e. the i in w_i.

Fig. 2. Reading $z = 6$ weights corresponding to 2 right neurons in each cycle. When traced back through π_W, this requires reading 6 left activation memories in permuted order.

instead $\lfloor i/z \rfloor = \lfloor j/z \rfloor$, where $\lfloor . \rfloor$ is the floor function, then the weights are in the same row of different memories.

Similar to the weights, all the activation and delta values of each layer are numbered and stored in separate banks of z memories each. For example, the left layer activations are numbered from a_0 for the first neuron to a_{p-1} for the last, and each activation memory would have p/z elements. The edges coming into a junction from the left pass through a weight interleaver (π_W) before getting connected to the right. For example, say 4 edges come out of the 1st neuron of a certain layer of a network which has a 100-neuron layer following it. These edges might connect to the 9th, 30th, 67th and 84th neurons of the following layer.

A single *cycle* of processing (say the kth) comprises accessing the kth cell in each of the z weight memories. This implies reading all z values from the kth row of the bank, which we refer to as *natural order* access, as shown in Fig. 1. The interleaver determines which neurons in the left layer are connected to those z edges. In general, these could be any z neurons in the left layer. So the activation memories are accessed in *permuted order*. Fig. 2 shows this through an example where z is 6 and fi is 3. Note that all the entries in the left activation memory bank are read fo times, since that many weights belong to the same neuron and share the same activation value. Each stage of processing where all the activations are read once is referred to as a *sweep*, which consists of p/z cycles. One complete operation such as FF consists of fo sweeps, i.e. $p \times fo/z = W/z$ cycles, which are collectively referred to as 1 *block cycle*.

B. Merits of our Architecture

Since there is significant data reuse between FF, BP and UP, we use *operational parallelization* to make all of them occur concurrently. Since every operation in a junction uses data generated by an adjacent junction or layer, we designed a *junction pipelining* architecture where all the junctions execute all 3 operations concurrently on different inputs from the training set. This enables our architecture to achieve a $3(L-1)$ times speedup for L layers. See [9] for a complete description.

Note that z can be set to any value as per the overall area-speed tradeoff desired. The number of clock cycles to process each junction can be made equal by adjusting z for each individually. This ensures an always full pipeline and no stalls. Thus, the size and complexity of the network is decoupled from the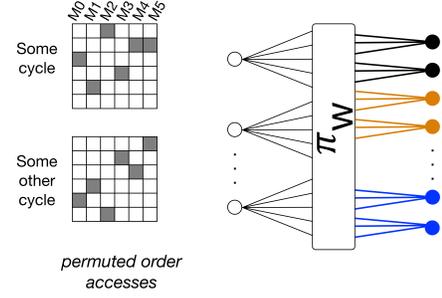 hardware resources available. Our architecture can be reconfigured to varying amounts of fan-out and sparsity, which makes it adaptable to a large class of DNNs. This speedup and flexibility gives us the potential to achieve *online* training, as compared to inference-only works such as [4].

III. INTERLEAVER REQUIREMENTS

An interleaver π operates on elements i from a list x with cardinality N and produces rearranged list elements $\pi(i)$. We will follow the convention that $x = \{0, 1, ..., N-1\}$. As an example, let $x = \{0, 1, 2, 3\}$, i.e. $N = 4$. Then $\pi(x) = \{\pi(0), \pi(1), \pi(2), \pi(3)\}$, such as $\{1, 3, 2, 0\}$. Interleaver patterns can be visualized by plotting $\pi(x)$ vs. x.

A. Clash Freedom

As mentioned before, the activation memories are accessed in permuted order. For any weight index i, the corresponding left activation index is $\lfloor \pi_W(i)/fo \rfloor$. The z activations read in a cycle should come from z different left neurons in order to achieve optimum spatial spread. Moreover, these z values should be stored in z different activation memories. Violating this condition leads to the same memory needing to be accessed more than once in the same cycle, i.e. a *clash*, which stalls processing. Notice that Fig. 2 is free from clashes since all the columns in permuted order accesses have exactly 1 shaded cell. Clash-freedom is mathematically expressed as:

$$\text{If} \quad \lfloor i/z \rfloor = \lfloor j/z \rfloor \tag{1a}$$

$$\text{Then we need} \quad \lfloor \pi_W(i)/fo \rfloor \% z \neq \lfloor \pi_W(j)/fo \rfloor \% z \tag{1b}$$

where $i \neq j$ is implicitly assumed here, and in the future. Equation (1) implies that for 2 weights w_i and w_j read in the same cycle, their left activations must be in different memories.

