



Accelerating Inference and Language Model Fusion of Recurrent Neural Network Transducers via End-to-End 4-bit Quantization

Andrea Fasoli, Chia-Yu Chen, Mauricio Serrano, Swagath Venkataramani,
George Saon, Xiaodong Cui, Brian Kingsbury, Kailash Gopalakrishnan

IBM Research, USA

{andrea.fasoli, swagath.venkataramani}@ibm.com
{cchen, mserrano, gsaon, cuix, bedk, kailash}@us.ibm.com

Abstract

We report on aggressive quantization strategies that greatly accelerate inference of Recurrent Neural Network Transducers (RNN-T). We use a 4 bit integer representation for both weights and activations and apply Quantization Aware Training (QAT) to retrain the full model (acoustic encoder and language model) and achieve near-iso-accuracy. We show that customized quantization schemes that are tailored to the local properties of the network are essential to achieve good performance while limiting the computational overhead of QAT.

Density ratio Language Model fusion has shown remarkable accuracy gains on RNN-T workloads but it severely increases the computational cost of inference. We show that our quantization strategies enable using large beam widths for hypothesis search while achieving streaming-compatible runtimes and a full model compression ratio of $7.6\times$ compared to the full precision model.

Via hardware simulations, we estimate a $3.4\times$ acceleration from FP16 to INT4 for the end-to-end quantized RNN-T inclusive of LM fusion, resulting in a Real Time Factor (RTF) of 0.06. On the NIST Hub5 2000, Hub5 2001, and RT-03 test sets, we retain most of the gains associated with LM fusion, improving the average WER by $>1.5\%$.

Index Terms: RNN-T, quantization, reduced precision, INT4, density ratio language model fusion

1. Introduction

End-to-end (E2E) models such as the Recurrent Neural Network Transducer (RNN-T) [1] map an acoustic feature sequence \mathbf{x} to a sequence of characters, sub-words, or words \mathbf{y} , directly defining the posterior probability $P(\mathbf{y}|\mathbf{x})$ without additional conditional independence assumptions. They dramatically simplify both training and decoding pipelines compared to competing approaches for Automatic Speech Recognition (ASR) and demonstrated excellent performance on standard benchmarks [2–6].

RNN-T models have several properties that make them attractive for ASR, such as the lack of a need for pre-existing alignments in training, the incorporation of a recurrent network Language Model (LM) (which is jointly trained with the acoustic encoder), and a computation which is monotonic in time suitable for streaming recognition [7, 8].

In order to boost accuracy, separate LMs, which provide an estimate of $P(\mathbf{y})$, can be trained on larger text-only corpora and combined with the RNN-T outputs using shallow fusion. In particular, Density Ratio fusion [9, 10] is an inference-only method which extends popular shallow fusion methods. Yet, shallow fusion can result in a dramatic increase in model size and computational time.

In this work, we examine mixed-precision quantization strategies, primarily at 4 bits, based on Quantization Aware

Training (QAT) to alleviate the challenges posed by LM fusion approaches. With QAT, the network is retrained while simulating low-precision operations, allowing it to learn to compensate for the error introduced by the quantization process. The resulting model can be efficiently compressed and accelerated at inference time, ideally with minimal or no loss in accuracy.

Previous works on the acceleration of E2E ASR models for inference tasks demonstrated minimal degradation with 8-bit quantization of the matrix multiplications (both weights and activations), applied to both the acoustic encoder and the LM [11–14]. In addition, ultra low precision (1 or 2 bits) representation of the LMs weights alone has been extensively investigated [15–18]. To the best of our knowledge, only [19] has reported on 4 bit quantization of weights and activations of an E2E ASR model (RNN-T), showing a 1.3% average WER degradation on the Hub5 2000 test set.

