AMPNet: Asynchronous Model-Parallel Training for Dynamic Neural Networks

Alexander L. Gaunt, Matthew A. Johnson, Maik Riechert, Daniel Tarlow[*],
Ryota Tomioka, Dimitrios Vytiniotis, Sam Webster

Microsoft Research
Cambridge, United Kingdom

June 23, 2017

Abstract

New types of machine learning hardware in development and entering the market hold the promise of revolutionizing deep learning in a manner as profound as GPUs. However, existing software frameworks and training algorithms for deep learning have yet to evolve to fully leverage the capability of the new wave of silicon. We already see the limitations of existing algorithms for models that exploit structured input via complex and instance-dependent control flow, which prohibits minibatching. We present an *asynchronous model-parallel* (AMP) training algorithm that is specifically motivated by training on networks of interconnected devices. Through an implementation on multi-core CPUs, we show that AMP training converges to the same accuracy as conventional synchronous training algorithms in a similar number of epochs, but utilizes the available hardware more efficiently even for small minibatch sizes, resulting in significantly shorter overall training times. Our framework opens the door for scaling up a new class of deep learning models that cannot be efficiently trained today.

1 Introduction

A new category of neural networks is emerging whose common trait is their ability to react in dynamic and unique ways to properties of their input. Networks like tree-structured recursive neural networks [34, 35] and graph neural networks (GNNs) [31, 21, 13] defy the modern GPU-driven paradigm of minibatch-based data management. Instead, these networks take a tree or a graph as input and carry out a computation that depends on their individual structures. We refer to this new class of models with dynamic control flow as *dynamic neural networks*.

Modern neural network frameworks are certainly capable of expressing dynamic networks. TensorFlow [1] introduces `cond`, `while_loop`, and other higher order functional abstractions, while Chainer [36], DyNet [25], and PyTorch [26] dynamically construct the computation graph using the control flow of the host language. However, *training* these networks with existing software frameworks and hardware can be painfully slow because these networks require highly irregular, non-uniform computation that depends on individual instances. This makes batching impractical or impossible, thus causing the cost of matrix-vector product to be dominated by the cost of loading the weights from DRAM – typically orders of magnitude slower than the peak compute on both CPUs and GPUs [1]. Moreover, these frameworks are not typically optimized with single instance batch size in mind. Dynamically unfolding the computation graph, for example, is a concern when there are not enough instances to amortize the cost for it. [2]

With limited batching, we show in this paper that a way to scale up dynamic models is by exploiting an extreme form of *model parallelism*, amenable to distributed execution on a cluster of interconnected compute devices. By model parallelism, we not only mean computing disjoint parts of the computational graph in parallel, but also computing sequential operations in the graph in a pipeline-parallel fashion [see e.g., 7].

---

[*]Currently at Google Brain

[1] For example, the TitanX GPU performs 10^{13} FLOPS but only 10^{11} floats/s can be brought into the chip due to memory bandwidth (480 GB/s).

[2]Recently proposed TensorFlow Fold [22] mitigates these issues with dynamic batching. (Section 7)

(a) Pipelined execution (#inflight 1)   (b) Pipelined parallel execution (#inflight 4)   (c) Asynchronous model parallelism
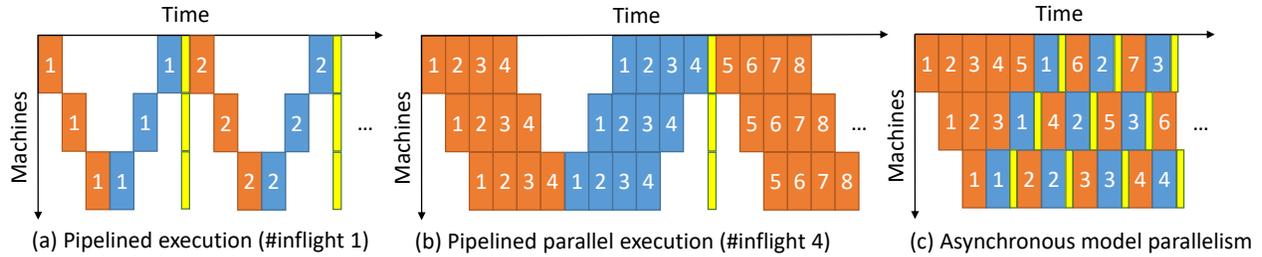
Figure 1: Gantt charts comparing pipelined synchronous model parallelism and asynchronous model parallelism. Orange, blue, and yellow boxes correspond to forward, backward, and parameter update operations, respectively. The numbers in the boxes indicate instance IDs.

Conventional (pipeline) model parallelism, however, can only maximize device utilization if we can keep the pipeline full at all times. Unfortunately, as we show in Figure 1, keeping a conventional pipeline full is at odds with convergence speed due to a decreased parameter update frequency; compare Figure 1 (a) and (b). This is analogous to the trade-off we face in batching. To overcome this problem, we propose *asynchronous model-parallel (AMP) training*, where we allow asynchronous gradient updates to occur, whenever enough gradients have been accumulated; see Figure 1 (c). With this design we aim for both high device utilization and update frequency.

In this setting, however, model parameters may be updated between the forward and the backward computation of an instance, introducing gradient "staleness". Despite staleness, we show that AMP training can converge fast with good hardware utilization. Specifically, our contributions are:

- We present the AMPNet framework for efficient distributed training of dynamic networks.

- We present an intermediate representation (IR) with explicit constructs for branching and joining control flow that supports AMP training. Unlike previous work that considers static computation graphs for static control flow (e.g., Caffe), and dynamic computation graphs for dynamic control flow (e.g., Chainer), our IR encodes a static computation graph to execute dynamic control flow[3]. As a consequence, training becomes easy to distribute and parallelize. Further, IR nodes can process forward and backward messages from multiple instances at the same time and seamlessly support simultaneous training and inference.

- We show that, thanks to explicit control flow constructs, our IR can readily encode *replicas*, a form of *data parallelism* (see Sec. 5). In addition, our IR includes operators for data aggregation which recover forms of *batching*. These features can further improve efficiency, even on CPUs.

- We show that AMP training converges to similar accuracies as synchronous algorithms but often significantly faster. (Sec. 6) For example on the QM9 dataset [30, 27] our implementation of gated graph sequence neural network (GGSNN) [21] on a 16 core CPU runs 9x faster than a (manually optimized) TensorFlow CPU implementation and 2.1x faster than a TensorFlow GPU implementation on the TitanX GPU, because it can better exploit sparsity. Though we do not aim to compete across-the-board with mature frameworks such as TensorFlow, our evaluation proves that AMPNet is particularly beneficial for dynamic networks.

In summary, our work demonstrates the benefits of AMP training and gives a novel way to design and deploy neural network libraries with dynamic control flow. Together these contributions open up new ways to scale up dynamic networks on interconnected compute devices. Inspired by the increasing investment and innovation in custom silicon for machine learning (i.e., FPGAs [11, 5] and ASICs [18]), we perform a simple calculation on the QM9 dataset that shows that AMPNet on a network of 1 TFLOPS devices can be 10x faster than our CPU runtime requiring only 1.2 Gb/s network bandwidth (Sec. 8).

2   Neural networks with complex and dynamic control flow

Below we highlight three models with dynamic control flow, which will be studied in depth in this paper:

---

[3]Our IR bears similarity to TensorFlow but we discuss differences in Section 7.
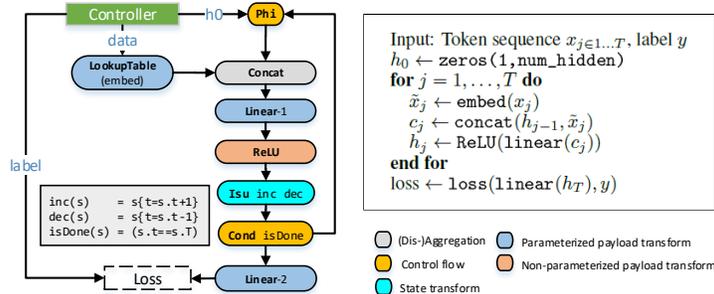
Figure 2: Variable-length RNN in IR and pseudocode (colors denote IR node types)

Variable-length RNNs iterate over the tokens of variable-length sequences. Pseudo-code for a simple RNN is given in Figure 2. The linear layer can be substituted with a more sophisticated unit such as a gated recurrent unit [9]. Though each instance has a different length, it *is* possible to add padding to enable batching. However this may lead to limited speedup due to variability in sequence lengths.