B. Ease of Memory Address Computation

The interleaver should be designed so that the addresses of the activation memories (accessed in permuted order) can be easily computed in any cycle. This can be done by defining a starting cell index – to be used in the first cycle of every sweep – for each activation memory. Cell indices for the following cycles are obtained by adding 1 each time to the starting index, and cycling back to the first cell after reaching the last.

	Mem 0	Mem 1	Mem 2	Mem 3	Mem 4	Mem 5	Mem 6	Mem 7
Row 0	a0	a1	a2	a3	a4	a5	a6	a7
Row 1	a8	a9	a10	a11	a12	a13	a14	a15
Row 2	a16	a17	a18	a19	a20	a21	a22	a23
Row 3	a24	a25	a26	a27	a28	a29	a30	a31

Fig. 3. Activation memory bank configuration for a left layer with $p = 32$ neurons when the junction following it has $z = 8$.

As a concrete example, assume $p = 32$, $fo = 2$, and $z = 8$. This leads to the activation memory mapping shown in Fig. 3. Let us define the starting cell indices for the 8 activation memories as $s = \{2, 0, 3, 1, 2, 0, 3, 1\}$. Then the cells read in the next cycle will be $(s + 1)\%4 = \{3, 1, 0, 2, 3, 1, 0, 2\}$, and so on until all 4 cycles in the sweep are completed. This can be mathematically expressed as:

$$\text{If} \qquad \lfloor i/z \rfloor \neq \lfloor j/z \rfloor \tag{2a}$$

$$\text{and} \qquad \lfloor \pi_W(i)/fo \rfloor \%z = \lfloor \pi_W(j)/fo \rfloor \%z \tag{2b}$$

Then we need

$$(\lfloor i/z \rfloor - \lfloor j/z \rfloor)\%(p/z) = \left(\left\lfloor \frac{\lfloor \pi_W(i)/fo \rfloor}{z} \right\rfloor - \left\lfloor \frac{\lfloor \pi_W(j)/fo \rfloor}{z} \right\rfloor \right) \%(p/z) \tag{2c}$$

Equations (2a) and (2b) consider 2 weights with indices i and j such that they are in different cycles and the left neurons to which they connect are in the same activation memory. Then, (2c) states that the difference in cycle numbers should be equal to the difference in activation memory row numbers. This leads to ease of address computation.

C. Optional Requirements – Spread and Dispersion

Spread is a standard interleaver metric which, when maximized, ensures that for 2 weights that are close together on the right (such as going to the same neuron), the neurons from which they come on the left are spaced well apart. Spread is classically defined [10] as:

$$\text{Spread} = \min\left(|i - j|\%N + |\pi(i) - \pi(j)|\%N\right) \tag{3}$$

Normalized dispersion, which we will simply refer to as dispersion is another standard metric measuring the randomness in the connection pattern. For example, if the 1st left neuron connects to the 10th, 20th and 30th right neurons, and the 2nd left neuron connects to the 11th, 21st and 31st right neurons, the pattern is quite regular and not well dispersed. Dispersion is classically defined [11] as the cardinality of the set

$$\mathbb{D} = \{(j - i, \pi(j) - \pi(i)) \mid 0 \leq i < j < N\} \tag{4}$$

divided by $N(N - 1)$. The effects of spread and dispersion on network performance are discussed in Section IV-D.

IV. INTERLEAVER DESIGN

A. Algorithm

Given the requirements of the DNN, we developed the following algorithm to design a suitable class of interleavers:

1) Let r be a random permutation of $[0, p/z - 1]$
2) Create list s with z elements according to:
   a) If $z \geq p/z$: Replicate r as many times as necessary to get z elements in s. If z is not an integral multiple of p/z, then fill the final few elements of s with the initial few elements of r. For example, if $r = \{2, 0, 3, 1\}$ and $z = 10$, then $s = \{2, 0, 3, 1, 2, 0, 3, 1, 2, 0\}$.
   b) If $z < p/z$: Take the 1st z elements of r
3) Create list t with p elements by concatenating s, $(s + 1)\%(p/z)$, ..., $(s + \frac{p}{z} - 1)\%(p/z)$. t acts as an activation interleaver (π_A), from which π_W can be obtained.
4) Let $t[x]$ denote the xth element of t. Then:

$$\pi_W(i) = (t[i\%p] \times z + i\%z) \times fo + \lfloor i/p \rfloor \tag{5}$$
$$\forall i \in [0, W - 1]$$

