

NRC Publications Archive Archives des publications du CNRC

Towards effective and efficient non-autoregressive decoding using block-based attention mask

Wang, Tianzi; Xie, Xurong; Li, Zhaoqing; Hu, Shoukang; Jin, Zengrui; Deng, Jiajun; Cui, Mingyu; Hu, Shujie; Geng, Mengzhe; Li, Guinan; Meng, Helen; Liu, Xunying

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. / La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

For the publisher's version, please access the DOI link below. / Pour consulter la version de l'éditeur, utilisez le lien DOI ci-dessous.

Publisher's version / Version de l'éditeur:

<https://doi.org/10.21437/Interspeech.2024-404>

Interspeech 2024, pp. 262-266, 2024-09-01

NRC Publications Archive Record / Notice des Archives des publications du CNRC :

<https://nrc-publications.canada.ca/eng/view/object/?id=9dad387d-f86d-4f4b-9d50-063ed037bd68>

<https://publications-cnrc.canada.ca/fra/voir/objet/?id=9dad387d-f86d-4f4b-9d50-063ed037bd68>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.



Towards Effective and Efficient Non-autoregressive Decoding Using Block-based Attention Mask

Tianzi Wang^{1,2*}, Xurong Xie^{2†}, Zhaoqing Li¹, Shoukang Hu³, Zengrui Jing¹, Jiajun Deng¹, Mingyu Cui¹, Shujie Hu¹, Mengzhe Geng⁴, Guinan Li¹, Helen Meng¹, Xunying Liu¹

¹The Chinese University of Hong Kong, ²Institute of Software, Chinese Academy of Sciences
³Nanyang Technological University, ⁴National Research Council Canada

twang@se.cuhk.edu.hk, xurong@iscas.ac.cn, xyliu@se.cuhk.edu.hk

Abstract

This paper proposes a novel non-autoregressive (NAR) block-based Attention Mask Decoder (AMD) that flexibly balances performance-efficiency trade-offs for Conformer ASR systems. AMD performs parallel NAR inference within contiguous blocks of output labels that are concealed using attention masks, while conducting left-to-right AR prediction and history context amalgamation between blocks. A beam search algorithm is designed to leverage a dynamic fusion of CTC, AR Decoder, and AMD probabilities. Experiments on the LibriSpeech-100hr corpus suggest the tripartite Decoder incorporating the AMD module produces a maximum decoding speed-up ratio of 1.73x over the baseline CTC+AR decoding, while incurring no statistically significant word error rate (WER) increase on the test sets. When operating with the same decoding real time factors, statistically significant WER reductions of up to 0.7% and 0.3% absolute (5.3% and 6.1% relative) were obtained over the CTC+AR baseline.

Index Terms: Speech recognition, Non-autoregressive decoder, Autoregressive decoder, Beam search, Conformer

1. Introduction

State-of-the-art Transformer based automatic speech recognition (ASR) systems represented by, for example, Conformer Encoder-Decoder models with joint CTC and attention costs [1–5], are often based on an autoregressive (AR) Transformer Decoder architecture. The resulting AR model inference follows a monotonic and left-to-right decoding strategy, which predicts one output token at a time. However, this limits the potential for model inference parallelization during beam search decoding for practical application scenarios that are not only performance-critical but also efficiency-sensitive.

Several methods have been studied to optimize the inference speed of ASR models with Encoder-Decoder architecture [6–22]. One general solution is to adopt non-autoregressive (NAR) Transformers based Decoder designs. NAR Transformers provide more powerful parallelization than AR architectures to improve ASR decoding efficiency, and have gained considerable attention in a series of prior researches in recent years [11–14, 18]. These include, but not limited to, the following categories: a) mask-based NAR [14–17] learn to fill the randomly masked training label tokens conditioned on the unmasked ones; b) alignment-refinement based approaches [18, 19] that aim to refine the initial CTC produced alignment by injecting noisy labels extracted from, for example, auto-encoders

in the Align-Refine method [18], or noise perturbed Encoder alignment posteriors in the Align-Denoise approach [19]; c) integrate-and-fire based approaches [13] which learn to align speech and labels using continuous integrate-and-fire modules; and d) hybrid AR+NAR Decoders that exploit their complementarity in combination [20–22].