Our main contributions are summarized as follows:

- We carefully assign quantizers and precisions to both weights and activations of the RNN-T. We stress that while weight quantization is sufficient for model compression, in order to achieve a significant acceleration, quantization of the activations is also a requirement. Our selections are driven by the minimization of WER and by the need to keep the process practical by limiting the retraining time imposed by QAT. Our scheme achieves near-iso-accuracy at 4 bits.
- In the context of Density Ratio LM fusion, we fully quantize the separate LMs and demonstrate that the accuracy gains due to fusion translate well to the E2E-quantized RNN-T with quantized additional LMs.
- Via hardware (HW) runtime simulations, we show that the workload can be effectively accelerated, enabling hypothesis search at beam width 16 to achieve a Real Time Factor (RTF, defined as processing time over audio duration) competitive with small beam width search using non-quantized models without LM fusion.

2. RNN-T with LM fusion

The base RNN-T (Fig. 1) consists of an acoustic encoder, a prediction network, and a joint network. The encoder comprises 6 bi-directional LSTM layers (input size 340 - including 100-dim i -vectors) and a linear layer that generates the latent acoustic representation h_t^{enc} , of dimension 256. The prediction network is a single unidirectional LSTM layer (hidden size 768), preceded by a small embedding layer, and followed by a linear layer with output h_u^{dec} also of size 256. Encoder and prediction network outputs are combined multiplicatively in a joint network [5], with an additional linear layer and log-Softmax over 46 output characters. In this configuration, the RNN-T comprises 57.2 M parameters (54.6 M in the encoder, 2.6 M in the prediction

network). We train for 20 epochs with 64-utterance batches on audio and character-level transcripts from the SWB-1 data collection, augmented with speed and tempo perturbation [20], SpecAugment [21], and Sequence Noise Injection [22]. We use the AdamW optimizer and the OneCycleLR learning rate scheduler.

For density ratio LM fusion, we train an external LM (LM_{EXT}) on the 2000 hours Switchboard + Fisher acoustic transcripts and a source LM (LM_{SRC}) on the 300 hours Switchboard transcripts. LM_{EXT} consists of an embedding layer, two unidirectional LSTM layers (hidden size = 2048), a bottleneck layer of size 256, and an output layer of size 46 (see Fig. 1). The total number of parameters is 51.0 M, almost $20\times$ larger than the original prediction network. LM_{EXT} is trained with Nesterov SGD for 40 epochs, at a constant learning rate (LR) = 0.03 for 20 epochs, then exponentially decreased by $1/\sqrt{2}$ /epoch. LM_{SRC} has the same building blocks and size as the RNN-T prediction network (2.6 M parameters) and is trained for 40 epochs.

At inference, the symbol predicted at the previous step y_{u-1} is fed back into the prediction network, LM_{EXT} , and LM_{SRC} , generating 3 prediction scores that are combined as [9]:

$$S = \log P(\mathbf{y}|\mathbf{x}) + \mu \log P_{ext}(\mathbf{y}) - \lambda \log P_{src}(\mathbf{y}) + \rho|\mathbf{y}| \quad (1)$$

where S is the combined score, P , P_{ext} , and P_{src} the probability of the original RNN-T network, LM_{EXT} , LM_{SRC} , respectively, and μ , λ , and ρ are the LM weights and character insertion reward, set to 0.7, 0.5, 0.2, respectively.

We use alignment-length synchronous beam search [23] with variable beam width (1 to 16), corresponding to the maximum number of hypotheses that are retained and evaluated at each iteration of the decoding loop. Decoding is performed on the Hub5 2000, Hub5 2001, and RT-03 test sets and Kaldi scoring is used for measuring WER.

3. Quantization strategies

3.1. Background

Quantization is the process of mapping high-precision floating point numbers into a lower bit representation. In particular, low-precision arithmetic, which expresses values as integers multiplied by a scaling factor, is well suited for AI-accelerating HW and enables faster and highly energy efficient execution of Multiply Accumulate (MAC) operations and data transfer. Linear Quantization (LQ) is the most widely adopted approach in QAT, as it allows for a direct implementation of integer arithmetic. Asymmetric LQ computes:

$$x_{int} = \text{clamp}(\lfloor \frac{x}{s} \rceil + z; 0; 2^b - 1) \quad (2)$$

$$x_q = s(x_{int} - z) \quad (3)$$

where x is the input to be quantized, x_{int} its integer version, x_q the discrete representation on the same scale of x , b the number of bits, $\lfloor \cdot \rceil$ the rounding operator, $s = \frac{\alpha_+ - \alpha_-}{2^b - 1}$ a "scale" that is function of the boundaries α_{\pm} , and $z = -\alpha_-/s$ is the "zero point", that aligns x_{int} to zero for $x = \alpha_-$. The clamp function restricts x_{int} range to $[0, 2^b - 1]$.