Consider the prior example from Section III-B. Say $r = \{2, 0, 3, 1\}$. Since $z \geq p/z$, $s = \{2, 0, 3, 1, 2, 0, 3, 1\}$. Since $p/z = 4$, $t = \{2,0,3,1,2,0,3,1,3,1,0,2,3,1,0,2,0,2,1,3,0,2,1,3,1,3,2,0,1,3,2,0\}$. There are 64 weights. Say we are in cycle 5, where one of the weights read is w_{45}. Using (5), $t[45\%32] = t[13] = 1$. This gives the row number in the left activation memory bank. The term $i\%z$ is the bank column, which is $45\%8 = 5$. Now the key purpose of the interleaver equation, which is to compute the addresses of the activation memory bank used in a cycle, is served. Since our architecture uses powers of 2 for all the key variables, operations such as multiplication, modulo and flooring reduce to simple bit shifts and bit selects.

The remainder of (5) serves the purely mathematical purpose of completely characterizing π_W as a permutation of 64 weights. Multiplying the bank row by $z = 8$ and adding the bank column gives the left neuron number from where the weight comes into the junction, i.e. $1 \times 8 + 5 = 13$. Multiplying this by $fo = 2$ takes us from the activation space to the weight space, while the final addition of $\lfloor 45/32 \rfloor = 1$ adds an offset to indicate that it's the 2nd weight from neuron 13. The final index of the weight on the left side is 27. Thus, $\pi_W(45) = 27$.

B. Meeting Requirements

Now we will prove that given the interleaver design equation (5), the requirements in (1) and (2) are satisfied.

1) Clash Freedom: *Proof:*
Since $W = p \times fo$, the $\lfloor i/p \rfloor$ term in (5) is in the range $[0, fo - 1]$. Then we get:

$$\lfloor \pi_W(i)/fo \rfloor = \left\lfloor \frac{(t[i\%p] \times z + i\%z) \times fo + \lfloor i/p \rfloor}{fo} \right\rfloor$$
$$= t[i\%p] \times z + i\%z \tag{6}$$

So $\quad \lfloor \pi_W(i)/fo \rfloor \% z = i \% z$ $\qquad(7)$

It is given from (1a) that $\lfloor i/z \rfloor = \lfloor j/z \rfloor$, but $i \neq j$ as usual. So it must be that $i\%z \neq j\%z$. This implies that:

$$\lfloor \pi_W(i)/fo \rfloor \% z \neq \lfloor \pi_W(j)/fo \rfloor \% z \qquad(8)$$

which satisfies (1b). $\blacksquare$

2) Ease of Memory Address Computation: *Proof:* Firstly, note that using (6) and (7), (2b) can be written as $i\%z = j\%z$. Secondly, using (6):

$$\left\lfloor \frac{\lfloor \pi_W(i)/fo \rfloor}{z} \right\rfloor = \left\lfloor \frac{t[i\%p] \times z + i\%z}{z} \right\rfloor$$
$$= t[i\%p] \qquad(9)$$

So the right hand side of (2c) can be written as $(t[i\%p] - t[j\%p])\%(p/z)$. t is constructed by concatenating s repeatedly with some changing offset added to it every time. Using this, and the fact that s has z elements, we get:

$$t[i\%p] = (s[i\%z] + \lfloor (i\%p)/z \rfloor)\%(p/z) \qquad(10)$$

So the modified right hand side of (2c) now becomes:

$$(t[i\%p] - t[j\%p])\%(p/z) = \{(s[i\%z] + \lfloor (i\%p)/z \rfloor)\%(p/z)$$
$$- (s[j\%z] + \lfloor (j\%p)/z \rfloor)\%(p/z)\}\%(p/z) \quad(11)$$

We will use 2 mathematical theorems in this proof. Given any 3 positive integers a, b and c, firstly:

$$(a \pm b)\%c = (a\%c \pm b\%c)\%c \qquad(12)$$

Secondly, if b is an integral multiple of c:

$$\lfloor (a\%b)/c \rfloor = \lfloor a/c \rfloor \% (b/c) \qquad(13)$$

Using (12) and the fact that $i\%z = j\%z$, (11) becomes:

$$(11) = (s[i\%z] + \lfloor (i\%p)/z \rfloor - s[j\%z] - \lfloor (j\%p)/z \rfloor)\%(p/z)$$
$$= (\lfloor (i\%p)/z \rfloor - \lfloor (j\%p)/z \rfloor)\%(p/z) \qquad(14)$$