Efforts on developing high-performance and low-latency NAR based ASR models require several key challenges to be addressed. First, NAR models' intrinsic lack of monotonic sequence modelling constraints leads to their large performance gap against state-of-the-art ASR systems based on AR designs. Despite recent attempts to mitigate such modelling deficiency of NAR ASR systems using, for example, Mask-CTC and its improved variants [14, 15], or iterative alignment-refinement approaches [18, 19], their performance gap against AR counterparts still exists. Second, there is a notable lack of effective and efficient one-pass beam search decoding algorithm that are purpose-designed for NAR Decoders and their combination with CTC and AR ones. Prior researches in this direction either deployed standalone NAR Decoders in a later rescoring pass within a multi-pass decoding framework [20], while the first pass recurrent neural network transducer (RNN-T) decoding serves to produce an initial set of hypotheses, or vice versa when the NAR Decoders are used in the initial N-best generation before AR rescoring [22].

To this end, this paper proposes a novel non-autoregressive block-based attention mask decoder (AMD) that flexibly balances performance-efficiency trade-offs for Conformer ASR systems. The AMD decoder leverages both parallel NAR inference and monotonic left-to-right AR prediction. Parallel NAR inference is performed within contiguous blocks of output labels that are concealed using attention masks, while monotonic and left-to-right AR prediction and history context amalgamation are conducted between blocks. The AMD based NAR decoder is jointly trained with the standard CTC module and attention-based AR Transformer. A beam search algorithm is designed to leverage a dynamic fusion of CTC, AR Decoder and AMD probabilities. In addition to fixed size attention-masking blocks during NAR inference, mixed size blocks were also explored to facilitate cold start monotonic inference (block size = 1) for the initial N labels of each speech segment, before switching to parallel label prediction (block size > 1) for the remaining labels.

Experiments were conducted on the benchmark LibriSpeech-100 dataset using smaller Conformer ASR models trained on the 100-hr data only, or larger ones [23] further incorporating WavLM SSL features and 960-hr data based pre-training. The tripartite Decoder incorporating the additional AMD module produces a maximum decoding speed up ratio of 1.73x over the baseline CTC+AR decoding, while

* Work done while at Institute of Software, Chinese Academy of Sciences during the internship

† Corresponding author




Figure 1: The proposed Conformer ASR system architecture using a tripartite Decoder that includes the proposed non-autoregressive block-based attention mask decoder (AMD) (left, dashed yellow line) in addition to the CTC module (centre, purple) and attention-based AR Decoder (right, dashed blue line).

incurring no statistically significant word error rate (WER) increase on the test sets on average. When operating with the same decoding real time factors, statistically significant WER reductions of up to 0.7% and 0.3% absolute (5.2% and 6.1% relative) were obtained over the CTC+AR baseline.

The main contributions of this paper are as follows:

- This paper presents a novel block-based attention mask decoder (AMD) that flexibly balances performance-efficiency trade-offs for Conformer ASR systems in recognition time. For the first time, this NAR decoder allows: 1) decoding time speed up via non-autoregressive, parallel inference without increasing ASR WERs; and 2) statistically significant WER reductions over AR decoders when operating with the same decoding real time factors. In contrast, large performance degradation were often observed in prior researches when alternative forms of NAR Decoders [14, 15, 19], e.g. Mask-CTC, were used, or failed to provide the WERs at equivalent real-time factors for fair comparison [13].
- A novel beam search algorithm is designed to leverage a dynamic fusion of CTC, AR Decoder and AMD probabilities. In contrast, prior researches have largely deployed NAR Decoders in a more time-consuming multi-pass decoding framework [20, 22, 24–26].

2. Hybrid Attention Encoder-Decoder Based ASR System

A hybrid CTC-attention Encoder-Decoder (AED) based ASR system considered in this paper consists of three parts [1]: a shared Conformer-based Encoder, a CTC Decoder and an attention-based AR Decoder. An interpolated CTC and AR cost function is used for model training:

$$\mathcal{L}_{c-a} = \gamma_1 \mathcal{L}_{CTC} + \gamma_2 \mathcal{L}_{AR} \quad (1)$$

where γ_1 and γ_2 are the weights for CTC and AR Decoder during training respectively. The decoding stage follows an label synchronous autoregressive beam search [1] to find the hypothesis by taking into account both the CTC and AR Decoder probabilities. At the i -th step of the decoding process, the score of the partially decoded hypothesis $\mathbf{h}_{\leq i}$ is given by:

$$\alpha_{c-a}(\mathbf{h}_{\leq i}) = \lambda_1 \alpha_{CTC}(\mathbf{h}_{\leq i}) + \lambda_2 \alpha_{AR}(\mathbf{h}_{\leq i}) \quad (2)$$

where λ_1 and λ_2 are the weights for CTC and AR Decoder during decoding respectively. The CTC score α_{CTC} is expressed as the cumulative probability of all token sequences sharing $\mathbf{h}_{\leq i}$ as the common history context given by

$$\alpha_{CTC}(\mathbf{h}_{\leq i}) = \log P_{CTC}(\mathbf{h}_{\leq i}, \dots | \mathcal{X}) \quad (3)$$

where \mathcal{X} denotes the Encoder outputs. Meanwhile, the AR Decoder probability of the output token y_i is conditioned on $\mathbf{h}_{< i}$ and \mathcal{X} as $P_{AR}(y_i | \mathbf{h}_{< i}, \mathcal{X})$. The AR Decoder score $\alpha_{AR}(\mathbf{h}_{\leq i})$ is given by:




Figure 2: Inference using: a) an AR Decoder; b) a Mask Prediction Decoder; c) the proposed AMD. ‘msk’ refers to input mask tokens, **M** is the attention mask within contiguous blocks for parallel inference via AMD, ‘B’ refers to the block size. Tokens highlighted in green background denote Decoder inputs at the current inference step. Tokens highlighted in purple denote predicted tokens at current step. Tokens in dashed boxes in a) and c) represent those from the previous inference step. Tokens with “-” denote those obtained from CTC prediction.

$$\alpha_{AR}(\mathbf{h}_{\leq i}) = \log \prod_{j=1}^i P_{AR}(y_j | \mathbf{h}_{< j}, \mathcal{X}) \quad (4)$$

3. Block-Based Attention Mask Decoder

The proposed tripartite Decoder integrates a novel non-autoregressive block-based attention mask decoder (AMD) (Fig. 1, left, dashed yellow line) alongside the standard CTC module (Fig. 1, center, purple) and attention-based AR Decoder (Fig. 1, right, dashed blue line).

3.1. Block-Based Attention-Masking Decoder

Fig. 2 illustrates example inference using either: a) a AR Decoder; b) a Mask Prediction Decoder; or c) the proposed AMD. The AR decoder in Fig. 2 assigns the probability of each output token in a strictly autoregressive left-to-right and monotonic manner. The NAR Mask-CTC Decoder [14] performs parallel inference over multiple input tokens that are concealed using special “msk” symbols, e.g. y_2 and y_4 , while the sequential constraint between the masked input tokens ignored. It is also possible to restrict Mask-CTC prediction to be applied only to a smaller subset of the output tokens that are carefully selected using CTC scores during recognition time, in order to mitigate the performance loss when used to predict all the tokens [14] (See contrast between Sys. 2 and 3, Tab. 1).

In contrast, the proposed AMD model leverages both parallel NAR inference and AR prediction as illustrated in Fig. 2(c). Parallel NAR inference is performed within each contiguous block of output tokens concealed via block-based attention masks **M**, e.g. y_2 and y_3 , while sequential AR prediction and history context accumulation are made left-to-right between the blocks. To prevent information leakage in AMD, the same input token position being attention-masked also requires its respective token embedding layer outputs to be set as zero. In this paper, AMD is used to predict all tokens without any subset filtering as in Mask-CTC prediction.

Given an input token sequence \mathbf{h} and and the Encoder outputs \mathcal{X} , the AMD probability $P_{AMD}(y_j | \mathbf{h}, \mathcal{X})$ of the j -th token which are in an attention-mask block, j -th $\in [i, i + B - 1]$,

$$P_{AMD}(y_j | \mathbf{h}, \mathcal{X}) = P_{AMD}(y_j | \mathbf{h}_{< i}, \mathbf{h}_{> i+B-1}, \mathcal{X}) \quad (5)$$

where B is the size of the attention-mask block.

Ground truth text labels are used as the input during AMD training. In order to improve the AMD model’s generalization to all the tokens in the training data, and also different inference block sizes (including using mixed sized blocks of Sec. 3.2), diversity is injected into the training process by repeatedly performing inference over all the tokens of each sentence us-