Symmetric LQ, which relies on a single boundary α and no zero point, is a less flexible representation but naturally maps zero inputs exactly to an integer in x_{int} and, from a HW perspective, avoids the overhead incurred by performing MAC operations between values solely quantized with asymmetric LQ [24].

The quantization process establishes an inherent trade-off between clipping error and rounding error. The former arises

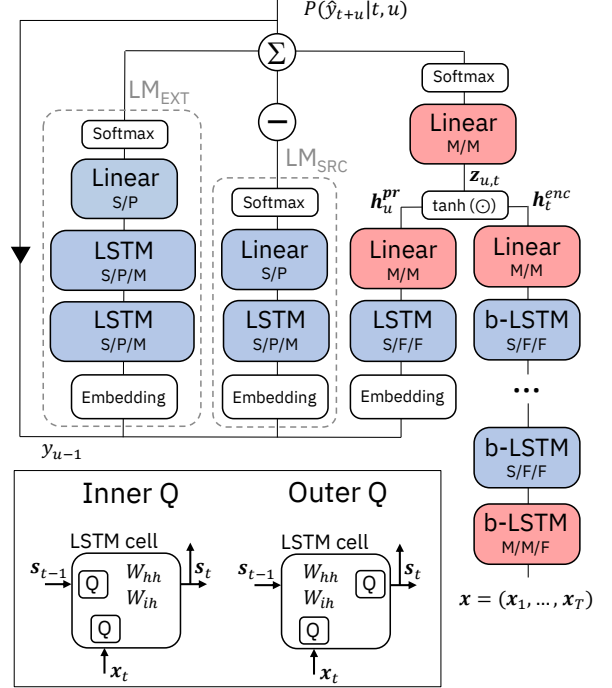


Figure 1: Architecture and main quantization scheme of RNN-T with LMs. Red cell: 8 bits; blue cell: 4 bits. Layers are labelled with their respective quantizers (Linear: weight/activation; LSTM weight/input/hidden state). Abbreviations: F = FIX, M = MAX, P = PACT, S = SAWB, b-LSTM = bidirectional LSTM. Inset: Inner vs. Outer quantization scheme of inputs and hidden states of LSTM cell

by the approximation of the input values x that lay outside the range $[\alpha_-, \alpha_+]$. The latter relates to the residual between each input x and its rounded form x_q . For inputs within the clipping range, the residual is upper bounded by $s/2$, which increases as the $[\alpha_-, \alpha_+]$ range widens.

Different quantizers can be chosen to determine the most appropriate α_{\pm} for a given weight or activation tensor. The simplest, herein named FIX, relies on a predetermined choice for α_{\pm} which remain fixed during QAT and inference. This is the fastest quantizer for QAT as it does not utilize any analysis of the distribution statistics nor gradient propagation on α_{\pm} . However, it does not adapt to the input distribution and the choice of α_{\pm} may not be well motivated and sub-optimal. FIX is only applicable in specific cases, as discussed later.

A more common strategy, herein called MAX, is to set α_{\pm} to the extremes (min and max) of the distribution. MAX captures the whole range of values with no clipping error but suffers in the presence of outliers that increase the rounding error.

Statistics Aware Weight Binning (SAWB) [25] is a symmetric quantizer for weights that leverages the distribution statistics to identify the boundary α that most closely approximates the optimal α^* that minimizes the mean square error (a combination of clipping and rounding error).