Using (12), (13) and the fact that p is an integral multiple of z, the left hand side of (2c) becomes:

$$(\lfloor i/z \rfloor - \lfloor j/z \rfloor)\%(p/z)$$
$$= \{(\lfloor i/z \rfloor \%(p/z)) - (\lfloor j/z \rfloor \%(p/z))\}\%(p/z)$$
$$= (\lfloor (i\%p)/z \rfloor - \lfloor (j\%p)/z \rfloor)\%(p/z) \qquad(15)$$

which equals the right hand side of (2c), as obtained in (14). Thus, the requirement in (2c) is satisfied. $\blacksquare$

C. Variations

The *basic* π_W described so far has excellent spread, but poor dispersion. We experimented with the following variations, all of which still satisfy the properties of clash-freedom and ease of memory address computation:

1) Start Vector Shuffle (SV): This only applies when $z > p/z$. Instead of simply replicating r in step 2a of the interleaver design algorithm, random permutations of $[0, p/z - 1]$ could be concatenated together to form s. In other words, there are several different r vectors. For example, $s = \{2, 0, 3, 1, 3, 0, 1, 2, 1, 0\}$.

2) Sweep Starter Shuffle (SS): This only applies when $fo > 1$. No algorithm change is required. Every time a new sweep is started, a new r, s and t are generated. For example, for the 1st sweep, $s = \{2, 0, 3, 1, 2, 0, 3, 1, 2, 0\}$, for the 2nd sweep, $s = \{0, 3, 2, 1, 0, 3, 2, 1, 0, 3\}$, and so on. This leads to a revised interleaver equation:

$$\pi_W(i) = (t_k[i\%p] \times z + i\%z) \times fo + \lfloor i/p \rfloor \qquad(16)$$
$$\forall i \in [0, W-1], \forall k \in [0, fo-1]$$

i.e. every sweep has a new t.

3) Memory Dither (MD): Equation (5) reveals that for any cycle, the weight read from the ith weight memory ($i \in [0, z-1]$) will always trace back to a left activation value stored in the ith activation memory. This trait can be removed and dispersion increased by replacing the 'activation memory number generating' term $i\%z$ in (5) with $v_k[i\%z]$, where v_k is a random permutation of $[0, z-1]$ for the kth cycle. In other words, the weight read from the ith weight memory will, in general, trace back to a different activation memory every cycle. However, the total z weights read in the kth cycle will always trace back to z different activation memories since v_k is a permutation, i.e. it has no repeated elements. This means that clash freedom is preserved as no activation memory needs to be accessed more than once in the same cycle. The revised interleaver equation is:

$$\pi_W(i) = (t[i\%p] \times z + v_k[i\%z]) \times fo + \lfloor i/p \rfloor \qquad(17)$$
$$\forall i \in [0, W-1], \forall k \in [0, W/z-1]$$

4) Meeting Requirements: Note that the proofs in Section IV-B were for a general t vector, which only has to satisfy the property that it is formed by concatenating $s + o$, where the offset o starts from 0 and goes up to $p/z - 1$.

- For SV, s is not constructed by repeating r, but this doesn't change the generality of t. In particular, t is still a p-element vector which specifies the complete order of accessing the activation memory bank during a sweep. So the proofs hold.
- For SS, t is different every sweep, but it's still constructed in the same way every sweep – by adding offsets of some vector s. So we can do our analysis within a sweep by keeping t fixed, which meets the requirements. Since 1 sweep is 1 complete access of the activation memory bank, this means that every individual sweep meets the requirements, so the interleaver meets the requirements as a whole.
- For MD, note that $v_k(\cdot)$ for every cycle is a bijective function mapping domain $= [0, z-1]$ to range $= [0, z-1]$, which means that provided weights i and j are both read in the kth cycle:

$$(i\%z = j\%z) \Leftrightarrow (v_k[i\%z] = v_k[j\%z]) \qquad(18)$$

For MD, eq. (7) becomes:

$$\lfloor \pi_W(i)/fo \rfloor \% z = v_k[i\%z] \qquad(19)$$

Interleaver Variant	π_W Spread	π_A Spread	π_W Disp.	π_A Disp.
Basic	**18.28**	**8**	0.04	0.1
MD	7.48	4.1	0.22	0.5
SS	9.7	**8**	0.07	0.1
SS+MD	6.5	4	0.37	0.5
SV	6.6	2.64	0.08	0.19
SV+MD	7.31	3.74	0.23	**0.52**
SV+SS	5.05	2.54	0.09	0.19
SV+SS+MD	5.7	3.47	**0.39**	**0.52**

Fig. 5. Corresponding π_A patterns for Fig. 4. Interleaver size $= p = 64$.