Algorithm 1 NAR Beam Search with CTC-AR-AMD Tripartite Decoder

L_{\max}	Maximum allowed hypothesis length size (e.g. 512)
B	AMD NAR parallel inference block size (e.g. 2)
K_{main}	main top-K hypotheses beam for tripartite decoder (e.g. 10)
K_{AMD}	local top-K hypotheses beam in AMD search (e.g. 12)
$\mathcal{H}_{\text{main}}^{\text{main}}$	Top K_{main} hypotheses by CTC-AR-AMD tripartite decoder
$\mathcal{H}_{\text{AMD}}^{\text{AMD}}$	Top K_{AMD} hypotheses by CTC-AMD decoding up to slot j
\mathbf{h}^{CTC}	CTC greedy search 1-best hypothesis

```

1:  $\text{hyps} = \{\text{sos} : 1.0\}$ ;  $\text{curr.hyps} = \{\text{sos} : 1.0\}$ ;  $\mathbf{h}^{\text{CTC}}$ 
2: # each fixed size attention-mask block of  $B$  label slots
3: for  $i = 1 \dots L_{\max}$ ; step =  $B$ , do
4:   for  $\mathbf{h}_{\leq i} \in \mathcal{H}_{\text{main}}^{\text{main}}$  do in parallel,
5:     # each of  $B$  in-block label slots
6:     for  $j = i, \dots, i+B-1$  do in parallel,
7:       # Top  $K_{\text{AMD}}$  labels predicted for each of  $B$  in-block slots
8:        $\mathcal{Y}_j = \{y_j | y_j \in \text{Topk}(P_{\text{AMD}}(y_j | \mathbf{h}_{\leq i}, \mathbf{h}_{>i+B-1}, \mathcal{X}), K_{\text{AMD}})\}$ 
9:       # Union with  $i^{\text{th}}$  label of CTC hypothesis  $\mathbf{h}^{\text{CTC}}$ 
10:       $\mathcal{Y}_j = \mathcal{Y}_j \cup \mathbf{h}_j^{\text{CTC}}$  AMD in-block Decoding
11:      for  $j = i, \dots, i+B-1$  do, # each of  $B$  in-block label slots
12:        # Connecting (o) previous paths  $\mathcal{H}_{j-1}^{\text{AMD}}$  up to label slot
13:         $j-1$  located immediately before slot  $j$  with the
14:         $K_{\text{AMD}}+1$  in-block label predictions  $\mathcal{Y}_j$  at position  $j$ 
15:         $\mathcal{H}_j = \{\mathbf{h}_{\leq j-1} \circ y_i | \mathbf{h}_{\leq j-1} \in \mathcal{H}_{j-1}^{\text{AMD}} \text{ and } y_i \in \mathcal{Y}_j\}$ 
16:        # CTC+AMD scores for expanded paths up to slot  $j$ 
17:        for  $\mathbf{h}_{\leq j} \in \mathcal{H}_j$  do in parallel,
18:           $\alpha_{\text{c-m}}(\mathbf{h}_{\leq j}) = \lambda_1 \alpha_{\text{CTC}}(\mathbf{h}_{\leq j}) + \lambda_2 \alpha_{\text{AMD}}(\mathbf{h}_{\leq j})$ 
19:          # Pruning to top  $K_{\text{AMD}}$  partial hypotheses up to slot  $j$ 
20:           $\mathcal{H}_j^{\text{AMD}} = \text{Topk}(\alpha_{\text{c-m}}(\mathbf{h}_{\leq j}), K_{\text{AMD}})$ 
21:           $\mathbf{h}_{\leq j} \in \mathcal{H}_j$  CTC+AMD in-block Decoding
22:        # Tripartite decoder hypotheses rescoring by adding AR scores
23:        for  $\mathbf{h}_{\leq i+B-1} \in \mathcal{H}_{i+B-1}^{\text{AMD}}$  do in parallel,
24:           $\alpha_{\text{c-m-a}}(\mathbf{h}_{\leq i+B-1}) = \alpha_{\text{c-m}}(\mathbf{h}_{\leq i+B-1}) + \lambda_3 \alpha_{\text{AR}}(\mathbf{h}_{\leq i+B-1})$ 
25:          # Pruning to top  $K_{\text{main}}$  hypotheses leading up to last in-block slot
26:           $\mathcal{H}_{\text{main}}^{\text{main}} = \text{Topk}(\alpha_{\text{c-m-a}}(\mathbf{h}_{\leq i+B-1}), K_{\text{main}})$ 
27:           $\mathbf{h}_{\leq i+B-1} \in \mathcal{H}_{i+B-1}^{\text{AMD}}$  CTC+AR+AMD Reranking
28: return  $\mathcal{H}_{\text{main}}^{\text{main}}$ 

```

ing attention-mask blocks of 4 different sizes that are randomly sampled within the range of $[1, L]$, where L is the length of each sentence. The AMD training loss is defined as the probabilities accumulated over all tokens in a sentence using a block size B_n , before being marginalized over different block sizes:

$$\mathcal{L}_{\text{AMD}} = - \sum_{n=1}^4 \log \prod_{j=1}^L P_{\text{AMD}}(y_j | \mathbf{h}_{\leq i}, \mathbf{h}_{>i+B_n-1}, \mathcal{X}) \quad (6)$$

The 3-way interpolated \mathcal{L}_{CTC} , \mathcal{L}_{AR} , and \mathcal{L}_{AMD} losses was used to form a tripartite loss:

$$\mathcal{L}_{\text{c-m-a}} = \gamma_1 \mathcal{L}_{\text{CTC}} + \gamma_2 \mathcal{L}_{\text{AR}} + \gamma_3 \mathcal{L}_{\text{AMD}} \quad (7)$$

where γ_1 , γ_2 and γ_3 are the weights for CTC, AR Decoder and AMD during training respectively. This combined loss is used to train the Conformer ASR model featuring the additional AMD module in Fig. 1.

