



A GPU-based WFST Decoder with Exact Lattice Generation

Zhehuai Chen^{1,3}, Justin Luitjens², Hainan Xu³, Yiming Wang³, Daniel Povey³, Sanjeev Khudanpur³

¹SpeechLab, Department of Computer Science and Engineering, Shanghai Jiao Tong University

²NVIDIA, USA

³Center for Language and Speech Processing, Johns Hopkins University

chenzhehuai@sjtu.edu.cn, jluitjens@nvidia.com, dpovey@gmail.com,
{hXu31, yiming.wang, khudanpur}@jhu.edu

Abstract

We describe initial work on an extension of the Kaldi toolkit that supports weighted finite-state transducer (WFST) decoding on Graphics Processing Units (GPUs). We implement token recombination as an atomic GPU operation in order to fully parallelize the Viterbi beam search, and propose a dynamic load balancing strategy for more efficient token passing scheduling among GPU threads. We also redesign the exact lattice generation and lattice pruning algorithms for better utilization of the GPUs. Experiments on the Switchboard corpus show that the proposed method achieves identical 1-best results and lattice quality in recognition and confidence measure tasks, while running 3 to 15 times faster than the single process Kaldi decoder. The above results are reported on different GPU architectures. Additionally we obtain a 46-fold speedup with sequence parallelism and multi-process service (MPS) in GPU. **Index Terms:** ASR, Decoder, Parallel Computing, WFST, GPU, Lattice Processing, Confidence Measure

1. Introduction

Recent advances in deep learning based automatic speech recognition (ASR) invoke growing demands in large scale speech transcription and analysis. The main computation of a common speech transcription system includes: acoustic model (AM) inference and weighted finite-state transducer (WFST) decoding. To reduce computation of the AM inference, researchers have proposed a variety of efficient forms of AMs, including novel structures [1, 2], quantization [3, 4], frame-skipping [5, 6] and end-to-end systems [7, 8]. Meanwhile, algorithmic improvement, e.g. pruning [9, 10] and lookahead [11, 12] is the common way to speed up decoding.

GPU-based parallel computing is another potential direction which utilizes a large number of units to parallelize the computation. As most of the computation takes the form of matrix operation, the AM inference can be sped up by sequence batching [13] similar to that during training [14]. Nevertheless, GPU-based decoding is less prevalent, despite its success in small language model (LM) [15]. Two reasons hamper its wide application: i) it is hard to utilize large LMs since GPUs have smaller memory. ii) the solution is not general, e.g. with specific AM or LM requirements, and lack of lattice generation.

This work is an extension to Kaldi [16], applying GPU parallel computing in WFST decoding. It is a general-purpose

offline decoder¹, which does not impose special requirements on the form of AM or LM, and works for all GPU architectures. To utilize large LMs in the 2nd pass and support rich post-processing, our design is to decode the WFSTs and generate exact lattices [17]. The work is open-sourced², and will be compatible with all released Kaldi recipes.

The rest of the paper is organized as follows. Prior work is compared in Section 2. Our main contributions are described in Sections 3: i) implement token recombination as an atomic GPU operation with no precision loss. ii) propose a load balancing strategy for token passing scheduling among GPU threads. iii) redesign parallel versions of exact lattice generation and pruning. Experiments are conducted on the Switchboard corpus in Section 4, followed by conclusion in Section 5.

2. Comparisons with Prior Work

[18] introduced a parallel Viterbi algorithm in CPU, which concentrates on partitioning the WFSTs to trade off between load balancing and synchronization overhead. To cope with similar issues in the context of GPU architectures, We give general LM-independent solutions in Section 3.3. [15, 19] first utilized GPUs in ASR decoding and showed significant speedup. However, the limited GPU memory becomes a bottleneck for large vocabulary tasks. A series of works were proposed in [20, 21, 23] and attempted to cope with large LM issue by an on-the-fly rescoring [10] algorithm. Advantages of our method include: i) general solution and better performance, by generating an exact lattice in the first pass decoding followed by a second pass rescoring, which always obtains slightly better results [22]. The resultant lattice also benefits other speech processing tasks, e.g. confidence measure. ii) faster decoding speed. This work benefits from less synchronization overheads, better load balancing, and sequence parallelism. Moreover, it also requires less data transfer compared with [23]. Finally, to the best of our understanding, [24] is the only open-source project in this field, but it only implemented the basic Viterbi decoding without combining AM posteriors and beam [9], which cannot be applied to ASR.