Tree-structured neural networks are powerful models used for parsing of natural language and images, semantic representation, and sentiment analysis [33, 4, 34, 35]. They require evaluation of (potentially multiple) trees with shared parameters but different topology for each instance. Even if one only needs to evaluate a single computational tree per instance as in [34, 35], the tree is instance-specific and batching requires nontrivial planning [22]. A simple form of tree neural network performs a bottom up traversal of the instance, starting from an embedding of the leaves. At each level the values from the child nodes are concatenated and sent through a specialized unit (e.g. LSTM). The result is then propagated further up the tree. Backpropagation over the tree structure is known as backpropagation through structure [14].

Graph neural networks [31, 21, 13] combine both the temporal recurrence in variable length RNN and recurrence over the structure in tree RNN. GNNs can be seen as performing aggregation/distribution operations over a general graph structure with shared parameters.

Apart from the models above, there exist many recently proposed models with flexible control flow (e.g. hierarchical memory networks [6], neural programmer interpreters [29], adaptive computation networks [15, 12], and networks with stochastic depth [16]), to which our framework can be applied.

3   Asynchronous model-parallel training

The basic idea behind AMP training is to distribute a computation graph across compute nodes and communicate activations. For training, the nodes of the computation graph exchange forward or backward messages. Parameterized computations (e.g. fully-connected layers) can individually accumulate gradients computed from backwards messages. Once the number of accumulated gradients since the last update exceeds a threshold `min_update_frequency`, a local update is applied and the accumulator gets cleared. The local parameter update occurs without further communication or synchronization with other parameterized computations. The staleness of a gradient can be measured by the number of updates between the forward and backward computation that produces the gradient. Small `min_update_frequency` may increase gradient staleness. On the other hand, large `min_update_frequency` can reduce the variance of the gradient but can result in very infrequent updates and also slow down convergence. In addition, `max_active_keys` controls the maximum number of *active instances* that are in-flight at any point in time. By setting `max_active_keys` = 1 we restrict to single-instance processing.[4] More in-flight messages generally increase hardware utilization, but may also increase gradient staleness. We have implemented an AMPNet runtime for multi-core CPUs, the details of which are given in Appendix A. Section 6 demonstrates the effects of these parameters.

---

[4]Note this is usually, but not always, equivalent to synchronous training. For example, a single instance can be comprised of a stream of messages (e.g. tree nodes in a tree RNN) and depending on the model some updates may occur asynchronously, even if all the messages in-flight belong to a single instance.

4 A static intermediate representation for dynamic control flow

Overview   Motivated by the need to distribute dynamic networks on networks of interconnected devices and apply AMP training, we have designed a static graph-like *intermediate representation* (IR) that can serve as a target of compilation for high-level libraries for dynamic networks (e.g. TensorFlow or our own frontend), and can itself admit multiple backends (e.g. the multi-core CPU runtime that we consider in detail in this paper, or a network of accelerators). The key feature of our IR is that it is a static graph, but can execute dynamic and instance-dependent control flow decisions.

A neural network model is specified by (i) an IR graph, and (ii) a specialized controller loop that pumps instances and other data – e.g. initial hidden states – and is responsible for throttling asynchrony.

Each IR node can receive and process either forward messages (from its predecessor in the IR graph) or backward messages (from its successors). During training, forward propagation is carried out by passing forward messages through the IR graph. Each message consists of a *payload* and a *state*. The payload is typically a tensor, whereas the state is typically model-specific and is used to keep track of algorithm and control flow information. For example, in a variable-length RNN the state contains the instance identifier, the current position in the sequence, and the total sequence length for the instance. The final loss layer initiates the backward propagation through the IR graph. An invariant of our IR is that for every forward message that is generated by a node with a specific state, this node will eventually receive a backward message with the same state. Depending on `max_active_keys` (Section 3) multiple forward or backward messages can be in-flight, from one or more instances.

In the rest of this section we discuss the most important IR nodes along with their operational semantics, and show how they are used in the example models from the previous section.

Payload transformations   *Parameterized payload transform* (PPT) nodes can be used to encode, for instance, fully connected layers. They apply a transform in the forward pass, but also record the activation in order to use it to compute gradients in the backward pass. An activation is recorded by keying on the state of the message, and hence this state must include all necessary information to allow the node to process multiple messages from potentially different instances without conflating the activations. We require specifications of the forward and the backward transformation, the operation to produce a new gradient, as well as the state keying function to be used. A PPT node may decide to independently apply accumulated gradients to update its parameters. For transformations that do not involve parameters (e.g. ReLUs) our IR offers a simpler *non-parameterized payload transform*.

Loops, state, and control flow   A *condition* node (`Cond` f) is parameterized by a function f that queries the *state* (but not the payload) of the incoming message and, based on the response, routes the input to one of the successor nodes. A *join* node (`Phi`) propagates the messages it receives from each of its ancestor nodes but records the origin so that in the backward pass it can backpropagate them to the correct origin. Like PPT nodes, a `Phi` node must be parameterized over the keying function on the state of the incoming message. An *invertible state update* node (`Isu` f f^{-1}) is parameterized by two functions f and f^{-1} that operate on the state of a message, and satisfy $f^{-1}(f(x)) = x$.

Figure 2 shows how to encode an RNN. The controller pumps sequence tokens into a lookup table – just a PPT node, where the parameter is the embedding table and is also being learned. The controller also pumps labels to the loss layer (dashed boxes are compound graphs whose details we omit), and an initial hidden state h_0 for every sequence. Message states contain the sequence time-step. Following the embedding, messages are concatenated (`Concat` node, see next paragraph) with the hidden state, and the result goes into a linear node followed by a ReLU activation. The `Isu` node increments the time-step, and the conditional node tests whether the end of the sequence has been reached. Depending on the answer it either propagates the hidden state back to `Phi`, or pushes the hidden state to the final linear and loss layers. In backward mode, the gradient is propagated inside the body of the loop, passes through the `Isu` (which decrements the time-step), and reaches the `Phi` node. The `Phi` node will (based on information from the forward phase) either backpropagate to the `Cond` node, or to the controller. Hence the loop is executed in both the forward and backward direction.