3.2. Beam Search Using AMD

AMD Decoding differs from the above training procedure mainly by replacing the ground truth input labels to be block attention-masked in training using the label hypotheses obtained from CTC greedy search. The AMD probabilities are dynamically fused with the CTC and AR Decoder scores in a novel beam search algorithm (Algorithm 1) tailored designed for AMD in this paper.

In addition to the above mentioned NAR style parallel inference within each block (line 5-8, Algorithm 1), there are several further important details of this algorithm to note:

(1) Left-to-right sequential AR prediction requires commensurate history context accumulation between the moving blocks. This is achieved by connecting the existing partial sequence hypotheses leading up to a $(j-1)$ -th label slot, immediately before the current in-block j -th label being predicted, with the top-K output token labels for the j -th label (line 11-15).

(2) To improve hypotheses coverage, the CTC greedy search 1-best hypothesis is progressively merged with the AMD decoding outputs (line 9-10).

(3) To ensure efficiency during AMD search, top-K pruning is initially applied locally to the output tokens for each in-block slot using AMD scores alone (line 8, via K_{AMD}). It is further applied to the iteratively expanded partial sequence hypotheses using interpolated CTC+AMD scores with λ_1 and λ_2 as the weights of CTC and AMD respectively (line 16-18, via K_{AMD})¹, and finally to the tripartite CTC+AR+AMD decoder scores re-ranked hypotheses with λ_3 as the weight for AR Decoder (line 21-23, via K_{main})². The CTC score can be calculated efficiently following [1], and the AMD score for each partial hypothesis sequence $\mathbf{h}_{\leq i}$, $\alpha_{\text{AMD}}(\mathbf{h}_{\leq i})$ is

$$\alpha_{\text{AMD}}(\mathbf{h}_{\leq i}) = \log \prod_{j=1}^i P_{\text{AMD}}(y_j | \cdot) \quad (8)$$

where the AMD probability $P_{\text{AMD}}(y_j | \cdot)$ is defined in Eqn. (5).

To achieve more flexible balance between performance and efficiency, in addition to using fixed size attention-masking blocks during NAR inference, **mixed size blocks** were also explored to allow cold start monotonic inference (block size $B=1$) for the initial N labels of each speech segment to be more slowly built-up, before switching to faster parallel label prediction ($B>1$) for the remaining labels in the same sentence.

4. Experiments

Experiment Setup and Baseline Systems: Experiments were conducted using (1) smaller Conformer end-to-end models³ with 80-dim log-mel filter-banks and 3-dim pitch as input, and (2) larger Conformer models⁴ further incorporating WavLM [27] SSL features and 960-hr data based pre-training. All models are of 2 Convolution blocks followed by 12 Conformer blocks for encoder, and 6 Transformer decoder blocks. The decoder output vocabulary comprised 5,000 byte pair encoding (BPE) tokens derived from the transcripts of respective training sets. For all models with a CTC+AR Decoder, an interpolated CTC+AED (3:7) loss function was used for training. Speed perturbation [28] and SpecAugment [29] were also applied. Real time factors (RTFs) of the smaller Conformer models were measured on an Intel Xeon 5317, while RTFs of the larger ones were measured on an NVIDIA A40. Matched pairs sentence-segment word error (MAPSSWE) [30] based statistical significance test was performed at a significance level $\alpha = 0.05$.

Experiments Using Greedy Search: Tab. 1 presents the performance of AMD tripartite and baseline CTC+AR Decoders using greedy search. Following trends can be observed: (a) Bipartite CTC+AMD (Sys. 4-5) and tripartite CTC+AR+AMD (Sys. 6-12) Decoders achieve lower WER but slower speeds compared to Mask-CTC with an MLM decoder (Sys. 2-3), which refills tokens partially. (b) The CTC+AMD Decoder achieves 1-best WER and RTF similar to the CTC+AR baseline (Sys. 4 vs. 1). (c) The proposed tripartite Decoder achieves a maximum speed-up of 1.73x over the baseline CTC+AR Decoder, with no significant WER changes on the averaged WERs across “clean” and “other” test sets using fixed-size decoding (Sys.9 vs. 1). Mixed-size decoding allows more accurate AR in-

¹CTC and AMD scores weights empirically set as 0.3 and 0.3

²CTC-AMD and AR scores weights empirically set as 0.6 and 0.4

³feedforward dim = 1024, # attention head = 4, attention head dim = 256, convolution kernel size = 15.