3. Parallel Viterbi Beam Search

The proposed system works in a 2-pass decoding framework [25] to tackle the language model size problem and enable rich post-processing based on exact lattices.

This work was partially supported by DARPA LORELEI Grant No HR0011-15-2-0024 and NSF Grant No CRI-1513128. The first author was supported by the National Key Research and Development Program of China under Grant No.2017YFB1002102, the China NSFC project (No. 61573241) and the Shanghai International Science and Technology Cooperation Fund (No. 16550720300).

¹Recent advances in CPU decoding make it easier to build a real-time decoder [2, 6] for online streaming applications. Thus our goal is to transcribe tons of offline audios.

²<https://github.com/chenzhehuai/kaldi/tree/gpu-decoder>

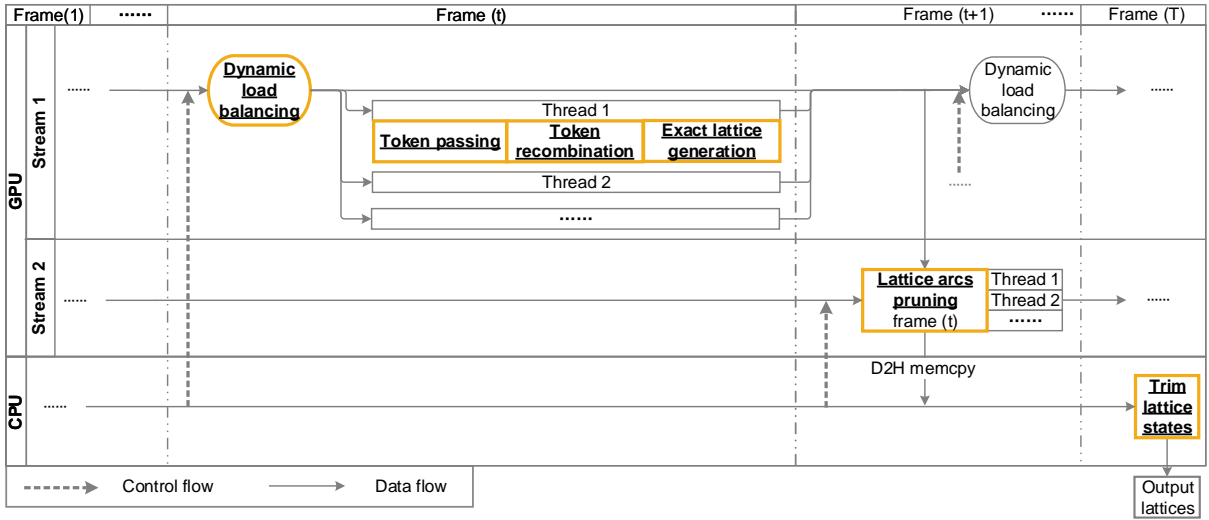


Figure 1: Framework of Parallel Viterbi Beam Search with Exact Lattice Processing

3.1. System Outline

Figure 1 shows our framework, with two GPU concurrent streams performing decoding and lattice-pruning in parallel launched by CPU asynchronous calls. Namely, at frame $(t + 1)$, stream 2 prunes the lattice generated by stream 1 at frame (t) . The procedure of decoding is similar to the CPU version [16] but works in parallel with specific designs discussed in the following sections. Load balancing controls the thread parallelism over both WFST states and arcs.

3.2. GPU-based Token Recombination

The synchronization overhead in Viterbi search comes from the token recombination. In ASR decoding, Viterbi search is reformulated to *token passing algorithm* [25], where a token represents a partial match up to frame t and each WFST state at frame t holds a single moveable token. At each frame, the Viterbi path is obtained by a *token recombination* procedure, where a *min* operation is performed on each state over all of its incoming arcs (e.g. state 7 in Figure 2 and the incoming arcs from state 2, 5 and 7), to compute the best cost and the corresponding predecessor of that state.