Aggregation and disaggregation   Our IR offers several constructs for aggregation and disagreggation; the most important ones are outlined below, and their behavior is summarized in Figure 3. `Concat`, `Split`, and `Bcast` perform concatenation, partition, and broadcast of incoming messages as their names suggest. `Group` can group together several incoming messages based on their state. The output message contains a tensor composed of the
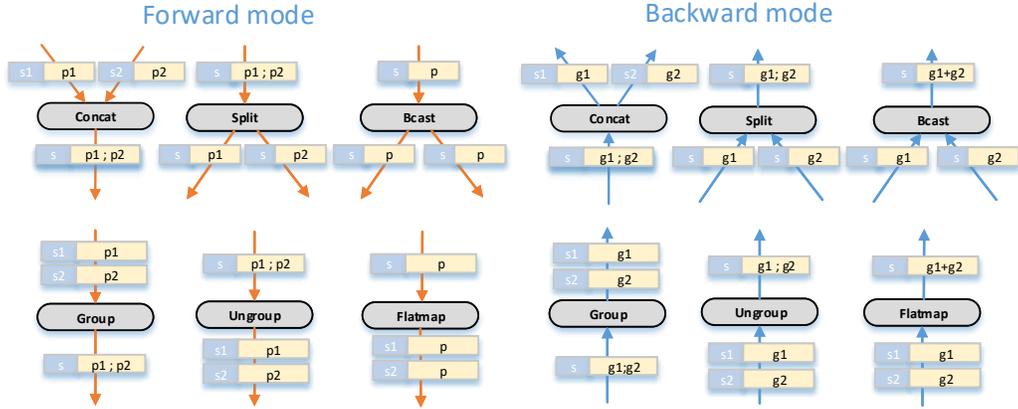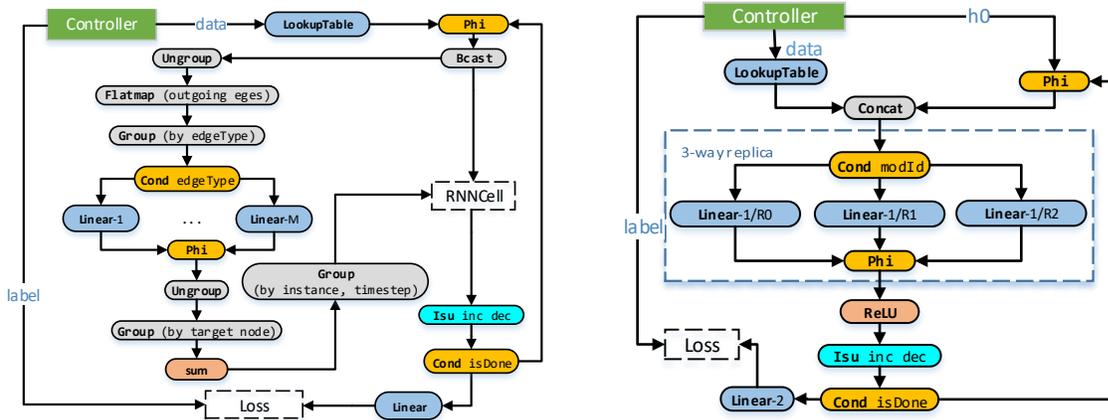
Figure 3: (Dis-)aggregation combinators (forward mode, left; backward mode, right)

(a) Gated Graph Sequence Neural Network. RNNCell is a placeholder for a recurrent structure (e.g. GRU, LSTM), the details of which we omit.

(b) IR graph for an RNN with replicas.

Figure 4: IR graphs for Gated Graph Sequence Neural Network and RNN with replicas.

input payloads, whereas the state is a function of the incoming states. In forward mode `Group` (and also `Concat`) must key on this new state to cache the states of the original messages, so as to restore those in the backward phase. `Ungroup` is a symmetric version of `Group`. `Flatmap` creates a sequence of outgoing messages per incoming message, with replicated payload and new states given by a state generation function that is a parameter of the node. The node keys on the outgoing states and caches the incoming state and number of expected messages, so as to sum all the gradients and restore the original state in backward mode.

Figure 4(a) describes a GNN that combines aggregation on the structure of a graph instance with an outer iteration. The iteration controls in effect the locality of information propagation. The controller pumps data, as before, to a lookup table and labels for this instance to the loss layer. The lookup table emits payloads that are matrices where each row corresponds to the embedding of an instance node, and states that contain the current iteration counter, the instance id, and a reference to the graph structure. The messages are broadcast and ungrouped so that each outgoing message corresponds to each node of the graph instance. Next, each message is goes through a `Flatmap` node that replicates the payload for each *outgoing edge* and creates states that record the incoming node, outgoing node, and type of that edge, resulting in a stream of messages, one for each edge in the graph. Next, all edges are grouped by edge type and each group is sent to a designated linear layer. Each group is then dismantled back and edges are re-grouped by their target node. Each group is passed through a non-parameterized payload transformation that sums together all payloads. The result is a stream of messages where each message contains

5

an aggregated value for a graph node. Finally, we group back all these aggregated values and send the result to the RNNCell for another outer iteration. We note that this constitutes a form of batching – the information about all nodes is batched together before been sent to the RNNCell.

5   Interaction with data parallelism and replicas

Pipeline-style parallelism can often be augmented with forms of data parallelism. Consider the RNN in Fig. 2. The only heavy operation (Linear-1) in the body of the loop is going to act as a bottleneck for computation. One solution is to split the linear layer into smaller tiles and compute them in parallel. This is expressible in our IR but the linear operation needs to be large enough to benefit from tiling in this way. Another approach is to replicate the linear layer in full. Fortunately this requires only minimal new machinery – we can replicate the linear layer and place the replicas inside `Cond` and `Phi` nodes as in Figure 4(b). Different instances or messages from the same instance but with different position in the sequence can be processed in an (pipeline-)parallel fashion using 3 replicas in this case. To enable parameters to be shared among the replicas, we have implemented infrequent end-of-epoch replica synchronization (averaging) to keep the communication cost negligible, as well as a message-passing protocol asynchronous trigger of whole-replica group synchronization, but found that infrequent synchronization was sufficient for fast convergence.

6   Experiments

We evaluate AMPNet using the dynamic models introduced in Section 2. For completeness, we additionally consider a simple multi-layer perceptron (MLP) as an example of a static network with instances that are easy to batch. For each model we select a dataset and compare the throughput and convergence profile of AMPNet against traditional training schemes implemented in TensorFlow.

MLP: MNIST   As preliminary task, we train a 4-layer perceptron with ReLUs on MNIST [20]. We choose 784-dimensional hidden units, and we affinitize the 3 linear operations on individual workers (or threads; see Appendix A). Both AMP runtime and TensorFlow use SGD with learning rate 0.1 and batch size of 100.

RNN: List reduction dataset   As a starting point for experiments on networks with complex control flow we use a synthetic dataset solved by a vanilla RNN. Specifically, we train an RNN to perform reduction operations on variable length lists of digits. Each training instance is a sequence of at most 10 tokens: The first token indicates which of 4 reduction operations [5] is to be performed, and the remaining tokens represent the list of digits. The output is the result of the calculation rounded modulo 10. The dataset consists of 10^5 training and 10^4 validation instances.

We present this task as a classification problem to a vanilla RNN with ReLU activation and a hidden dimension of 128. All parameterized operations are affinitized on individual workers. We bucket training instances into batches of 100 sequences (in the baseline and in AMPNet).

Tree-LSTM: Stanford Sentiment Treebank   As a non-synthetic problem, we consider a real-world sentiment classification dataset [34] consisting of binarized constituency parse trees of English sentences with sentiment labels at each node. Following Tai et al. [35], we use 8,544 trees for training, 1,101 trees for validation, and 2,210 trees for testing.

We use a Tree LSTM for this classification task based on the TensorFlow Fold [22] benchmark model. Both the AMP and Fold models are trained following [35] with the additional architectural modifications proposed by [22, 32]. Furthermore, we split our Tree-LSTM cell into Leaf LSTM and Branch LSTM cells. This does not affect the expressiveness of the model because the LSTM cell receives either zero input (on branch) or zero hidden states (on leaves); i.e., the two cells do not share weights except for the bias parameters, which are learned independently in our implementation. We compare the time to reach 82 % fine grained (5 classes) accuracy (averaged over all the nodes) on the validation set.

---

[5]The operations considered in our toy dataset act on a list L and are expressed in python syntax as: `mean($L$)`, `mean($L$[0::2])-mean($L$[1::2])`, `max($L$)-min($L$)` and `len($L$)`.

Table 1: Time to convergence to target validation accuracy. The time to convergence can be broken down into number of epochs and the throughput (instances/s). The target accuracy is shown inside parentheses next to each dataset. `mak` is a short-hand for `max_active_keys` defined in Sec. 3; `mak = 1` corresponds to synchronous training for MNIST and minimal asynchrony arising from just one in-flight instance for other models with recursive structures.