⁴feedforward dim = 2048, # attention head = 8, attention head dim = 512, convolution kernel size = 15.

Table 1: Performance of (a) greedy search baseline with CTC + autoregressive (AR) Decoder (Sys. 1, 13); (b) Mask-CTC Decoder [14] implemented in the form of CTC incorporated with masked language model (MLM), with approx. 10% (Sys. 2, selected acc. CTC scores) and all (Sys. 3) of the tokens refilled by MLM Decoder; (c) bipartite CTC + AMD based Decoder (Sys. 5-6); (d) the proposed tripartite Decoder with fixed decoding block size (Sys. 4-9, 14-16), and with variable size decoding (Sys. 10-12). “N-B” denotes first N tokens are decoded in an AR manner, followed by non-AR decoding of remaining tokens with block size = B. “1-best” stands for the 1-best WER. † and ‡ denote no statistically significant WER difference over Sys. 1 and 13 respectively on the average (Ave.) WER, ◊ and ▷ denote RTF similar to Sys. 1 and 13 respectively.

Sys.	Encoder	Decoder	Weights	Block Size	1-best			RTF
					clean	other	Ave.	
1	Conformer (100-hr)	CTC+AR	0.7:0.3	1	7.8	18.9	13.3	0.26
2		CTC+MLM [14] (10%)	-	-	7.5	20.6	14.0	0.07
3		CTC+MLM [14] (all)	-	-	10.7	21.1	15.9	0.10
4		CTC+AMD	0.5:0.5	1	7.6	18.7	13.2	0.26
5		CTC+AMD	0.5:0.5	2	8.9	20.9	13.4	0.17
6		CTC+AR+AMD	0.4:0.3:0.3	1	6.7	17.9	12.3	0.41
7				2	7.1	18.8	12.9	0.31
8				4	7.3	19.1	13.2 [†]	0.17
9				8	7.5	19.4	13.4 [†]	0.15
10				10-2	7.2	18.6	12.9	0.27 [◊]
11				10-4	7.2	18.7	12.9	0.26 [◊]
12		30-8	7.0	18.3	12.6	0.27 [◊]		
13	WavLM feat.	CTC+AR*	0.7:0.3	1	3.8	6.1	4.9	0.069
14	+ Conformer	CTC+AR+AMD*	0.4:0.3:0.3	1	3.2	5.5	4.3	0.102
15	(960-hr pt.+			4	3.4	5.8	4.6	0.064 [▷]
16	100-hr ft.) [23]			16	3.7	6.1	4.9 [‡]	0.048

* The CTC and AR Decoder parameters were initialized with the pre-trained model [23] and were frozen during fine-tuning. The AMD Decoder parameters were initialized with [23] and fine-tuned to LibriSpeech-100hr.

ference to be used for the first N tokens of each sentence, before switching to faster NAR parallel prediction. By doing so performance is weighed more against efficiency. Notably, mixed-size decoding systems with RTFs similar to the CTC+AR baseline (Sys. 10-12) achieve statistically significant WER reductions of up to 0.7% absolute (5.3% relative) (Sys. 12 vs. 1). (d) Similar trends are observed on larger models incorporating WavLM SSL features and 960-hr pre-training. The tripartite Decoder with the AMD module (Sys. 14-16) achieves a maximum speed-up of 1.44x over the CTC+AR baseline (Sys. 16 vs. 13) without a significant WER increase. At the same decoding speed, WER reductions of up to 0.3% absolute (6.1% relative) are obtained over the CTC+AR baseline (Sys. 15 vs. 13).

Experiments Using Beam Search: Performance of AMD Decoders using beam search are in Tab. 2. The results suggest the following trends, (a) the tripartite Decoder consistently speeds up decoding without significant WER increase. For LibriSpeech-100hr trained models, it achieves a maximum speed-up of 1.59x (Sys. 6 vs. 1). Models with WavLM SSL feat. and 960-hour pretraining see a 1.53x speed-up (Sys. 11 vs. 9). (b) Mixed-size decoding allows the tripartite Decoder to achieve similar or even slightly better performance than the CTC+AR baseline at the same RTF (Sys. 8 vs. 1, Sys. 12 vs. 9). (c) For systems that exhibit similar 1-best WER (Sys. 6-8 vs. 1, Sys. 11-12 vs. 9), the tripartite Decoder consistently larger Oracle WERs than the CTC+AR baselines. Further analysis of this phenomenon is conducted in the next section.