Algorithm 1 Thread-level Token Recombination (Inputs: accumulated cost, an out-going WFST arc and a current token)

```

1: procedure RECOMBINE(cost, arc, curTok)
2:   oldTokPack = state2tokPack[arc.next_state]
3:   curTokPack = packFunc(cost,arc.id) ▷ pack into uint64
4:   ret = atomicMin3(oldTokPack,curTokPack)
5:   if ret > curTokPack then ▷ recombine
6:     perArcTokBuf[arc.id] = *curTok ▷ store token

```

The original CPU algorithm conducts recombination in serial. We discuss how to implement GPU-based token recombination in this section and leave how to parallelize it to the next section. A naive implementation is by adding critical sections [27] to recombine tokens in serial. Such design is inefficient and causes unexpected deadlocks on pre-Volta GPUs. [15] performs a per-state reduction on the token passing

³atomicMin(*address, val) [26]. Computes $\min(*address, val)$, writes the result to *address*, and returns the original **address*.

results, which requires extra burden of bookkeeping memory to eliminate write conflicts. [20] encodes all token data in 32 bits and does the recombination by atomic operations. This method loses precision and makes the decoder dependent on models.

We propose Algorithm 1, which is a general method for GPU token recombination with no precision loss. The algorithm is performed per GPU thread, and runs in parallel for every arc, e.g. the arc from state 2 to state 7 in Figure 2 is processed by a thread following this algorithm (details on parallelism are in the following section). We pack the cost and the arc index into an uint64 to represent the token before recombination, with the former one in the higher bits for comparison purpose. In each frame, we save the token information in an array whose size is the number of arcs. This ensures there are no write conflicts between threads since each arc can be accessed at most once in each frame. After passing all tokens, we aggregate survived packed tokens, unpack them to get arc indexes, and store token information from the former array to token data structures exploiting thread parallelism.

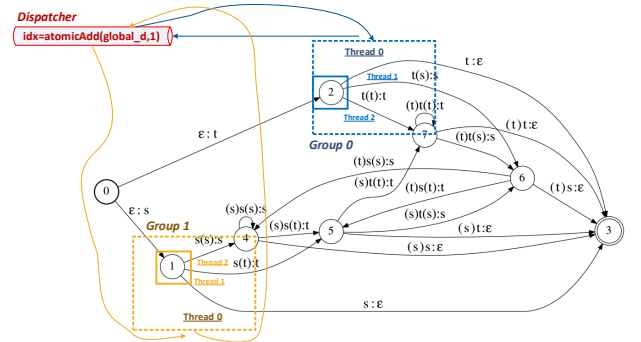


Figure 2: Example of Dynamic Load Balancing. The dashed box denotes a CUDA cooperative group and different groups are with different colors. Each group is controlled by Thread 0 of it. After processing all the forward links of a state, Thread 0 accesses the dispatcher and the next token is dynamically decided by atomic operation. Group 0 and 1 work in parallel.

3.3. Parallel Token Passing

Another issue in parallel algorithm is load balancing. For each state, we traverse their out-going arcs in parallel, till we reach a final state. Because WFST states might have different numbers of out-going arcs, the allocation of states and arcs to threads can result in load imbalance. Unlike in [15, 18] where the authors redesigned the WFST structure, this work does not require any structural change on AM, lexicon or LM. We propose *static load balancing*, which starts off by assigning roughly equal number of arcs to each thread, before processing the corresponding arcs. This design requires the accumulated sum of the arc numbers of active tokens⁴, which requires extra computation and kernel launch time.

Motivated by [28], we also introduce *dynamic load balancing* in Algorithm 2. We use a dispatcher in charge of global scheduling, and make N threads as a group ($N = 32$) to process all arcs from a single token. When the token is processed, the group requests from the dispatcher a new token. We implement task dispatching as an atomic operation. Figure 2 shows an example. We compare static and dynamic load balancing in Section 4.2.