	AMP				TensorFlow		
	mak	time (s)	epochs	inst/s	time (s)	epochs	inst/s
MNIST (97%)	1	130	4	1949			
	4	44 (3x)	4	5750	**34.5**	3	5880
List reduction (97%)	1	82.9	9	12k			
	4	69.7 (1.2x)	9	14k			
	16	64.9 (1.3x)	9	14k			
(2 replicas)	4	33.7 (2.5x)	10	32k			
(4 replicas)	8	**23.9 (3.5x)**	14	66k	46	7	18k
Sentiment (82%)	1	305	3	88			
	4	230 (1.3x)	3	117			
	16	**201 (1.5x)**	3	133	208	5	265
bAbI 15 (100%)	1	12.2	7	319			
	16	**5.8 (2.1x)**	6	662	6.3	5	350
QM9 (4.6)	4	29k	93	400			
	16	**14k (2.1x)**	69	640	129k	59	58

GNN: Facebook bAbI 15 & QM9 datasets   We verify our GNN implementation using a toy logic deduction benchmark (bAbI task 15 [37]) and study a real-world application for GNNs: prediction of organic molecule properties from structural formulae in the QM9 dataset [30, 27]. GNNs have previously been applied to these tasks in [21] and [13] respectively.

For the bAbI 15 dataset we inflate each graphs from the default 8 nodes to 54 nodes to increase the computational load, but we preserve the two-hop complexity of the deduction task. The architecture of the model follows [21] with a hidden dimension of 5, and 2 propagation steps.

For the QM9 dataset we concentrate on prediction of the norm of a molecule's dipole moment using a regression layer build on the propagation model from [21] (corresponding to the simplest setting in [13]). We use a hidden dimension of 100 and 4 propagation steps, initializing the graph nodes (atoms) following [13]. The molecules contain up to 29 atoms and in a TensorFlow baseline we bucket molecules into batches of 20 with atom counts differing by at most 1 within a batch. Following [13], we report regression accuracies in multiples of a target accuracy from the chemistry community.

Results   On MNIST, Table 1 shows 3x speedup from synchrony (`max_active_keys` = 1) to asynchrony (`max_active_keys` = 4). This is almost ideal as the first three linear layers are the heaviest operations. As we can see in the fourth column of the table, mild asynchrony has negligible effect on the convergence while greatly improving throughput and time to convergence.

The list reduction dataset demonstrates the power of *replicas*. As there is only one heavy operation (Linear-1, Figure 2), the speedup from asynchrony is mild (1.3x). However we get 2.5x and 3.5x speedup for 2 and 4 replicas, respectively, which is nearly ideal. Again, the # of epochs to convergence is not affected by increasing `max_active_keys`. The slowdown in convergence for 4 replicas is due to the increased *effective* minibatch size – also commonly observed in data parallel training.

Next the sentiment tree-RNN dataset shows that our runtime is competitive without batching to TensorFlow Fold [22] using dynamic batching of batch size 100. It is worth mentioning that our runtime allows us to specify different `min_update_frequency` parameter for each parameterized operation. We set this parameter to 1000 for the embedding layer, which is initialized by Glove vectors, and 50 for all other layers. This greatly reduced gradient staleness in the embedding layer.

Finally bAbI 15 (54 nodes) and QM9 datasets demonstrates the importance of sparsity. Note that the TensorFlow implementation of GGSNN [21] implements the message propagation and aggregation over the input graph as a dense $NH \times NH$ matrix multiplication where N is the number of nodes and H is the hidden state dimension. Since each input graph has a unique connectivity, this matrix needs to be constructed for each instance. By contrast, we handle this by message passing and branching as we described in Section 4. As a result we get roughly 9x speedup on QM9 against TensorFlow implementation on CPUs with the same number of threads. Our runtime was also

faster than a GPU TensorFlow implementation by 2.1x. AMPNet and TensorFlow implementation were comparable on the small bAbI 15 (54 nodes) dataset.

Asynchrony   The degree of asynchrony is controlled by hyperparameters `min_update_frequency` and `max_active_keys`. In Fig. 5 we use an 8-replica RNN model on the list reduction dataset to investigate how these parameters affect the data and time required to converge to 96% validation accuracy. We find, in analogy with minibatch size in traditional systems, that `min_update_frequency` must neither be too large nor too small. Increasing `max_active_keys` (increasing asynchrony) monotonically increases performance when the number of keys is similar to the number of individually affinitized heavy operations in the model 8 in this case). Increasing `max_active_keys` significantly beyond this point produces diminishing returns.
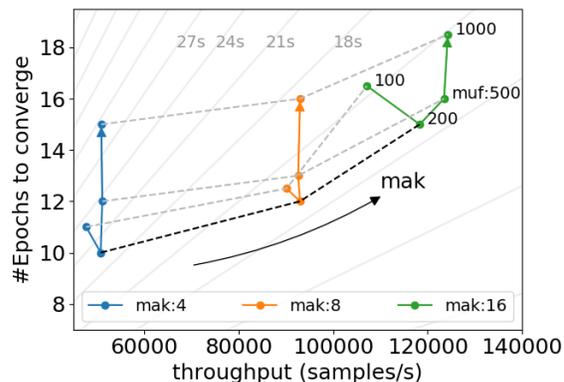

Figure 5: Performance of an 8-replica RNN model on the as a function of asynchrony hyperparameters. Solid gray lines show constant convergence time trajectories. muf stands for `min_update_frequency`.

7   Related Work

Chainer [36], DyNet [25], and PyTorch [26] belong to a new class of deep learning frameworks that define the computation graph dynamically per-instance by executing the control flow of the host language (e.g. Python) that can limit cross-instance parallelism and has a cost that is difficult to hide when the minibatch size is small[see 25]. By contrast our IR graph is static so it is easier to distribute, optimize, and pipeline-parallelize across instances.

Theano [2] and TensorFlow (TF)[1] provide powerful abstractions for conditional execution (`ifelse` in Theano and `cond` in TF) and loops (`scan` and `while_loop`, respectively); TF also provides higher-order functions, such as `map`, `foldl`, `foldr`, and `scan`. The main difference between AMPNet and the above frameworks is that AMPNet is streaming and asynchronous whereas Theano is non-streaming and synchronous. Although not designed for streaming, TF can support streaming programmatically as it exposes first-class queues, as well as data prefetching with so called *input pipelines*. In our IR, all the queuing is implicit and stream-based execution is the default. TF additionally does support static description of dynamic control flow and state update, but we depart from the classic dataflow architecture that TF follows [3]: First, instead of having nodes that represent mutable reference cells, we encapsulate the state with which a message should be processed through the graph in the message itself. Second, because we encapsulate algorithmic state in the messages, we do not introduce the notion of control dependencies (which can be used to impose a specific execution order on TF operations). Our choices complicate algorithmic state management from a programming point of view and make the task of designing a high-level compiler non-trivial, but allow every node to run asynchronously and independently without a scheduler and without the need for control messages: For example, nodes that dynamically take a control flow path or split the data simply consult the state of the incoming message, instead of having to accept additional control inputs. For "small" states (e.g. nested loop counters or edge and node ids) this might be preferable than out-of-band signaling. Our IR can implement loops by simply using state-update, conditional, and phi nodes, because the state accompanies the payload throughout its lifetime, whereas TF introduces specialized operators from timely dataflow [24] to achieve the same effect.

TensorFlow Fold (TFF) [22] is a recent extension of TensorFlow that attempts to increase batching for TF dynamic networks and is an interesting alternative to our asynchronous execution. TFF unrolls and merges together (by depth) the computation graphs of several instances, resulting in a batch-like execution. TFF effectiveness greatly depends on the model – for example, it would not batch well for random permutations of a sequence of operations, whereas our IR would very succinctly express and achieve pipeline parallelism through our control-flow IR nodes.

Asynchronous data parallel training [28, 10, 8] is another popular approach to scale out optimization by removing synchronization, orthogonal to and combinable with model-parallel training. For example, convolutional layers are more amenable to data-parallel training than fully connected layers, because the weights are smaller than the activations. Moreover, when control flow differs per data instance, data parallelism is one way to get an effective minibatch size > 1, which may improve convergence by reducing variance. The impact of staleness on convergence

[28] and optimization dynamics [23] have been studied for data parallelism. It would be interesting to extend those results to our setting.