Analysis of AMD Lattice Density and Oracle WER: Fig. 3 further explores the relationship between (a) decoding lattice density and (b) the average Oracle WER obtained from the 100-best hypotheses on “clean” and “other” test sets. Lattice density here measures the average number of unique predictions for

Table 2: Performance of (a) baseline CTC+AR beam search (Sys. 1, 9), (b) the proposed tripartite Decoder with fixed decoding block (Sys. 2-5, 10-11), and (c) with variable size decoding (Sys. 6-8, 12). Beam size for all systems was set to 10. “Oracle” denotes the Oracle WER obtained from 100-best hypotheses. † and ‡ denote no significant statistically significant WER difference over Sys. 1 and 9 respectively on the average (Ave.) WER, ◊ and ▷ denote RTF similar to Sys. 1 and 9 respectively. Other naming conventions follow Tab. 1.

Sys.	Encoder	Decoder	Weights	Block Size	1-best			Oracle		RTF	
					clean	other	Ave.	clean	other		
1	Conformer (100-hr)	CTC+AR	0.7:0.3	1	6.8	17.5	12.1	4.3	13.6	1.19	
2					6.5	17.4	11.9	4.4	13.7	1.33	
3					6.7	17.9	12.3	4.9	15.9	0.74	
4					6.9	18.3	12.6	5.2	15.6	0.45	
5					7.1	18.6	12.8	5.4	15.8	0.31	
6					5-2	6.7	17.8	12.2 [†]	4.8	14.7	0.75
7					10-2	6.6	17.7	12.1 [†]	4.7	14.6	0.84
8					50-2	6.5	17.5	12.0 [†]	4.4	13.9	1.21 [◊]
9	WavLM feat.	CTC+AR	0.7:0.3	1	3.1	5.4	4.2	2.1	3.8	0.124	
10	+ Conformer (960-hr pt.+	CTC+AR	0.4:0.3:0.3	1	3.0	5.3	4.1 [‡]	2.1	3.7	0.178	
11	100-hr ft.) [23]				8	3.3	5.5	4.3 [‡]	2.7	4.5	0.081
12	10-2				3.0	5.4	4.1 [‡]	2.3	4.1	0.119 [◊]	




Figure 3: AMD lattice density and oracle WERs computed using 100-best hypotheses over varying top-K beam size from 1 to 20.

each token in the ground truth text, as found in the 100-best list, while Oracle WER evaluates beam search performance by identifying the single hypothesis with the lowest WER within the 100-best list. As block size increases, the tripartite Decoder exhibits sparser lattices (lower density) and faster saturation compared to the baseline system (orange, green, and red curves vs. dashed blue curve). This explains why the improvements from the CTC+AR+AMD tripartite Decoder over the CTC+AR baseline when operating with greedy search under aggressive pruning of Tab. 1 are much larger than those found in beam search experiments of Tab. 2. This disparity on performance improvements may be due to in the premature pruning within attention-mask blocks of AMD, as implemented in line 21, Algorithm 1. We plan to investigate this in future work.

5. Conclusions

This paper proposes a non-autoregressive block-based Attention Mask Decoder (AMD) that flexibly balance performance-efficiency trade-offs. Experimental results on the LibriSpeech-100hr dataset demonstrate that the proposed tripartite Decoder incorporating the AMD module achieves a maximum decoding speed-up ratio of 1.73x over the baseline CTC-AR decoding, while while incurring no statistically significant WER increase on the “clean” and “other” test sets. When operating with the same decoding real time factors, statistically significant WER reductions of up to 0.7% and 0.3% absolute (5.3% and 6.1% relative) were obtained over the CTC-AR baseline. Future research will focus on further refining the AMD beam search algorithm and enhancing the quality of lattice or N-best outputs.

6. Acknowledgements

This research is supported by Hong Kong RGC GRF grant No. 14200021, 14200220, TRS grant No. T45-407/19N, Innovation & Technology Fund grant No. ITS/218/21, National Natural Science Foundation of China 62106255, and Youth Innovation Promotion Association CAS grant 2023119.