Algorithm 2 Grid-level Token Passing ($N=32$; Inputs: the current active token vector)

```
1: procedure DYNAMIC LOAD BALANCING (toks)
2:   group = cooperative_groups::tiled_partition(32)
3:   if group.thread_rank()==0 then ▷ rank 0 in each group
4:     i = atomicAdd(global_d,1) ▷ allocate new tokens
5:     i = group.shfl(i,0) ▷ rank 0 broadcasts i to whole group
6:     if i >= sizeof(toks) then return ▷ all tokens processed
7:     for arc in tok2arcs(toks[i]) do ▷ thread parallelism
8:       call Recombine(toks[i].cost+arc.cost, arc, toks[i])
```

3.4. Exact Lattice Generation

An exact lattice [17] is one that stores precise costs and state level alignments, which is crucial to LM rescoring performance, and enables rich post-processing⁵. Implementing the lattice processing in GPU is non-trivial. [23] simply decodes in GPU and generates lattice in CPU, which significantly slows down the decoder and results in overheads of device-to-host (D2H) memory copy. We thoroughly solve this problem by redesigning parallel versions of lattice processing algorithms [17] in the following.

In lattice generation, for each token passing operation, a lattice arc needs to be stored in GPU memory. We design a GPU vector to store arcs using *push.back* function implemented by an atomic operation, where the memory is pre-allocated.

```
// implementation of v.push_back(val)
int idx = atomicAdd(cnt_d, 1); // idx=cnt_d++
mem_d[idx] = *val; // store data
```

To reduce the overhead of atomic operations, we pre-allocate K vectors and randomly select one to *push.back* the lattice arc into. After token passing, arcs from K vectors are aggregated. We process lattice nodes in a similar manner⁶.

⁴A GPU DeviceScan operation over around 10K integers.

⁵Beside the confidence measure [29] examined in this paper, it also benefits tasks including minimum Bayes-risk decoding [30], system combination [31], discriminative training [32], etc. .

⁶In our preliminary experiments, this optimization speeds up lattice arcs by 10 times with $K = 32$. The speedup of lattice nodes is small.

3.5. Lattice Pruning

The parallel lattice pruning is based on the algorithm described in [33, 17] with necessary modifications for GPU parallelization. The original CPU version of lattice pruning iteratively updates extra costs of nodes and arcs until they stop changing, where extra cost is defined as the difference between the best cost including the current arc and the best overall path. Arcs with high extra costs are then removed, along with nodes that are not associated with any remaining arcs. In GPUs, we i) parallelize the iterative updating of nodes and arcs over GPU threads; ii) use a global arc vector to replace the linked lists in the old implementation, for its lack of random access features to enable parallel access; iii) implement the extra cost updating as an atomic operation to eliminate write conflicts among threads.

When a lattice arc is pruned, we do not physically remove the arc, as memory allocation is expensive. Instead, we do a final merging step to aggregate all remaining arcs using thread parallelism and the GPU vector proposed in Section 3.4. Memory consumption of this implementation is discussed in Section 4.3. Meanwhile, we do not prune lattice nodes because: i) we need a static mapping for each arc to trace the previous and the next nodes before and after D2H memory copy. We use frame index t and vector index i to trace a node, thus node positions in the vector cannot be changed. ii) the lattice is constructed in CPU by iterating remaining arcs, thus nodes are implicitly pruned. iii) node D2H copy is done in each frame asynchronously, which does not introduce overheads.

4. Experiments

4.1. Setup

In Switchboard 300 hours corpus, we have two acoustic model baseline setups: one is a “TDNN-LSTM-C” model described in [2] with lattice-free MMI (LF-MMI) objective [34] with sub-sampled frame rate. To eliminate the effect of sub-sampling, the other is a stacked bidirectional LSTM network [35] with cross entropy (CE) objective with the original frame rate.

Evaluation is carried out on the NIST 2000 CTS test set. A 30k-vocabulary tri-gram LM trained from the transcription of Switchboard corpus and a Fisher interpolated 4-gram LM are used for decoding and lattice rescoring respectively. Kaldi 1-best decoder and lattice decoder are taken as CPU baselines⁷.