Jaderberg et al. [17], like us, aim to to train different parts of a model in a decoupled or asynchronous manner. More precisely, their goal is to approximate a gradient with a *synthetic gradient* computed by a small neural network that is locally attached to each layer. Hence, the local gradient calculation becomes independent of other layers (except for the training of the gradient predictor network) and allows asynchronous parameter updates. This would be especially useful if the evaluation of the local network is cheaper than the computation of the real gradient; for example, if the computation of the real gradient required communication of forward/backward messages between devices.

8 Conclusion and Outlook

We have presented an asynchronous model-parallel SGD algorithm for distributed neural network training. We have described an IR and multi-core CPU runtime for models with irregular and/or instance-dependent control flow. Looking forward, we aim to deploy our system on specialized hardware. To give an idea of performant FPGA implementations of AMPNet, we perform a simple estimate of the peak throughput on the QM9 dataset running on a network of 1 TFLOPS FPGAs (see Appendix C for details). Our calculation shows that we achieve 6k graphs/s (10x compared to our CPU runtime) on the QM9 dataset with 200 hidden dimensions and 30 nodes per graph on average. This only requires a very reasonable 1.2 Gb/s network bandwidth. Equally importantly, we plan to build a compiler that automatically deduces the information to be placed in the states and generates state keying functions from a higher-level description of the models. By unlocking scalable distributed training of dynamic models, we hope to enable exploration of this class of models that are currently only on the horizon but may become more mainstream in the future.

Acknowledgements

References

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

[2] R. Al-Rfou, G. Alain, A. Almahairi, C. Angermueller, D. Bahdanau, N. Ballas, F. Bastien, J. Bayer, A. Belikov, et al. Theano: A python framework for fast computation of mathematical expressions. *arXiv preprint arXiv:1605.02688*, 2016.

[3] Arvind and D. E. Culler. Dataflow architectures. *Annual Review of Computer Science*, 1(1):225–253, 1986. URL http://csg.csail.mit.edu/pubs/memos/Memo-261-1/Memo-261-2.pdf.

[4] S. R. Bowman, J. Gauthier, A. Rastogi, R. Gupta, C. D. Manning, and C. Potts. A fast unified model for parsing and sentence understanding. *arXiv preprint arXiv:1603.06021*, 2016.

[5] A. M. Caulfield, E. S. Chung, A. Putnam, H. Angepat, J. Fowers, M. Haselman, S. Heil, M. Humphrey, P. Kaur, J.-Y. Kim, et al. A cloud-scale acceleration architecture. In *Microarchitecture (MICRO), 2016 49th Annual IEEE/ACM International Symposium on*, pages 1–13. IEEE, 2016.

[6] S. Chandar, S. Ahn, H. Larochelle, P. Vincent, G. Tesauro, and Y. Bengio. Hierarchical memory networks. *arXiv preprint arXiv:1605.07427*, 2016.

[7] X. Chen, A. Eversole, G. Li, D. Yu, and F. Seide. Pipelined back-propagation for context-dependent deep neural networks. In *Interspeech*, pages 26–29, 2012.

[8] T. M. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman. Project Adam: Building an efficient and scalable deep learning training system. In *OSDI*, volume 14, pages 571–582, 2014.

[9] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

[10] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le, et al. Large scale distributed deep networks. In *Advances in neural information processing systems*, pages 1223–1231, 2012.

[11] C. Farabet, Y. LeCun, K. Kavukcuoglu, E. Culurciello, B. Martini, P. Akselrod, and S. Talay. Large-scale FPGA-based convolutional networks. *Scaling up Machine Learning: Parallel and Distributed Approaches*, pages 399–419, 2011.

[12] M. Figurnov, M. D. Collins, Y. Zhu, L. Zhang, J. Huang, D. Vetrov, and R. Salakhutdinov. Spatially adaptive computation time for residual networks. *arXiv preprint arXiv:1612.02297*, 2016.

[13] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.

[14] C. Goller and A. Kuchler. Learning task-dependent distributed representations by backpropagation through structure. In *Neural Networks, 1996., IEEE International Conference on*, volume 1, pages 347–352. IEEE, 1996.

[15] A. Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.

[16] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger. Deep networks with stochastic depth. In *European Conference on Computer Vision*, pages 646–661. Springer, 2016.

[17] M. Jaderberg, W. M. Czarnecki, S. Osindero, O. Vinyals, A. Graves, and K. Kavukcuoglu. Decoupled neural interfaces using synthetic gradients. *arXiv preprint arXiv:1608.05343*, 2016.

[18] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, et al. In-datacenter performance analysis of a tensor processing unit. *arXiv preprint arXiv:1704.04760*, 2017.

[19] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL http://arxiv.org/abs/1412.6980.

[20] Y. LeCun and C. Cortes. The MNIST database of handwritten digits, 1998.

[21] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.

[22] M. Looks, M. Herreshoff, D. Hutchins, and P. Norvig. Deep learning with dynamic computation graphs. *arXiv preprint arXiv:1702.02181*, 2017. Sentiment Tree Bank example: https://github.com/tensorflow/fold/blob/master/tensorflow_fold/g3doc/sentiment.ipynb.

[23] I. Mitliagkas, C. Zhang, S. Hadjis, and C. Ré. Asynchrony begets momentum, with an application to deep learning. *arXiv preprint arXiv:1605.09774*, 2016.

[24] D. G. Murray, F. McSherry, M. Isard, R. Isaacs, P. Barham, and M. Abadi. Incremental, iterative data processing with timely dataflow. *Commun. ACM*, 59(10):75–83, Sept. 2016. ISSN 0001-0782.

[25] G. Neubig, C. Dyer, Y. Goldberg, A. Matthews, W. Ammar, A. Anastasopoulos, M. Ballesteros, D. Chiang, D. Clothiaux, T. Cohn, et al. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*, 2017.

[26] PyTorch core team. PyTorch. URL http://pytorch.org/.

[27] R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1, 2014.

[28] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 693–701, 2011.

[29] S. Reed and N. De Freitas. Neural programmer-interpreters. *arXiv preprint arXiv:1511.06279*, 2015.

[30] L. Ruddigkeit, R. Van Deursen, L. C. Blum, and J.-L. Reymond. Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17. *Journal of chemical information and modeling*, 52(11):2864–2875, 2012.

[31] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.

[32] S. Semeniuta, A. Severyn, and E. Barth. Recurrent dropout without memory loss. *arXiv preprint arXiv:1603.05118*, 2016.

[33] R. Socher, C. C. Lin, C. Manning, and A. Y. Ng. Parsing natural scenes and natural language with recursive neural networks. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 129–136, 2011.

[34] R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, C. D. Manning, A. Y. Ng, C. Potts, et al. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, volume 1631, page 1642, 2013. Dataset: `https://nlp.stanford.edu/sentiment/treebank.html`.

[35] K. S. Tai, R. Socher, and C. D. Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.

[36] S. Tokui, K. Oono, S. Hido, and J. Clayton. Chainer: a next-generation open source framework for deep learning. In *Proceedings of workshop on machine learning systems (LearningSys) in the twenty-ninth annual conference on neural information processing systems (NIPS)*, 2015.

[37] J. Weston, A. Bordes, S. Chopra, A. M. Rush, B. van Merriënboer, A. Joulin, and T. Mikolov. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*, 2015.

A   AMPNet runtime implementation

We have implemented an AMPNet runtime for multi-core CPUs. Our runtime spawns multiple *workers* each associated with a hardware thread and hosting one or more IR nodes – in a more general setting each worker corresponds to a compute device. To remain faithful to a distributed environment communication is only through message passing. Each worker is equipped with a multiple-producer single-consumer queue that can accept messages for any IR node hosted on that worker.

The main worker loop periodically offloads messages from the concurrent queue to a worker-local priority queue that assigns higher priority to backward messages. Backward prioritization is designed for situations when multiple IR nodes with a dependency on the IR graph end up hosted on the same worker. As a consequence, backpropagation can complete faster and new instances can be pumped in by the controller. We dequeue the top message and invoke the forward or backward method of the target IR node. These methods may update internal IR node state (such as cache the state of the incoming message and wait for more messages) or post new forward or backward messages.