7. References

- [1] S. Watanabe, T. Hori, S. Kim *et al.*, “Hybrid ctc/attention architecture for end-to-end speech recognition,” *IEEE J. Sel. Top. Signal Process.*, 2017.
- [2] P. Guo, F. Boyer, X. Chang *et al.*, “Recent developments on espnet toolkit boosted by conformer,” in *ICASSP*, 2021.
- [3] S. Kim, T. Hori, and S. Watanabe, “Joint ctc-attention based end-to-end speech recognition using multi-task learning,” in *ICASSP*, 2017.
- [4] T. N. Sainath, R. Pang, D. Rybach *et al.*, “Two-pass end-to-end speech recognition,” *INTERSPEECH*, 2019.
- [5] Z. Yao, D. Wu, X. Wang *et al.*, “Wenet: Production oriented streaming and non-streaming end-to-end speech recognition toolkit,” *INTERSPEECH*, 2021.
- [6] S. Gandhi, P. von Platen, and A. M. Rush, “Distil-whisper: Robust knowledge distillation via large-scale pseudo labelling,” *arXiv preprint arXiv:2311.00430*, 2023.
- [7] Z. Li, T. Wang, J. Deng *et al.*, “Lossless 4-bit Quantization of Architecture Compressed Conformer ASR Systems on the 300-hr Switchboard Corpus,” in *INTERSPEECH*, 2023.
- [8] Y.-m. Qian and X. Xiang, “Binary neural networks for speech recognition,” *Frontiers of Information Technology & Electronic Engineering*, vol. 20, no. 5, pp. 701–715, 2019.
- [9] Y. Wang, J. Li, and Y. Gong, “Small-footprint high-performance deep neural network-based speech recognition using split-vq,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4984–4988.
- [10] Y. He, T. N. Sainath, R. Prabhavalkar *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.
- [11] E. G. Ng, C.-C. Chiu, Y. Zhang *et al.*, “Pushing the limits of non-autoregressive speech recognition,” *INTERSPEECH*, 2021.
- [12] T. Wang, Y. Fujita, X. Chang *et al.*, “Streaming end-to-end asr based on blockwise non-autoregressive models,” *INTERSPEECH*, 2021.
- [13] Z. Gao, S. Zhang, I. McLoughlin *et al.*, “Paraformer: Fast and accurate parallel transformer for non-autoregressive end-to-end speech recognition,” *INTERSPEECH*, 2022.
- [14] Y. Higuchi, S. Watanabe, N. Chen *et al.*, “Mask ctc: Non-autoregressive end-to-end asr with ctc and mask predict,” *INTERSPEECH*, 2020.
- [15] Y. Higuchi, H. Inaguma, S. Watanabe *et al.*, “Improved mask-ctc for non-autoregressive end-to-end asr,” in *ICASSP*, 2021.
- [16] W. Chan, C. Saharia, G. Hinton *et al.*, “Imputer: Sequence modelling via imputation and dynamic programming,” in *ICML*, 2020.
- [17] N. Chen, S. Watanabe, J. Villalba *et al.*, “Non-autoregressive transformer for speech recognition,” *IEEE Signal Processing Letters*, 2020.
- [18] E. A. Chi, J. Salazar, and K. Kirchhoff, “Align-refine: Non-autoregressive speech recognition via iterative realignment,” in *NAACL*, 2021.
- [19] N. Chen, P. Zelasko, L. Moro-Velázquez *et al.*, “Align-denoise: Single-pass non-autoregressive speech recognition,” in *INTERSPEECH*, 2021.
- [20] Y. Sudo, M. Shakeel, B. Yan *et al.*, “4d asr: Joint modeling of ctc, attention, transducer, and mask-predict decoders,” *INTERSPEECH*, 2023.
- [21] Z. Tian, J. Yi, J. Tao *et al.*, “Hybrid autoregressive and non-autoregressive transformer models for speech recognition,” *IEEE Signal Process. Lett.*, 2022.
- [22] W. Wang, K. Hu, and T. N. Sainath, “Deliberation of streaming rnn-transducer by non-autoregressive decoding,” in *ICASSP*, 2022.
- [23] X. Chang, T. Maekaku, P. Guo *et al.*, “An exploration of self-supervised pretrained representations for end-to-end speech recognition,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 228–235.
- [24] X. Song, Z. Wu, Y. Huang *et al.*, “Non-autoregressive transformer asr with ctc-enhanced decoder input,” in *ICASSP*, 2021.
- [25] W. Qi, Y. Gong, J. Jiao *et al.*, “BANG: Bridging Autoregressive and Non-autoregressive Generation with Large Scale Pretraining,” in *ICML*, 2021.
- [26] X. Liang, L. Wu, J. Li *et al.*, “JANUS: Joint autoregressive and non-autoregressive training with auxiliary loss for sequence generation,” in *EMNLP*, 2022.
- [27] S. Chen, C. Wang, Z. Chen *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [28] T. Ko, V. Peddinti, D. Povey *et al.*, “Audio augmentation for speech recognition,” in *INTERSPEECH*, 2015.
- [29] D. S. Park, W. Chan, Y. Zhang *et al.*, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *INTERSPEECH*, 2019.
- [30] L. Gillick and S. J. Cox, “Some statistical issues in the comparison of speech recognition algorithms,” in *ICASSP*, 1989.