1-best results and lattice quality are both examined. For fair comparisons, lattice densities (*lat.den.*, measured by arcs/frame) [25] are kept identical for CPU and GPU decoding results. We report word error rate (WER), lattice rescored WER (*+rescored*), lattice oracle WER (OWER) [36] in recognition tasks. Normalized cross entropy (NCE) [37] is reported to show the quality of word level confidence measure after lattice rescoring. A decision tree is trained from the posterior histogram to map the posterior probabilities to confidence scores [38]. Real-time factor (RTF) is taken to evaluate the decoding speed and the speedup factor (Δ) versus corresponding CPU baselines. Because of the focus of this work, we report RTFs excluding AM inference in the tables. Overall RTFs are given when analyzing the tables. As the current goal is to build a fast speech transcription system, latencies of online streaming applications have not been taken into account. All experiments are conducted on *E5-2686 v4 @ 2.30GHz* with 1 socket (8 threads). A Tesla V100 is used in default and different architectures from Kepler to Volta are examined in Section 4.3.

⁷*decode-faster-mapped* and *latgen-faster-mapped* in Kaldi toolkit.

4.2. Performance and Speedup

Table 1 compares the proposed GPU lattice decoder with the CPU baseline under the same beam. All indicators are very close if not identical. The slight differences might be caused by different orders in which states are visited during decoding.

Table 1: *1-best and Lattice Performance (beam=14).*

system	lat. den.	WER	+rescored	OWER	NCE
CPU	30.3	15.5	14.3	11.2	0.322
GPU	30.2	15.5	14.3	11.2 ⁸	0.328

Decoding RTF speedups of both 1-best and lattice decoders are shown in Table 2. We obtain a 15-fold and 9.7-fold single-sequence speedups for 1-best and lattice decoders respectively. Most of the speedups come from parallel traversals of WFST states and arcs. The second most significant improvement is from lattice pruning. Beside the parallelism benefit similar to that of token passing, the GPU lattice pruning reduces the data amount of D2H memory copy and separate kernel launch costs. We obtain a 46-fold speedup if sequence parallelism is utilized with GPU MPS, which reduces context switching of multi-process [26]. For comparison, similar multi-process parallelism is applied in CPU but only results in a 1.8-fold speedup.

Although the speedup of atomic operation is not significant, it removes the critical section discussed in Section 3.2, which enables the use of older GPU architectures, shown in the next section. Meanwhile, dynamic load balancing is a comparable and general solution versus both static load balancing, shown in the 4th row, and the graph partition in [18].

To make a fair comparison of the overall RTFs, we take a separate GPU to do single-sequence inference for both baseline and the GPU decoder. The single-sequence overall RTFs of baseline and the proposed 1-best decoder are 0.18 and 0.03⁹. Deep integration of inference and decoding can be future topics.

Table 2: *Speedup of the Proposed Method (beam=14).*

system	search		+ lattice	
	RTF	Δ	RTF	Δ
CPU	0.16	1.0X	0.27	1.0X
+ 8-sequence (1 socket) ¹⁰	-	-	0.15	1.8X
GPU	0.016	10X	0.080	3.3X
+ atomic operation	0.015	11X	0.077	3.5X
+ dyn. load balancing	0.011	15X	0.075	3.6X
+ lattice prune	-	-	0.028	9.7X
+ 8-sequence (MPS)	0.0035	46X	0.0080	34X

4.3. Analysis

Figure 3 shows the proposed decoder works well in varieties of ASR systems and consistently improves the speed.

- GPU architectures.

The proposed system can be used in GPU architectures older than Kepler. It achieves a 3-fold speedup in K20 versus CPU.

- Acoustic model frame rates.

⁸OWER in this dataset is imprecise as we do not normalize the text in its scoring using *glm*. In LibriSpeech [34], we consistently observe that OWER is 23% of WER (3.4 v.s. 14.7) for both CPU and GPU.

⁹The AM inference can be further improved by sequence batching [13], which speeds up 10 times in our preliminary trials.

¹⁰The result is obtained from *latgen-faster-mapped-parallel* in Kaldi. 1-best decoder do not have such implementation.




Figure 3: *LM Size, Frame Rates and Architectures Comparison.*

We examine an original frame rate (10 ms) acoustic model, denoted in the dashed line. It works consistently better than the CPU baseline with reduced frame rate. It is worth noting that the 10 ms system is 3 times slower than its reduced frame rate counterpart, as the speech stream is processed in serial. Thus frame rate reduction [5, 6] is crucial even in GPU decoders.

- Language model sizes and memory consumption.