How to update the parameters using the gradients is a configuration option that selects amongst a range of optimization algorithms. We have implemented runtime configuration options for selecting several well-known schemes such as (momentum-)SGD and Adam [19], and for controlling the training hyper-parameters.

B   Details of the experimental results

We provide more details of the experiment and analysis in this section. All experiments were carried out on machines with 16 cores and 112 GB of RAM. The validation curves were averaged over at least 20 independent runs. The time/epoch to reach a target accuracy was calculated as *median* of the time an algorithm takes to reach the target accuracy over the repetitions. We found this approach to be more reliable than reporting the time/epoch when the averaged accuracy reaches the target. Table 2 show both the training and validation throughputs we obtained with AMPNet and our TensorFlow baselines.

B.1   MNIST

Figure 6(a) shows the validation accuracy vs. time, validation accuracy vs. epochs, and throughputs of synchronous and asynchronous versions of AMPNet as well as TensorFlow. The throughput greatly increases from synchronous (`max_active_keys` = 1) to asynchronous (`max_active_keys` = 4) while the speed of convergence (middle panel) is hardly affected for mild amount of asynchrony. Taking higher `max_active_keys` = 8 increase throughput only very little (because there is no more work) and seems to rather make the convergence more unstable. This is due to the fact that our current scheduler is greedy and pumps in a forward message whenever the first layer is unoccupied, which leads to large gradient staleness. Clearly a better scheduling will remove this sensitivity.

B.2   List reduction dataset

Similarly Figure 6(b) shows the validation accuracy vs. time and the number of epochs, and throughputs of the methods we discussed in the main text on the list reduction dataset. We first notice that increasing the asynchrony from synchronous (`max_active_keys=1`) to `max_active_keys` = 4 and `max_active_keys` = 16 affects the convergence very little at least in average. However, there is also very little speedup without introducing replicas as we discussed in the main text. Increasing the number of replicas increases the throughput almost linearly from 15k sequences/s (synchronous) to 30k sequences/s (2 replicas) and over 60k sequences/s (4 replicas). Convergence is almost unaffected for 2 replicas. This was rather surprising because the parameters of the replicas are only synchronized after each epoch as we described in Sec. 5. A slight slow-down in convergence can be noticed for 4 replicas. Since even `max_active_keys` = 16 has almost no effect on the convergence without replicas, this is not due to asynchrony. We also tried to synchronize more frequently but this did not help. Thus we believe that the slow-down is due to the increase in the effective minibatch size resulting in reduced number of updates per epoch, which is commonly observed in data parallel training.

Table 2: Training and validation throughputs.

	number of instances		AMP			TensorFlow	
	train	valid	mak	train inst/s	valid inst/s	train inst/s	valid inst/s
MNIST (97%)			1	1948	6106		
	60k	10k	4	5750	19113	5880	8037
List reduction (97%)			1	12k	41k		
			4	14k	53k		
			16	14k	53k		
(2 replicas)			4	32k	99k		
(4 replicas)	100k	10k	8	66k	181k	18k	43k
Sentiment (82%)			1	88	326		
			4	117	568		
	8511	1101	16	133	589	265	1583
bAbI 15 (100%)			1	319	733		
	100^6	1000	16	662	1428	350	1093
QM9 (4.6)			4	400	970		
	117k	13k	16	640	1406	57.5	104

B.3 Sentiment Tree Bank dataset

Figure 6(c) shows the averaged fine grained validation accuracy for the tree RNN model with different `max_active_keys` on the Stanford Sentiment Tree Bank dataset. Interestingly although TensorFlow Fold achieves higher throughput, AMPNet converges faster (in terms of the number of epochs). This speedup is mainly due to the fact that we are not batching and updating whenever we have accumulated 50 gradients (except for the lookup table node that updates every 1000 gradients); 50 gradients correspond to roughly 2 trees. The reason for the lower throughput compared to TensorFlow Fold is that we are only grouping the leaf operations and not the branch operations. Grouping the branch operations is possible by extending our IR nodes and we are actively working on it.

Figure 6(d) shows the same information for fixed `max_active_keys` = 16 and different `min_update_frequency`. We can see that as we increase `min_update_frequency` from the originally used 50 to larger values, the peak of the validation accuracy shifts later and lower becoming closer to the curve obtained by TensorFlow Fold. This is consistent with the parallels between `min_update_frequency` and minibatch size we drew in Section 6. The `min_update_frequency` parameter has marginal influence on the training throughput.

B.4 bAbI 15 (54 nodes) and QM9 dataset

Figures 6(e) and 6(f) show that GGSNN can tolerate relatively large `max_active_keys` = 16. In particular, on the more challenging QM9 dataset taking `max_active_keys` = 16 increased the throughput significantly from 152 graphs/s (synchronous) to 640 graphs/s.

C Throughput calculation for the GGSNN model for QM9

Suppose that the hidden dimension H is sufficiently wide so that the speed of matrix-vector product dominates the throughput of the system compared to element-wise operations, such as sigmoid and tanh; we take $H = 200$ in the calculation below.

In an idealized scenario, pipeline parallel execution of the network consists of roughly 3 stages per time step. In the first stage, all four $H \times H$ linear nodes corresponding to different edge types execute in parallel. In the second stage, the two $2H \times H$ linear nodes (#9 and #12) inside the GRU cell corresponding to update and reset gates execute in parallel (see Fig. 7,). Finally, the last $2H \times H$ linear node in the GRU cell immediately before the Tanh node executes. We would need at least 7 devices that executes these linear nodes in a pipelined parallel fashion. The memory requirement for each device is 4 times the size of the $H \times H$ or $2H \times H$ weight matrix, which consists of the

---

[6]We sample 100 fresh samples for every epoch.

parameter, gradient buffer, and two slots for the statistics that need to be accumulated in the Adam optimizer. This would be 1.2MB for $H = 200$ and float32.

The throughput of training this model is either limited by the speed of the GRU block or that of the linear nodes corresponding to edges. The number of operations in the forward and backward passes per time step can thus be estimated as

$$\text{fwdop} = 2 \cdot \max(2NH^2, EH^2/C),$$
$$\text{bwdop} = 6 \cdot \max(2NH^2, EH^2/C),$$

where N and E are the average number of nodes and edges per instance, respectively and C is the number of edge types, which is 4 in this task. We assume that the backward operation is 3 times more expensive than the forward operation because it requires matrix transpose, matrix multiplication, and gradient accumulation.

Moreover, in an idealized scenario, we can expect that each neural network node alternates between forward and backward (we thank Vivek Seshadri for pointing this out).

Thus we can estimate the throughput of training this model on a network of 1 TFLOPS devices (e.g., Arria 10) as

$$\text{throughput (samples/s)} = 0.5 \cdot \frac{10^{12}}{(\text{fwdop} + \text{bwdop}) \cdot 4},$$
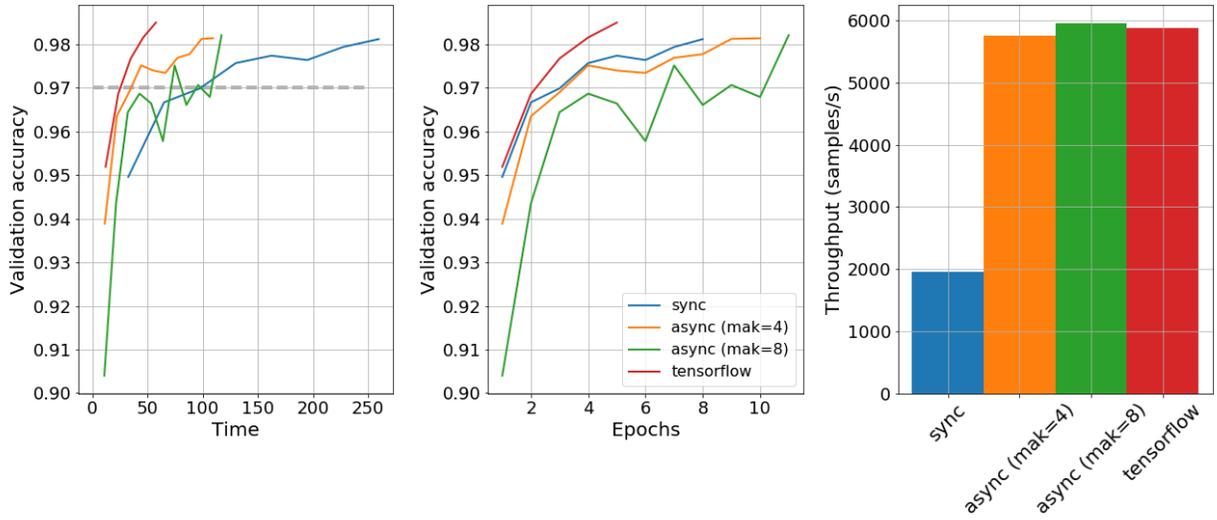
where the last 4 is the number of propagation time steps and 0.5 accounts for all the other operations and communication overhead we ignored in the calculation.

