



# Eloquent: A More Robust Transmission Scheme for LLM Token Streaming

Hanchen Li, Yuhan Liu, Yihua Cheng, Siddhant Ray, Kuntai Du, Junchen Jiang  
University of Chicago

## ABSTRACT

To render each generated token in real-time for users, the Large Language Model (LLM) server generates tokens one by one and streams each token (or group of a few tokens) through the network to the user right after generation, which we refer to as *LLM token streaming*. However, under unstable network conditions, the LLM token streaming experience could suffer greatly from stalls since one packet loss could block the rendering of later tokens even if the packets containing them arrive on time. With a measurement study, we show that current applications suffer from increased stalls under unstable networks.

For this emerging token streaming problem in LLM Chatbots that differs from previous multimedia and text applications, we propose a novel transmission scheme, called *Eloquent*, which puts newly generated tokens as well as currently unacknowledged tokens in the next outgoing packet. This ensures that each packet contains some new tokens and, in the meantime, is independently rendered when received, avoiding the aforementioned stalls caused by missing packets. Through simulation under various networks, we show *Eloquent* reduces stall ratio (proportion of token rendering wait time) by 71.0% compared to the retransmission method commonly used by real chatbot applications and by 31.6% compared to the baseline packet duplication scheme. By tailoring *Eloquent* to fit the token-by-token generation of LLM, we enable the Chatbots to respond like an eloquent speaker for users to better enjoy pervasive AI.

## CCS CONCEPTS

• **Networks** → **Application layer protocols**; *Network measurement*; • **Computing methodologies** → *Natural language processing*;

## KEYWORDS

Real-Time Communication, Large Language Models, Token Streaming

### ACM Reference Format:

Hanchen Li, Yuhan Liu, Yihua Cheng, Siddhant Ray, Kuntai Du, Junchen Jiang University of Chicago. 2024. Eloquent: A More Robust Transmission Scheme for LLM Token Streaming. In *SIGCOMM Workshop on Networks for AI Computing (NAIC '24)*, August 4–8, 2024, Sydney, NSW, Australia. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3672198.3673797>



This work is licensed under a Creative Commons Attribution International 4.0 License.

NAIC '24, August 4–8, 2024, Sydney, NSW, Australia

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0713-1/24/08.

<https://doi.org/10.1145/3672198.3673797>

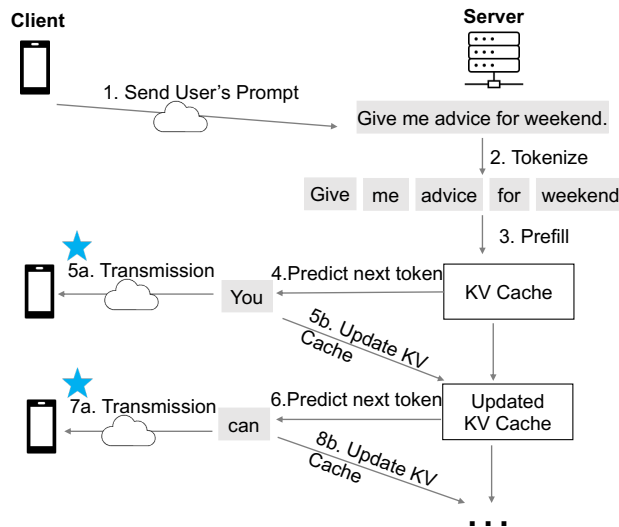


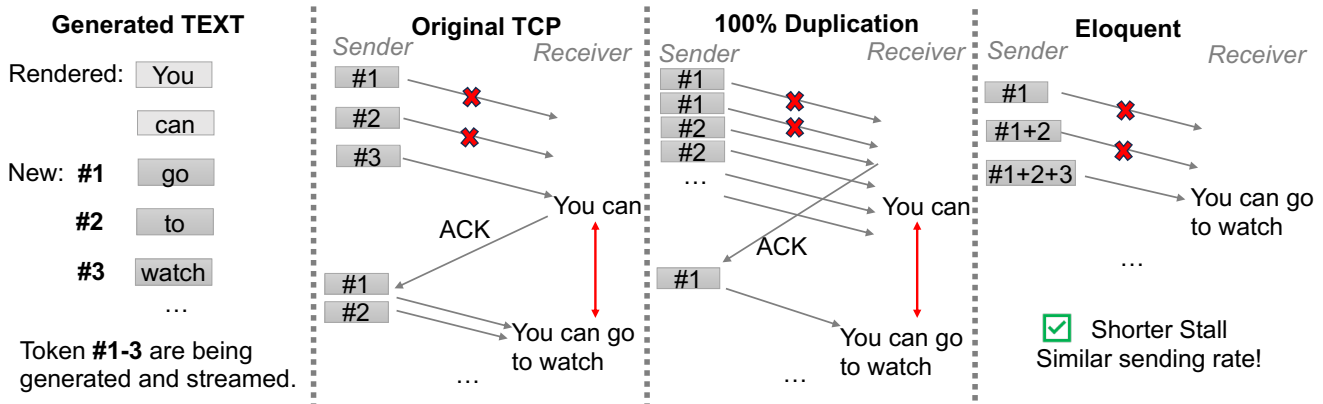
Figure 1: LLM Chatbot Token Streaming Pipeline. We aim to improve the transmission part (Blue Starred 5a, 7a)