With Parameterized Clipping Activation (PACT) [25], the boundaries α_{\pm} are independently learned during QAT. Learned quantizers require tracking the quantization operations, computing the corresponding gradients, and updating α_{\pm} . Consequently, they are more computationally expensive compared to non-learned quantizers like FIX, MAX, or SAWB. The gap in

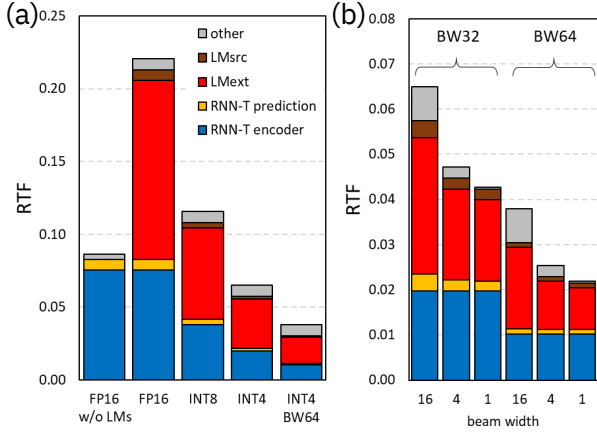


Figure 2: Real Time Factor (RTF) of quantized RNN-T decoding with density ratio LM fusion on simulated custom HW: (a) impact of reduced precision, at beam width 16 (BW64 = 64 Gbps bandwidth); (b) beam width-dependence of 4-bit model, at 32 and 64 Gbps. Legend applies to both figures.

computation time may be onerous in the QAT of large ASR models, due to the repeated calls to the activation quantizers.

3.2. Model quantization

Fig. 1 presents our main quantization scheme, as applied to an RNN-T model inclusive of LM fusion (LM_{EXT} and LM_{SRC}). Weights (without biases) and activations are quantized across the whole network. INT4 quantization is shown in blue, INT8 in red. The quantizers are indicated by the abbreviations within each cell (F = FIX, M = MAX, P = PACT, S = SAWB). For LSTM layers, 3 quantizers are used, for weights, inputs, and hidden states. The LSTM cell states are not quantized as they only participate in element-wise computations and their contribution to the total operations is minimal. In the linear layers, weights and the inputs are quantized. To limit HW overhead during MAC operations [24], symmetric quantizers are used for the weights across the whole network. Re-training the RNN-T with QAT is computationally expensive. Quantization of activations is critical because RNN models require quantizers at every time step. This is not the case for the weights, which can be quantized just once per utterance. Therefore, practical considerations drive the selection towards less expensive activation quantizers, provided that accuracy is not impacted.

With reference to Fig. 1, all the bidirectional LSTM layers of the RNN-T acoustic encoder *except the first* are quantized at 4 bits. We use SAWB to determine an optimal α boundary for the weights, as the rounding error introduced by the MAX quantizer at 4 bits is detrimental to accuracy. For inputs and hidden states we use the FIX quantizer in a shared configuration across time steps. In the LSTM layers, inputs and hidden states are the results of the computation taking place at the previous layer or time step, respectively, and in particular of the Hadamard product of two activation functions: $\tanh(\cdot) \odot \sigma(\cdot)$. Consequently, each tensor element is in $[-1,1]$ (before dropout scaling). Therefore, we initialize the α_{\pm} boundaries of these FIX quantizers to ± 1 . We also note that the upper boundary obtained with a MAX quantizer when utilized in this scenario typically falls in the interval $[0.9,1]$, which further supports the selection of the chosen boundaries of the FIX quantizer.

The first layer of the network has major impact on the accuracy, while accounting for $< 10\%$ of parameters and computa-

tions of the encoder. Therefore, we use 8-bit quantization on this layer, with a MAX quantizer for both weights and inputs. 8-bit MAX provides sufficient granularity to limit the rounding error in the presence of outliers, while not restricting the range to an arbitrary value. On the other hand, 8-bit FIX with $\alpha_{\pm} = \pm 1$ remains appropriate for the hidden states. On the prediction side, we quantize the LSTM layer to 4 bits using SAWB for weights, and FIX for inputs ($\alpha_{\pm} = \pm 1.25$) and hidden states ($\alpha_{\pm} = \pm 1$). The remaining linear layers are quantized at 8 bits using MAX for both weights and activations. At 4 bits, they bring substantial degradation while accounting for a small fraction of parameters and computations ($< 1\%$).