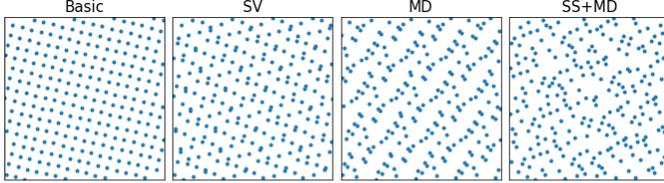

Fig. 4. Various π_W patterns using parameters $p = 64$, $fo = 4$ and $z = 16$. Interleaver size $= p \times fo = 256$.

Fig. 6. Classification accuracy obtained using different interleavers by training sparse networks on MNIST, CIFAR10 and Morse datasets for 10 epochs.

for the kth cycle. When proving clash-freedom, we consider weights i and j read in the same cycle. Given $i \neq j$, eq. (18) leads to $v[i\%z] \neq v[j\%z]$. Thus, eq. (8) holds and clash-freedom is satisfied.

However, ease of memory address computation *does not hold*. This is because the v permutation changes across cycles. As an example, assume $v_0 = \{0, 1, 2, 3, 4, 5, 6, 7\}$, $v_1 = \{2, 7, 3, 0, 6, 5, 1, 4\}$, and $s = \{2, 0, 3, 1, 2, 0, 3, 1\}$. So in cycle 0, the weight read from the 0th weight memory will lead to activation memory 0 row 2, however, the weight read from the 0th weight memory in cycle 1 will *not* lead to activation memory 0 row 3, instead it will lead to activation memory 2 row 0. So, while memory dither leads to clash freedom and increases the randomness and the number of possible clash-free patterns, it does not lead to ease of memory address computation.

D. Analysis and Results

Table I lists average spread and dispersion (disp.) over 100 iterations of all possible variations of π_W and corresponding π_A. Some of the patterns are shown in Figs. 4 and 5. Note that the basic π_W is the most linear, which leads to maximum spread and minimum dispersion. SS offers lesser spread and more dispersion for π_W, but no effect is observed on π_A. This is because SS affects different sweeps which have different weights, but same activations. SV offers slight increase in dispersion, but severe reduction in spread. This is because the SV pattern has lines with slope identical to basic, but each line is permuted, leading to left neurons getting bunched up. Introducing MD leads to big increases in dispersion, which are further increased for π_W when combined with SS. This is observed in the figures, where the MD patterns are irregular.
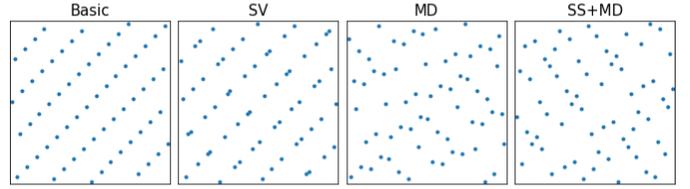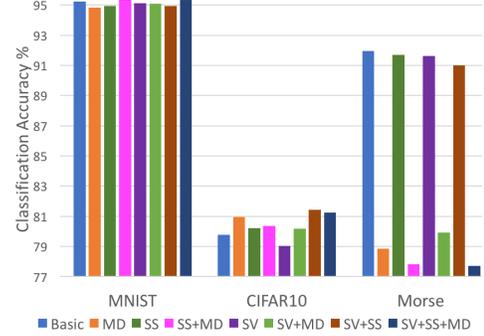
Fig. 6 shows results of all the possible interleaver variations implemented on networks trained for 10 epochs, with classification accuracy on validation data used as the performance metric. We used 3 datasets of different dimensionalities:

- MNIST handwritten digit classification – A 2D dataset where each input has width and height. The network has 1024 input, 64 hidden and 16 output neurons. Both junctions have $fo = 8$. This means that there are 8704 total weights, which is 13% of FC. z for the input-hidden junction is 512, and 32 for the hidden-output junction.
- CIFAR10 image classification – A 3D dataset where each input also has a number of features. We used standard convolutional and pooling layers for feature extraction, and then 2 sparse junctions in between layers of size 4096, 512 and 16. Fan-outs are 8 and 4, and z values are 2048 and 128. Overall density of the sparse junctions is 1.654%, and that of the entire network is 36.3%.
- Morse code symbol classification [12] – A 1D dataset where each input has 64 values representing dots, dashes and spaces in Morse code, and there are 64 output classes representing different characters. We created this dataset to rigorously test the limits of sparsity. The network used has 64 input and output neurons, and 1024 hidden neurons. Fan-outs are 384 and 24, and z values are both 64, leading to an overall density of 37.5%.