## 1 INTRODUCTION

Large Language Models (LLMs) are growing more and more popular [4–6]. One of the most important use cases is the Chatbots that utilize LLM models to take in user prompts and send back real-time responses to answer user queries. [7–9]. Significant effort has been made in the machine learning and system community to generate better responses more efficiently [16, 22, 25, 36, 39, 41]. However, the *streaming* of these generated tokens to users is also critical. Having intermittent stalls during transmission could break the smoothness of interaction and damage user experience.

In a token streaming pipeline of LLM Chatbot, tokens are generated in the cloud data center after the data center receives the prompt. After each token is generated, they are sent over the wide-area network to reach the client device. If some packets are dropped or significantly delayed in the network, the client will need to wait for packets to be retransmitted from the sender. This could cause an undesired stall for token streaming since all later tokens need to wait for previous ones to arrive at the receiver before being rendered.

As AI services become ubiquitous, people could use these LLM Chatbot services anywhere, including places where intermittent connections happen, such as on railway [29, 40] or using a lossy wifi link [42]. To observe how current applications perform under such scenarios, we conducted a measurement study to understand how an unstable network affects ChatGPT Streaming API. By reading the packet traces, we see that these applications rely on retransmission to handle packet loss, which aggravates the long stall. Moreover,



**Figure 2:** An illustrative example. When the packets containing the first two tokens are lost, Eloquent significantly reduces stall with the similar overall sending rate, compared to TCP or duplicating each packet twice. We use #1, #2, #3 to represent the newly generated tokens in later columns and only show ACKs that could signal retransmission.

we see an interesting pattern—after some packets are lost, the subsequent packets containing newer tokens may arrive earlier than the retransmission packet. Even so, no more new tokens could be rendered as they are blocked by the lost tokens. Every token must still wait for retransmission before rendering. However, under an unstable network, the Round Trip Time (RTT) could be long [18]. This means that all retransmission mechanisms in TCP will suffer from a long response delay, and this causes stall [20, 30]. A natural solution is to add redundancy. Unfortunately, sending more duplicated packets together with the original token could fail if the connection was lossy during that time. Sending Forward Error Correction (FEC) packets at the end of a token group will only help after FEC packets arrive at the end, which could already be too late.

Although text streaming is not entirely new, but token streaming *differs* in a few key aspects. For instance, collaborative editors (like Overleaf or Google Docs) also stream keystrokes from users. However, compared to LLM’s token generation speed, the typical rate at which users manually type words is arguably slower [11], making TCP retransmission a sensible solution to handle packet loss without backlogging words on the sender side. On the other hand, virtual assistants do automatically generate responses. Still, until recently, their text generation models were fast enough to produce the entire response before sending the response to the user as a whole [1]. In contrast, LLMs generate tokens *autoregressively*, so every token has to be sent out to users as soon as possible in order to ensure a smooth user experience.

To tackle this new *token streaming* problem under unstable networks, we proposed a novel transmission scheme, Eloquent, that provides a smooth token streaming experience and transforms Chatbot into a more eloquent speaker under various network conditions by adding smart redundancy. Demonstrated in Fig 2, Eloquent adds *unacked tokens*, defined as tokens whose packets have been sent but have not been ACK’ed by the receiver, into the packet of the newly generated token. It prepares each packet in such a way that each received packet contains sufficient information for the rendering of the new token contained inside, preventing the retransmission blocking issue discussed above.

Through our simulation under various network conditions, Eloquent can reduce stall ratio<sup>1</sup> by 71.0% compared to the retransmission method used by TCP, and by 31.6% with less total data sent than baseline packet duplication method.

Admittedly, there remain challenges in deploying Eloquent inside real-world protocols like RTP or QUIC. We hope this work sparks more discussion around taking system components outside datacenters into consideration in the LLM serving pipeline. For this end, we list several future research directions in §5 along this path.