The Fisher interpolated 4-gram LM is pruned to different sizes and compiled to HCLG [9] for comparison (13MB, 62MB, 196MB, 258MB). The GPU decoder works consistently better than the CPU baseline. Meanwhile, our current implementation decodes at most 11GB WFST in a 12GB TITAN GPU. Details of memory optimization can be referred to our recent code revision. Although working in a 2-pass framework, the memory consumption can be further optimized in future. Moreover, CPU speed changes slightly with WFST sizes from 196MB to 258MB. Because of per-transition beam pruning, the number of tokens is not linear with WFST size. Our GPU decoder always passes tokens through arcs in parallel and does not have the phenomenon. Better scheduling strategies can be future work.

- Profiling.

Figure 4 compares the profiling result of the baseline CPU decoder¹¹ and our final system. The GPU decoder significantly speeds up both token passing and lattice processing, which take up most of the decoding time in CPU. Meanwhile, the “other” part in GPU includes kernel launch time and some synchronization pending. Optimizing them can be future topics.




Figure 4: *Profiling of the Decoders (best viewed in color).*

5. Conclusions

In this work, we describe initial work on an extension of the Kaldi toolkit that supports WFST decoding on GPUs. We design parallel versions of decoding and lattice processing algorithms. The proposed system significantly and consistently speeds up the CPU counterparts in a varieties of architectures.

The implementation of this work is open-sourced. Future works include the deep integration with sequence parallel acoustic inference in GPUs.

¹¹Lattice generation statistics in CPU is included in token passing.