We observed that interleaver variations have negligible effect on classification accuracy of MNIST and CIFAR10 datasets. Note that for these datasets, the distinction between output classes is well pronounced. In MNIST for example, an image of a handwritten 7 is very different from a handwritten 0. Moreover, since the inputs in CIFAR10 are pre-processed by convolutional and pooling layers, the relative importance
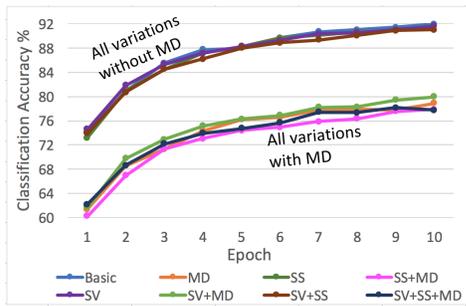
Fig. 7. Classification accuracy vs. epochs obtained using different interleavers by training a 37.5% dense network on the Morse dataset.

of the final classification layers is reduced.

For the Morse dataset, however, a clear trend of high dispersion hurting performance is observed. The 4 variations with MD have dispersion ≥ 0.5 and barely reach 80% accuracy, while the ones without MD have dispersions ≤ 0.2 and achieve $\geq 90\%$ accuracy. This dichotomy is further highlighted in Fig. 7. We hypothesize that this is due to the Morse dataset having lower redundancy compared to the other 2 since it has less input neurons and more output classes with little distinction between them. We are currently working on theories to better explain the link between dataset redundancy and high dispersion of junction connection patterns degrading performance.

V. CONCLUSION

This work presents a new way to design DNNs in hardware by interleaving edges between neurons and processing a programmable number of edges in parallel. The interleaver needs to be designed so as to achieve optimum network runtime efficiency on hardware. At the same time, performance needs to be maximized by selecting an interleaver with desirable metrics. We present an algorithm to satisfy interleaver requirements and investigate possible variations to it and their effects.

One limitation of these interleavers is that they characterize a single junction. To completely characterize a sparse network, it is desirable to have formulations which describe connection patterns in the whole network, such as which outputs connect to which inputs. We are currently working on the theory of adjacency matrices, which have elements corresponding to connections between any 2 neurons in any 2 layers, and exploring metrics which act as better proxies for performance.

REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1097–1105.
[2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015.
[3] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Jour. Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.
[4] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," in *Proc. ICLR*, 2016.
[5] W. Chen, J. T. Wilson, S. Tyree, K. Q. Weinberger, and Y. Chen, "Compressing neural networks with the hashing trick," in *Proc. ICML*, 2015, pp. 2285–2294.
[6] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Proc. NIPS*, 2015, pp. 1135–1143.
[7] X. Zhou, S. Li, K. Qin, K. Li, F. Tang, S. Hu, S. Liu, and Z. Lin, "Deep adaptive network: An efficient deep neural network with sparse binary connections," in *arXiv:1604.06154*, 2016.
[8] S. Dey, K.-W. Huang, P. A. Beerel, and K. M. Chugg, "Characterizing sparse connectivity patterns in neural networks," in *arXiv:1711.02131*, 2017, submitted for publication to ICLR 2018.
[9] S. Dey, Y. Shao, K. M. Chugg, and P. A. Beerel, "Accelerating training of deep neural networks via sparse edge processing," in *Proc. ICANN*. Springer, 2017, pp. 273–280.
[10] O. Y. Takeshita, "Permutation polynomial interleavers: An algebraic-geometric perspective," *IEEE Trans. Information Theory*, vol. 53, no. 6, pp. 2116–2132, June 2007.
[11] C. Corrada and I. Rubio, "Deterministic interleavers for turbo codes with random-like performance and simple implementation," in *Proc. 3rd Int. Symp. on Turbo Codes and Related Topics*, 2003, pp. 555–558.
[12] S. Dey, "Morse code dataset for artificial neural networks," Oct 2017. [Online]. Available: https://cobaltfolly.wordpress.com/2017/10/15/morse-code-dataset-for-artificial-neural-networks/