## 2 BACKGROUND AND MOTIVATION

### 2.1 Token Streaming

We provide a brief end-to-end overview of the current LLM Chatbot streaming pipeline in Fig. 1. First, the user enters a prompt, which is then sent to the cloud data center. In the data center, the prompt is split into tokens (partial or full words). These tokens then go through a prefill phase to process the prompt in a parallel fashion. After the prefill phase, the model iteratively generates the next token with the current sentence, appends the new token to the sentence, and repeats the process until an end-of-sentence (EOS) token is generated. In this output phase, the latest generated tokens are sent to the user through the network right after the generation, rather than waiting for all tokens to be generated, in order to provide a real-time streaming experience. The tokens are rendered on the screen in the order that they are generated. In this paper, we focus on the transmission step of *token streaming* under unstable network conditions.

### 2.2 Application Measurement

As LLM Chatbots are becoming more pervasive in our lives, it is important that the services deliver smooth experiences under different network conditions. However, as previous research and reports [2, 29, 31, 40, 42] have pointed out, intermittent connection could happen in many real-world scenarios such as on railway or using a lossy wi-fi link. In these cases, groups of packet might get

<sup>1</sup>Proportion of token wait time when inter-token gap exceeds 200ms

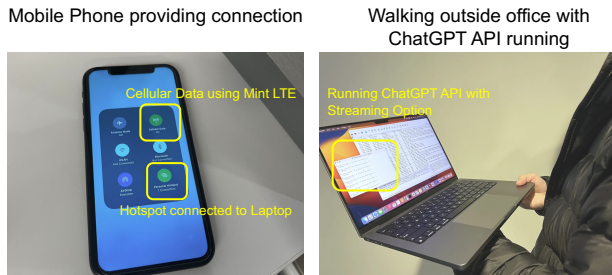


Figure 3: Our Network Measurement Testbed




Figure 4: Packet arrival and token rendering of one measured session. Retransmission packets blocking rendering are in circled in red. Blocked packets circled in blue.

lost during a short period of time due to a physical surface block or signal interference from other devices.

To learn about how current LLM Chatbot applications perform under these unstable network conditions, we conducted a measurement study on one of the most popular tools to build LLM Chatbot, ChatGPT API.

**Measurement Testbed:** We emulate an unstable network by connecting a Macbook 2023 M2 laptop to the hotspot of iPhone 11 using Mint LTE service in North America as shown in Fig. 3. We put the phone inside an office and walked around outside the office (within 30m). Meanwhile, on the laptop, we run an automated bash script that repetitively starts 30 GPT API sessions using the same prompt with the “streaming” option enabled. The prompt was set to “Count to 80, with a comma between each number and no newlines. E.g., 1, 2, 3, . . .” for controlled outputs. We record packet arrivals with tcpdump [10] with the interface flag set to “any”. To record token arrival times, we slightly modified the sample Python script provided in OpenAI documentation [3].

**Findings:** The packet trace shows that the data is transmitted through encrypted TLS protocol (over TCP). Performance-wise, under this unstable network condition, compared with the reference of connecting laptop directly to high-speed wifi in the office, the 95% token rendering stall<sup>2</sup> increased from 377ms to 3483ms and from 1501.4ms to 2924.4ms, respectively for model “GPT-3.5-turbo” and “GPT-4”, under an average network loss rate of 10.5% and 15.8%.

We provide a detailed analysis of a single trace to show why there is potential for improvement in Fig. 4. As the packet arrival graphs show on the left, some packets (red-circled) were lost and retransmitted. However, many packets containing later tokens arrived before the retransmission packets, which arrived much later

<sup>2</sup>We defined stall as inter-token gaps larger than 200ms. This follows from works in video streaming [15, 28].

at around the 9th second due to the unstable connection. The impact of this “retransmission blocking future token rendering” is severe, as shown on the right-hand side. No tokens were rendered from the moment these packets were lost until the retransmission came through in the 9th second. While many packets arrived after the loss, all of the tokens contained inside were blocked by the lost packets. Thus, a stall lasted almost 5 seconds purely due to the transmission glitch, even though new tokens kept arriving during that period.

## 2.3 Previous Work

Although the problem of token streaming is new, there have been prior works that aim to protect transmission against packet loss for other media forms.