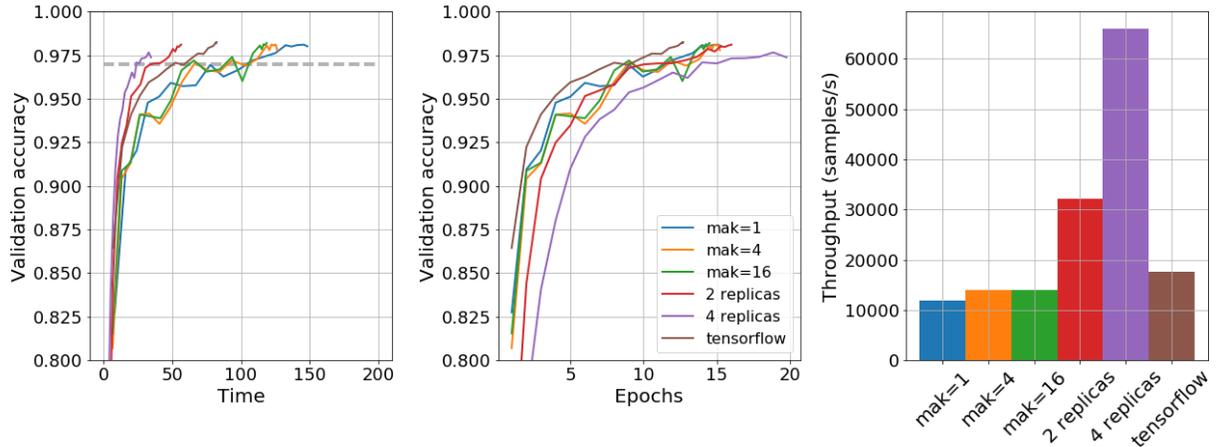For $H = 200$, $N = E = 30$ and $C = 4$ we obtain

$$\text{throughput (samples/s)} = 0.5 \cdot \frac{10^{12}}{64 \cdot NH^2} \simeq 6.5 \cdot 10^3 \text{ (samples/s)}.$$

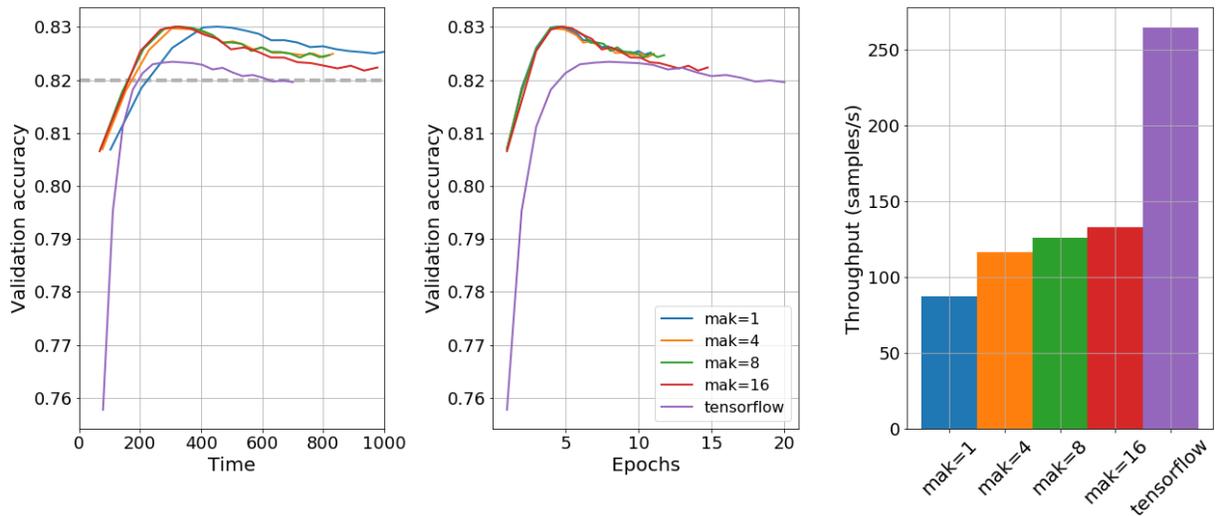The network bandwidth required in this scenario is