The re-training of the LM models is less computationally expensive and affords for more flexibility in the choice of quantizers. Therefore, we fully quantize both LM_{EXT} and LM_{SRC} at 4 bits, using SAWB for weights and PACT for inputs. With this configuration, we achieve a $7.6\times$ reduction in model size compared to the FP32 model, with 94.9% of parameters at 4-bits.

3.3. Quantization Aware Training

The RNN-T, LM_{EXT} , and LM_{SRC} are re-trained separately. The QAT process mirrors the non-quantized training, with a few important differences. First, we adjust the initial LR and LR schedule to ensure the model can properly learn to compensate for the error introduced by the quantization. Following [19], we lower the initial RNN-T LR to $4 \cdot 10^{-4}$ and decrease it linearly to a lower bound of 10^{-5} . Then, we shorten the QAT to < 15 epochs and train with batch size 768. For QAT of the LMs we retain the same LR schedule but lower the initial LR to $8 \cdot 10^{-4}$.

3.4. Inner vs. Outer quantization

It is common to conceptualize the quantization as “per layer”, where the quantizers convert both the weights and the input tensors into their discretized form. In RNNs, the time dependence also requires activation quantization within layers, across time steps. One option is to quantize the inputs x_t and the hidden states s_{t-1} as they “enter” the RNN cell, before performing the MAC with the corresponding weights W_{ih} and W_{hh} . We refer to this approach as “Inner” quantization (fig. 1 inset). The number of quantization calls throughout a network is $2 * T * L$ for 2 input tensors, with T the number of time steps and L the number of layers (counted twice for bi-dir RNNs).

This approach is inefficient. For each cell, the same output tensor s_t is provided to both the next time step and layer. Therefore, we quantize it once, at the cell output (fig. 1 inset). We call this process “Outer” quantization. Assuming s_{t-1} is zero at $t = 1$, the number of quantization calls that Outer quantization requires is reduced to $T + T * L$. The ratio between Outer and Inner quantizer calls is $(L + 1)/(2 * L)$, independent of the sequence length. Outer quantization reduces the number of calls to activation quantizers in the RNN-T encoder by 45%.

4. Experimental results

4.1. Runtime performance and acceleration

We perform a comparative evaluation of the runtime performance of the RNN-T with LM fusion. Following [19, 26, 27], we simulate the workload on a custom accelerator consisting of a coprocessor and attached CPU [28]. The compute-heavy data-parallel operations are executed on the coprocessor, while control-heavy operations such as hypothesis sorting are mapped to the CPU. We evaluate accelerations against an FP16 baseline

Table 1: Network acceleration (vs. FP16), compression ratio (vs. FP32), and WER at various quantization stages of RNN-T and LMs. Quantization is expressed as bits of weights/activations. Beam width (bm) is 16 unless otherwise stated. Δ is WER gap against a model with equal RNN-T quantization and no LM fusion. All networks are initialized with FP32 pre-trained weights

		Hub5 2000						Hub5 2001						RT 2003						
RNN-T	LM _{EXT}	LM _{SRC}	accel. (v. FP16)	size (MB)	compr. (v. FP32)	SWB	CH	avg	Δ	avg bm8	avg bm4	SWB	S2P3	S2P4	avg	Δ	SWB	FSH	avg	Δ
32/32	-	-		223.3		8.1	15.5	11.8		11.8	11.9	8.5	11.7	15.9	12.1		18.5	11.8	15.3	
32/32	32/32	32/32		432.5		6.3	13.1	9.8	-2.0	10.0	10.6	7.1	9.4	13.6	10.1	-2.0	15.4	9.5	12.5	-2.8
4/4	-	-	3.4	30.9	7.2	8.4	16.4	12.4		12.6	12.6	8.9	12.3	16.7	12.8		19.9	12.9	16.5	
4/4	32/32	32/32	1.4	240.1	1.8	6.8	14.3	10.6	-1.8	10.8	11.2	7.6	10.2	14.2	10.8	-2.0	17.0	10.5	13.8	-2.7
4/4	8/8	32/32	2.2	90.9	4.8	6.8	14.3	10.6	-1.8	10.8	11.3	7.6	10.5	14.3	10.9	-1.9	17.1	10.6	14.0	-2.5
4/4	4/4	32/32	3.1	66.0	6.6	6.9	14.5	10.7	-1.7	10.9	11.4	8.0	10.7	14.8	11.2	-1.6	17.6	10.9	14.4	-2.1
4/4	4/4	4/4	3.4	57.2	7.6	7.0	14.4	10.7	-1.7	11.0	11.3	8.1	10.6	14.9	11.3	-1.5	17.4	10.9	14.3	-2.2
2/4	-	-		18.3	12.2	9.0	17.4	13.2		13.3	13.3	9.6	13.1	17.5	13.5		21.0	13.7	17.5	
2/4	4/4	4/4		44.6	9.7	7.4	15.2	11.4	-1.8	11.5	11.9	9.8	12.5	17.1	13.2	-0.3	19.7	12.8	16.3	-1.2