The most popular method for loss-resilience is Forward Error Correction (FEC) [26, 33, 37]. There are different types of coding methods [26, 27, 37], and the fundamental idea is to encode an a-byte segment to (a+b) bytes such that when any a of the (a+b) bytes arrive, the receiver can reconstruct the original data. Moreover, FEC protects a file (video/audio frame) that spans multiple packets to recover partially delivered data. In token streaming, the minimal transmission unit is one single token, which is far less than the maximal size of one packet.<sup>3</sup> There have been works that target intermittent connectivity for text-based applications including Mosh for interactive remote shell [38]. Mosh predicts output of keystrokes at client-side before the server sends responses back. However, these user keystrokes are sent to the server by clients in the first place, and their output is easy to predict (the output of typing h,e,l,l,o is simply showing “hello” on the terminal). In token streaming, the texts are generated by the server-side LLM and sent back to the client and could hardly be predicted by the client. Cooperative document editor (ex. Google doc) is a similar application. But while LLM generation speed is 10 tokens/s [12], human typing speed is 30-45 words/min [11]. This makes online editing already have long gaps (~1s) between tokens. Virtual Assistants like Siri utilizes multi-path TCP [1] and only reveal content after the whole response has been received, which poses stronger assumption on client device and potentially leads to high time to first token (TTFT) if adapted to LLM Chatbots.

From the LLM systems, although there have been extensive work on scheduling [39, 41], memory management [16, 22, 36], speculative decoding [24], and QoE-aware inference [25], all of them solely optimize for the generation inside the cloud datacenter before the tokens are generated but does not include the delivery of the tokens. As long as the user is following the newest generated content of the chatbot and packets get lost in the middle, a long stall might still happen and damage user experience.

## 2.4 Unique Properties and Challenge

This *token streaming* problem bears some resemblance to the traditional streaming problems, but they have several distinct characteristics.

- Unlike video/audio streaming, increasing the bitrate for transmission does not increase rendering quality.

<sup>3</sup>Single token packet for Bard is 50-80 bytes; ChatGPT packet is 842 bytes for three tokens.

- Sequence needs to be strictly maintained so that one lost token could block future rendering. In contrast, video/audio may skip frame [17] under loss.
- The minimal bandwidth requirement for the network is much lower since text data are much smaller than audio (150Kbps) and video (>1Mbps).

Given these properties, we believe that in order to render tokens smoothly even under unstable networks, a robust token streaming method should prevent the problem of lost packet blocking subsequent arrived tokens rendering. For this to happen, a moderate increase in total sent data is acceptable. The main challenge we face is intermittent connectivity, while the average bandwidth could stay over 100Kbps even when the network is unstable [23]. Moreover, we could add redundancy without suffering from quality drops.

### 3 ELOQUENT DESIGN

We propose Eloquent, a novel transmission scheme that puts “unacked tokens” in the outgoing packets that contain the newly generated tokens, in order to reduce client token streaming stalls under an unstable network.

The key idea of Eloquent is to prepare packets so that the client can render each received packet independently. That is, each packet contains not only newly generated tokens but also enough information to render these new tokens. Thus, for each arrived new token, the client is able to render the current sentence all the way until this new token and would not potentially be blocked by any previously lost token. With this mechanism, retransmission-incurred stalls can be greatly reduced, especially under unstable network where round-trip-time (RTT) can be long.

#### 3.1 Illustrative Example

We showcase one example case of Eloquent comparing with TCP and naive redundancy in Fig. 2. In this scenario, we are streaming a response of “You can go to watch...” and the client has received tokens to render “You can”. The LLM model now outputs “go”, “to”, and “watch” sequentially with time gap in between. We assume the unstable network will drop the first two packets sent and successfully deliver the following packets.

In the current TCP retransmission logic, after the two packets containing tokens #1 and #2 are lost, token #3 successfully arrives at the receiver. The receiver sends an ACK back to the sender, and the sender retransmits the first two tokens to the receiver again to finally render the whole sentence.

In the second case, we added one duplication packet for each outgoing token. However, if the first two packets are lost, the first token is still lost during transmission. Although tokens #2 and #3 successfully arrive at the receiver, their rendering is blocked by token #1. The client still has to wait for the retransmission of token #1 to render the sentence.

In the last column, we demonstrate how our method could reduce the stall. After the first token is sent, the following packets will incorporate token #1 into the payload. Similarly token #2 is contained in later packets. Thus, when the packet containing the new token #3 arrives at the receiver, the receiver is able to render the whole sentence since the previous tokens needed to render #3 are

all contained in the packet. This greatly reduces the stall perceived on the client side. Moreover, the total size of sent data is similar to the original TCP and less than duplication since retransmission adds extra workload to the system.