6. References

- [1] J. Xue, J. Li, D. Yu, M. Seltzer, and Y. Gong, "Singular value decomposition based low-footprint speaker adaptation and personalization for deep neural network," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 6359–6363.
- [2] V. Peddinti, Y. Wang, D. Povey, and S. Khudanpur, "Low latency acoustic modeling using temporal convolution and lstms," *IEEE Signal Processing Letters*, vol. 25, no. 3, pp. 373–377, 2018.
- [3] I. McGraw, R. Prabhavalkar, R. Alvarez, M. G. Arenas, K. Rao, D. Rybach, O. Alsharif, H. Sak, A. Gruenstein, F. Beaufays et al., "Personalized speech recognition on mobile devices," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 5955–5959.
- [4] X. Xiang, Y. Qian, and K. Yu, "Binary deep neural networks for speech recognition," *Proc. Interspeech 2017*, pp. 533–537, 2017.
- [5] G. Pundak and T. N. Sainath, "Lower frame rate neural network acoustic models," in *Interspeech*, 2016, pp. 22–26.
- [6] Z. Chen, Y. Zhuang, Y. Qian, and K. Yu, "Phone synchronous speech recognition with ctc lattices," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 1, pp. 86–97, Jan 2017.
- [7] K. Audhkhasi, B. Ramabhadran, G. Saon, M. Picheny, and D. Nahamoo, "Direct acoustics-to-word models for english conversational speech recognition," *arXiv preprint arXiv:1703.07754*, 2017.
- [8] Z. Chen, Q. Liu, H. Li, and K. Yu, "On modular training of neural acoustics-to-word model for lvcstr," in *ICASSP*, April 2018.
- [9] M. Mohri, F. Pereira, and M. Riley, "Weighted finite-state transducers in speech recognition," *Computer Speech & Language*, vol. 16, no. 1, pp. 69–88, 2002.
- [10] T. Hori, C. Hori, and Y. Minami, "Fast on-the-fly composition for weighted finite-state transducers in 1.8 million-word vocabulary continuous speech recognition," in *Eighth International Conference on Spoken Language Processing*, 2004.
- [11] H. Soltau and G. Saon, "Dynamic network decoding revisited," in *Automatic Speech Recognition & Understanding, 2009. ASRU 2009. IEEE Workshop on*. IEEE, 2009, pp. 276–281.
- [12] D. Nolden, R. Schlüter, and H. Ney, "Search space pruning based on anticipated path recombination in lvcstr," in *Thirteenth Annual Conference of the International Speech Communication Association*, 2012.
- [13] P. R. Dixon, T. Oonishi, and S. Furui, "Harnessing graphics processors for the fast computation of acoustic likelihoods in speech recognition," *Computer Speech & Language*, vol. 23, no. 4, pp. 510–526, 2009.
- [14] K. Veselý, L. Burget, and F. Grézl, "Parallel training of neural networks for speech recognition," in *International Conference on Text, Speech and Dialogue*. Springer, 2010, pp. 439–446.
- [15] K. You, J. Chong, Y. Yi, E. Gonina, C. J. Hughes, Y.-K. Chen, W. Sung, and K. Keutzer, "Parallel scalability in speech recognition," *IEEE Signal Processing Magazine*, vol. 26, no. 6, 2009.
- [16] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz et al., "The kaldı speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. EPFL-CONF-192584. IEEE Signal Processing Society, 2011.
- [17] D. Povey, M. Hannemann, G. Boulianne, L. Burget, A. Ghoshal, M. Janda, M. Karafiát, S. Kombrink, P. Motlíček, Y. Qian et al., "Generating exact lattices in the wfst framework," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE, 2012, pp. 4213–4216.
- [18] C. Mendis, J. Droppo, S. Maleki, M. Musuvathi, T. Mytkowicz, and G. Zweig, "Parallelizing wfst speech decoders," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 5325–5329.
- [19] J. Chong, E. Gonina, K. You, and K. Keutzer, "Exploring recognition network representations for efficient speech inference on highly parallel platforms," in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [20] J. Kim, K. You, and W. Sung, "H-and c-level wfst-based large vocabulary continuous speech recognition on graphics processing units," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 1733–1736.
- [21] J. Kim, J. Chong, and I. Lane, "Efficient on-the-fly hypothesis rescoring in a hybrid gpu/cpu-based large vocabulary continuous speech recognition engine," in *Thirteenth Annual Conference of the International Speech Communication Association*, 2012.
- [22] H. Sak, M. Saraclar, and T. Güngör, "On-the-fly lattice rescoring for real-time automatic speech recognition," in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [23] J. Kim and I. Lane, "Accelerating large vocabulary continuous speech recognition on heterogeneous cpu-gpu platforms," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 3291–3295.
- [24] A. Argueta and D. Chiang, "Decoding with finite-state transducers on gpus," *arXiv preprint arXiv:1701.03038*, 2017.
- [25] S. Young, "The htk book version 3.4. 1," 2009.
- [26] "Cuda toolkit documentation," <http://docs.nvidia.com/cuda/>, accessed: 2018-03-17.
- [27] L. Lamport, "How to make a multiprocessor computer that correctly executes multiprocess program," *IEEE transactions on computers*, no. 9, pp. 690–691, 1979.
- [28] A. M. Alakeel, "A guide to dynamic load balancing in distributed computer systems," in *International Journal of Computer Science and Network Security (IJCSNS)*. Citeseer, 2010.
- [29] L. Mangu, E. Brill, and A. Stolcke, "Finding consensus among words: Lattice-based word error minimization," in *Sixth European Conference on Speech Communication and Technology*, 1999.
- [30] V. Goel and W. J. Byrne, "Minimum bayes-risk automatic speech recognition," *Computer Speech & Language*, vol. 14, no. 2, pp. 115–135, 2000.
- [31] J. G. Fiscus, "A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (rover)," in *Automatic Speech Recognition and Understanding, 1997. Proceedings., 1997 IEEE Workshop on*. IEEE, 1997, pp. 347–354.
- [32] D. Povey, "Discriminative training for large vocabulary speech recognition," Ph.D. dissertation, University of Cambridge, 2005.
- [33] A. Ljolje, F. Pereira, and M. Riley, "Efficient general lattice generation and rescoring," in *Sixth European Conference on Speech Communication and Technology*, 1999.
- [34] D. Povey, V. Peddinti, D. Galvez, P. Ghahramani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi," *Submitted to Interspeech*, 2016.
- [35] H. Sak, A. Senior, and F. Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [36] B. Hoffmeister, T. Klein, R. Schlüter, and H. Ney, "Frame based system combination and a comparison with weighted rover and cnc," in *INTERSPEECH*. Citeseer, 2006.
- [37] M. Siu and H. Gish, "Evaluation of word confidence for speech recognition systems," *Computer Speech & Language*, vol. 13, no. 4, pp. 299–319, 1999.
- [38] Z. Chen, Y. Zhuang, and K. Yu, "Confidence measures for ctc-based phone synchronous decoding," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4850–4854.