(instead of FP32) due to restrictions in our HW simulator.

Fig. 2 compares the estimated RTF at different precisions, for a single 152-frame sequence with beam width 16. In the absence of additional LMs, the FP16 inference workload is dominated by the acoustic encoder, which accounts for 87% of the RTF (fig. 2(a)). Smaller beam widths would only improve the decoding portion. Conversely, the model with LM fusion shows a strong contribution of the decoding loop, with the large LM_{EXT} accounting for 56% of the total RTF.

The quantization of the RNN-T with LM fusion brings about remarkable speedup in both the encoder and the decoder. RTF drops from 0.221 at FP16, to 0.116 on a full INT8 model, and further down to 0.065 with full INT4 quantization, improving the RTF of the original RNN-T without extra LMs (fig. 2(a)). At 4 bits, we estimate an acceleration of $3.8\times$ for the encoder, $3.6\times$ for the decoder (inclusive of prediction and LMs), and $3.4\times$ for the E2E model. The non-acceleratable operations account for 12% of the total 4-bit workload. Our simulations also reveal sub-optimal coprocessor occupancy, indicating a CPU-coprocessor communication bottleneck which would be mitigated by a higher bandwidth (BW) connection. The last column of fig. 2(a) shows the RTF when modeling a BW increase, from 32 to 64 Gbps. The higher BW is highly beneficial, yielding a $6.8\times$ acceleration of the decoding loop, $5.8\times$ for the whole network and an RTF of 0.038.

Due to the BW bottleneck, the computation favors large beam widths. As shown in fig. 2(b), the decoding RTF decreases sub-linearly at small beams, primarily as a result of the increased reuse of cached hypotheses. More remarkable is the improvement in non-acceleratable operations runtime, which becomes negligible at beam = 1 (greedy decoding). The higher BW improves RTF across all beam widths, allowing for a more considerable reduction in the decoder contribution, down from 0.038 at beam 16 to 0.022 at beam 1. However, as discussed in section 4.2, lower beam widths reduce transcription accuracy.

4.2. Accuracy and compression

Table 1 shows the WER results at different stages of the quantization. At beam width 16, the fusion process improves the non-quantized model WER by 2.0% while almost doubling the model size. As discussed in section 4.1, LM fusion also increases inference runtime dramatically.

INT4 quantization of the RNN-T alone (following the scheme shown in fig. 1 but without LM_{EXT} and LM_{SRC}) shows a small WER degradation across datasets (0.6% on average for

Hub5 2000) with a compression factor of $7.2\times$ (vs. FP32) and a $3.4\times$ acceleration (vs. FP16). Although implementation of the Density Ratio LM Fusion approach shows a remarkable WER improvement on the INT4 quantized RNN-T, it is apparent that it also calls for quantization of LM_{EXT} and LM_{SRC} to retain the gains in model size and acceleration.

We initialize the INT8 and INT4 LM_{EXT} with pre-trained FP32 weights, which we found helpful to minimize WER degradation. The 4-bit LM_{EXT} quantized as per Fig. 1, achieves comparable WER to the FP32 counterpart across benchmarks. To further boost compression and acceleration, we use an equivalent strategy to quantize the LM_{SRC} at 4 bits, with no change in performance. With both LMs fully quantized at 4 bits, we obtain a 1.7% WER improvement on Hub5 2000 compared to the model without LM fusion, while benefiting from a close-to-ideal $7.6\times$ compression ratio of the full model and a $3.4\times$ acceleration.