$$\text{network bandwidth (bits/s)} = 32 \cdot \text{throughput} \cdot \max(N, E) \cdot H = 1.2 \cdot 10^9 \text{ (bits/s)}.$$
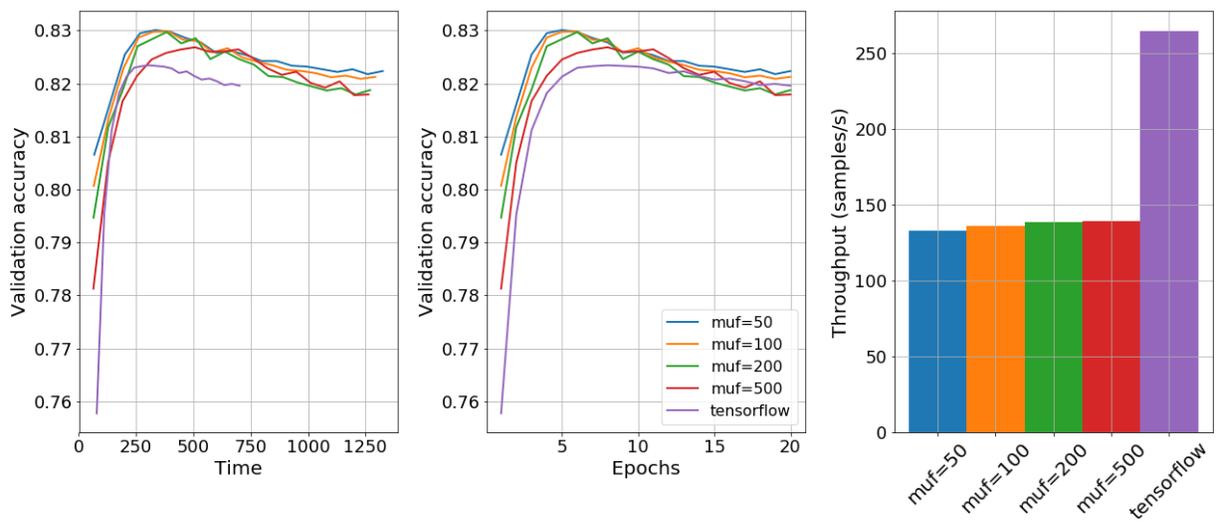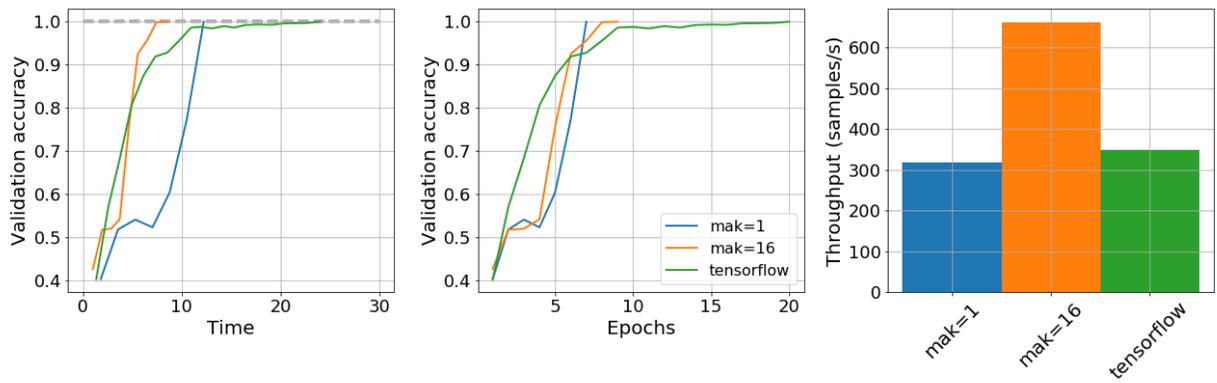
(a) MNIST dataset
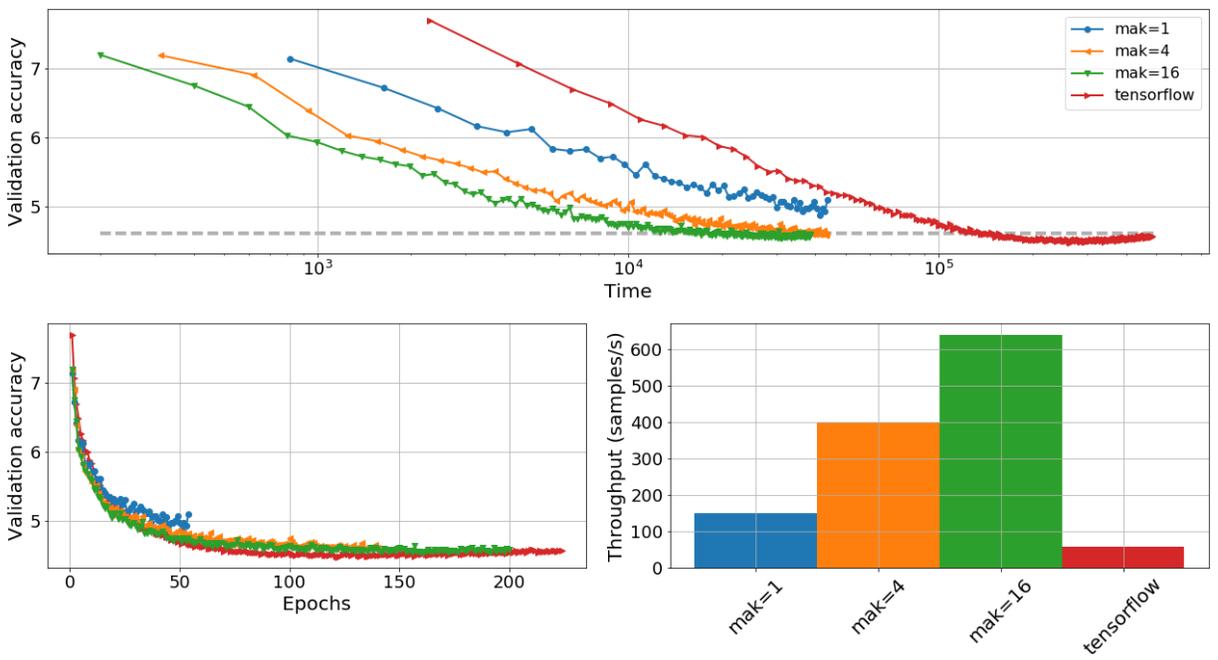

(b) List reduction dataset

(c) Sentiment Tree Bank (`min_update_frequency` = 50)

(d) Sentiment Tree Bank (`max_active_keys` = 16)
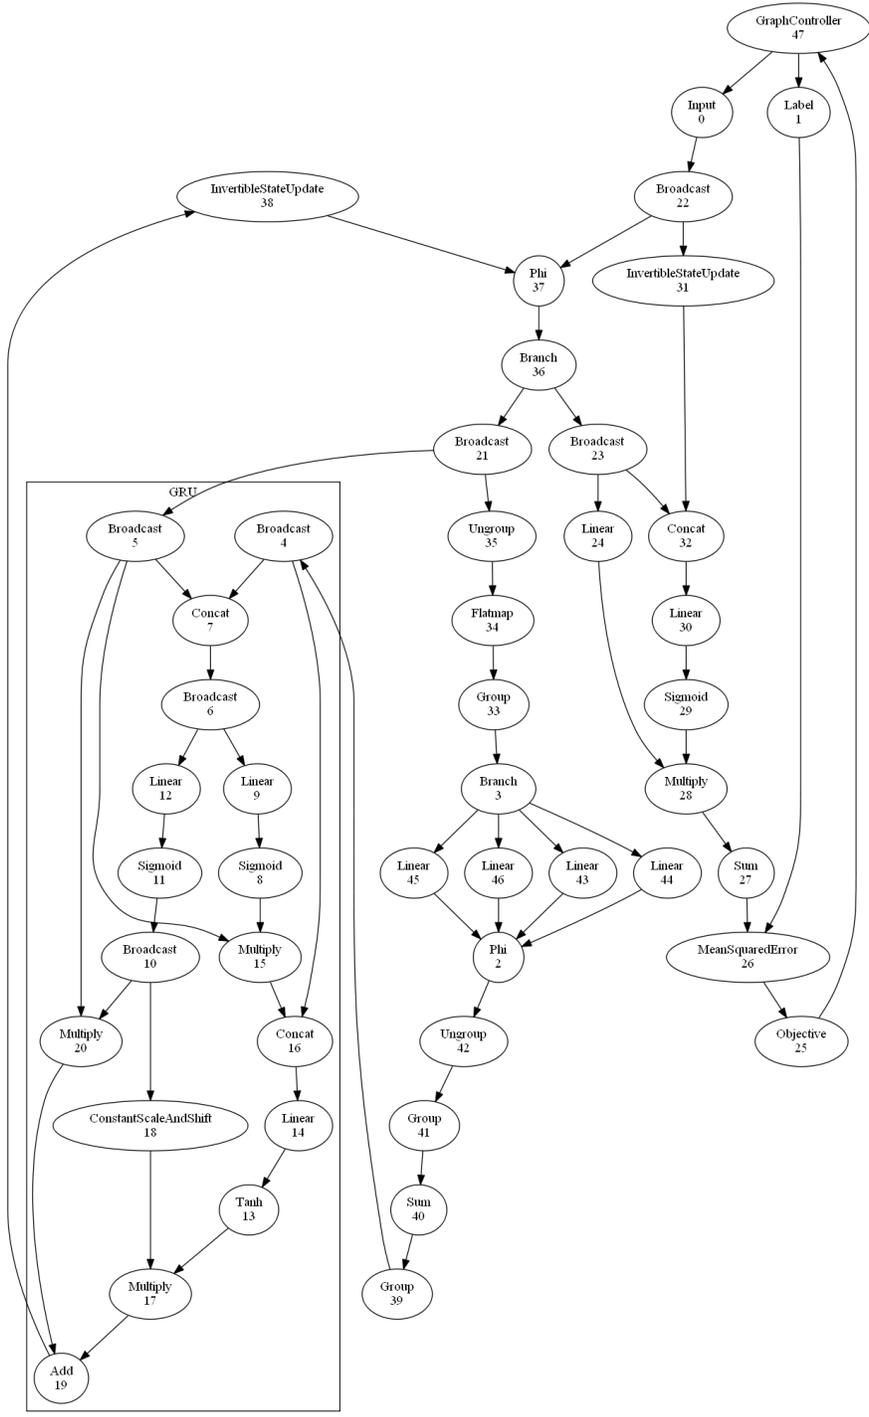
(e) bAbI 15 (large) dataset

(f) QM9

Figure 6: Convergence plots.

Figure 7: GGSNN IR graph for QM9.