In the beam width comparison, the INT4 model with LM fusion achieves at beam width 16 better WER than the full precision model without LMs at any beam widths, while providing a substantial reduction in model size and comparable RTF (as discussed in sec. 4.1).

More aggressive quantization of the RNN-T weights (INT2) shows increased but limited WER degradation in the model without LMs, $\leq 1\%$ across benchmarks compared to INT4 weights. On Hub5 2000, LM fusion has a similar impact on WER as in the 4-bit scenario ($\Delta = -1.8\%$). Here, we observe a trade off between WER, degraded from 10.7% to 11.4%, and compression ratio, improved from $7.6\times$ to $9.7\times$. The favorable impact of LM fusion on other benchmarks is less pronounced, showing that the error introduced by the 2-bit representation of the RNN-T weights can negatively affect the whole decoding process.

5. Conclusions

We presented a QAT approach at 4 bits which enables model compression and inference acceleration of RNN-T networks combined with external LM fusion methods with minimal WER degradation. We showed that an appropriate choice of quantizers that leverages the inherent boundaries of the LSTM activation functions and the connections between LSTM cells limits QAT model retraining time while maintaining excellent WER performance. The 4-bit quantized model with LM fusion improves the WER of the full precision model without LMs, while achieving comparable RTF (and vastly improved RTF compared to the full precision model with LMs) making it a viable option for real-time inference workloads.

6. References

- [1] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [2] Z. Tüske, G. Saon, K. Audhkhasi, and B. Kingsbury, "Single headed attention based sequence-to-sequence model for state-of-the-art results on switchboard-300," in *Proc. Interspeech*, 2020, pp. 551–555.
- [3] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [4] Y.-A. Chung, Y. Zhang, W. Han, C.-C. Chiu, J. Qin, R. Pang, and Y. Wu, "W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training," *arXiv preprint arXiv:2108.06209*, 2021.
- [5] G. Saon, Z. Tüske, D. Bolanos, and B. Kingsbury, "Advancing rnn transducer technology for speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5654–5658.
- [6] P. Guo, F. Boyer, X. Chang, T. Hayashi, Y. Higuchi, H. Inaguma, N. Kamo, C. Li, D. Garcia-Romero, J. Shi *et al.*, "Recent developments on esnet toolkit boosted by conformer," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5874–5878.
- [7] T. N. Sainath, Y. He, B. Li, A. Narayanan, R. Pang, A. Bruguier, S.-y. Chang, W. Li, R. Alvarez, Z. Chen *et al.*, "A streaming on-device end-to-end model surpassing server-side conventional model quality and latency," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6059–6063.
- [8] J. Li, R. Zhao, Z. Meng, Y. Liu, W. Wei, S. Parthasarathy, V. Mazalov, Z. Wang, L. He, S. Zhao, and Y. Gong, "Developing RNN-T Models Surpassing High-Performance Hybrid Models with Customization Capability," in *Proc. Interspeech 2020*, 2020, pp. 3590–3594.
- [9] E. McDermott, H. Sak, and E. Variari, "A density ratio approach to language model fusion in end-to-end automatic speech recognition," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 434–441.
- [10] E. Variari, D. Rybach, C. Allauzen, and M. Riley, "Hybrid autoregressive transducer (hat)," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6139–6143.
- [11] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang *et al.*, "Streaming end-to-end speech recognition for mobile devices," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6381–6385.
- [12] Y. Shangguan, J. Li, Q. Liang, R. Alvarez, and I. McGraw, "Optimizing speech recognition for the edge," *arXiv preprint arXiv:1909.12408*, 2019.
- [13] H. D. Nguyen, A. Alexandridis, and A. Mouchtaris, "Quantization aware training with absolute-cosine regularization for automatic speech recognition," in *Proc. Interspeech*, 2020, pp. 3366–3370.
- [14] S. Kim, A. Gholami, Z. Yao, N. Lee, P. Wang, A. Nrusimha, B. Zhai, T. Gao, M. W. Mahoney, and K. Keutzer, "Integer-only zero-shot quantization for efficient speech recognition," *arXiv preprint arXiv:2103.16827*, 2021.
- [15] A. Ardakani, Z. Ji, S. C. Smithson, B. H. Meyer, and W. J. Gross, "Learning recurrent binary/ternary weights," in *International Conference on Learning Representations*, 2019.
- [16] X. Liu, D. Cao, and K. Yu, "Binarized lstm language model," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, vol. 1, 2018, pp. 2113–2121.
- [17] K. Yu, R. Ma, K. Shi, and Q. Liu, "Neural network language model compression with product quantization and soft binarization," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2438–2449, 2020.
- [18] J. Xu, S. Hu, J. Yu, X. Liu, and H. Meng, "Mixed precision quantization of transformer language models for speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7383–7387.
- [19] A. Fasoli, C.-Y. Chen, M. Serrano, X. Sun, N. Wang, S. Venkataramani, G. Saon, X. Cui, B. Kingsbury, W. Zhang, Z. Tüske, and K. Gopalakrishnan, "4-Bit Quantization of LSTM-Based Speech Recognition Models," in *Proc. Interspeech 2021*, 2021, pp. 2586–2590.
- [20] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 3586–3589.
- [21] D. S. Park, W. Chan, Y. Zhang, C.-c. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proc. Interspeech*, 2019, pp. 2613–2617.
- [22] G. Saon, Z. Tüske, K. Audhkhasi, and B. Kingsbury, "Sequence noise injected training for end-to-end speech recognition," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6261–6265.
- [23] G. Saon, Z. Tüske, and K. Audhkhasi, "Alignment-length synchronous decoding for rnn transducer," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7804–7808.
- [24] H. Wu, P. Judd, X. Zhang, M. Isaev, and P. Micikevicius, "Integer quantization for deep learning inference: Principles and empirical evaluation," *arXiv preprint arXiv:2004.09602*, 2020.
- [25] J. Choi, S. Venkataramani, V. V. Srinivasan, K. Gopalakrishnan, Z. Wang, and P. Chuang, "Accurate and efficient 2-bit quantized neural networks," *Proceedings of Machine Learning and Systems*, vol. 1, pp. 348–359, 2019.
- [26] S. Venkataramani, J. Choi, V. Srinivasan, W. Wang, J. Zhang, M. Schaal, M. J. Serrano, K. Ishizaki, H. Inoue, E. Ogawa, M. Ohara, L. Chang, and K. Gopalakrishnan, "Deeptools: Compiler and execution runtime extensions for rapid ai accelerator," *IEEE Micro*, vol. 39, no. 5, pp. 102–111, 2019.
- [27] S. Venkataramani, V. Srinivasan, J. Choi, P. Heidelberger, L. Chang, and K. Gopalakrishnan, "Memory and interconnect optimizations for peta-scale deep learning systems," in *2019 IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC)*, 2019, pp. 225–234.
- [28] A. Agrawal, S. K. Lee, J. Silberman, M. Ziegler, M. Kang, S. Venkataramani, N. Cao, B. Fleischer, M. Guillorn, M. Cohen, S. Mueller, J. Oh, M. Lutz, J. Jung, S. Koswatta, C. Zhou, V. Zalani, J. Bonanno, R. Casatuta, C. Y. Chen, J. Choi, H. Haynie, A. Herbert, R. Jain, M. Kar, K. H. Kim, Y. Li, Z. Ren, S. Rider, M. Schaal, K. Schelm, M. Scheuermann, X. Sun, H. Tran, N. Wang, W. Wang, X. Zhang, V. Shah, B. Curran, V. Srinivasan, P. F. Lu, S. Shukla, L. Chang, and K. Gopalakrishnan, "9.1 a 7 nm 4-core ai chip with 25.6 tflops hybrid fp8 training, 102.4 tops int4 inference and workload-aware throttling," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64, 2021, pp. 144–146.