---

#### Algorithm 1: Eloquent Sender Logic

---

```

unacked ← empty
while EOS Token not Acked do
  Sleep until newly generated token or
  (unacked_buffer ≠ empty and slept for 200ms)
  unacked ← unacked \ new_acked_tokens
  if size of unacked ≤ max size of packet then
    packet ← token concatenated with unacked
  else
    remainingSize ← max size of packet – size of token
    partial_unacked ← empty
    for i ← 0 to len(unacked) do
      if sizeof(partial_unacked) + sizeof(unacked[i]) <
      remaining_size then
        partial_unacked.append(unacked[i])
      else
        Break
    end if
  end for
  packet ← packet concatenated with partial_unacked
  end if
  Send packet
  Append token into unacked
end while

```

---

#### 3.2 Specific Workflow

We provide the pseudo-code for Eloquent scheme in algorithm 1 and algorithm 2. The core logic of the sender is to maintain a unacked token buffer that contains the all tokens that were sent but have not received ACK. Every time a new token is generated, we visit this buffer to put as many unacked tokens as maximal size allows into the packet with the new token to prepare for the rendering of new token. The receiver follows a simpler logic. After receiving a packet, it extracts the tokens inside, renders the tokens if all previous tokens have arrived, and sends back ACKs to the sender. This scheme could be implemented inside RTP [35] or QUIC [19], since we allow sent data to never arrive at client side.

#### 3.3 Limitation

Admittedly, our method has limitations. The most notable limitation is that sometimes not all unacked tokens could fit into one packet. Mathematically, this happens when:

$$G \times (T - 1) \leq 2 \times RTT + L,$$

where  $G$  is the time gap between two generated tokens,  $T$  is the maximum number of tokens in one packet,  $RTT$  is the round trip time, and  $L$  is the packet loss time.

As measured in [12], most of the time current GPT4 throughput is around 10 tokens/s. Since human normal reading speed is 4-6 words per second (fast reading speed is 8-15 words per second) [32],

---

**Algorithm 2: Eloquent Receiver Logic**

---

```
acked_list ← empty list
while EOS was not received do
  Sleep until receives Packet
  token_list ← tokens in packet
  for token in token_list do
    if token not acked before then
      acked_list ← acked_list append token
    end if
  end for
  Render new tokens according to acked_list
  Send ACK for packet
end while
```

---

LLM service providers have less incentive to purchase more GPUs to drastically increase throughput beyond this level when more users contend for the computing resources in the future. On the other hand, it is possible to reuse the meta-data inside each packet to save spaces for the token data. In worst case, if the unacked tokens cannot all fit in the new packet, we only add the earlier unacked tokens into the packet in order to make sure that some tokens can still be independently rendered after arrival.

## 4 EVALUATION

### 4.1 Setup


We measured our performance in an end-to-end Python simulation platform, which takes in generation timestamps of tokens, network channel parameters, as well as a sender/receiver pair, and outputs rendering time for each token. We report two metrics in our report: redundancy rate and stall ratio. Redundancy rate is the extra amount of data sent (including retransmission) besides the original tokens. We borrowed the idea of stall ratio from video streaming [15, 28] and defined it to be the amount of time that client side is waiting for tokens to render when stall is longer than 200ms. As discussed in §2.4, an ideal method should prioritize reducing stall ratio and try to have lower redundancy rate afterwards.

We implemented a pair of Eloquent sender/receiver in ~ 70 lines of code and set the maximal number of tokens in one packet to be 10. We also implemented a group of packet duplication senders, which send multiple copies of every packet to prevent loss (similar to FEC in video streaming).


In our experiments, We set input generation sequence to generate new token every 100ms as the current GPT-4 generation throughput is around 10 tokens/s [12]. For each setting, we run 30 30-second sessions.

### 4.2 Experiment Results

In this part, we utilize a two-state Markov model [21, 34] to simulate unstable network conditions like wireless connections. We set two state: lossy state, when packets are lost with 90% possibility; and good states, when no packets are lost. Each state lasts for 100ms. There are two parameters  $p$  and  $q$ :  $p$  is the probability to transform from good state to lossy state and  $q$  is the probability to change in the opposite direction. On average, the channel enters lossy state



**Figure 5: Eloquent reduces stall ratio under 15% loss rate. Duplication-X are with rate 2x, 3x, 4x, 5x from left to right.**



**Figure 6: Eloquent performs well under various loss patterns**

for  $\frac{100}{q}$  ms after staying in good state for  $\frac{100}{p}$  ms. Our default case, where  $p = 0.9$  and  $q = 0.5$ , has 200ms bad connections every 1000ms good connection on average. This gives a packet loss rate of 15%, similar to what we observed in the measurement in §2.2.

We start with the default setting and set round trip time (RTT) to 400ms and 100ms in Fig. 5. As shown on the left, Eloquent is able to reduce stall ratio by 71.0% compared with the original TCP baseline. This is because Eloquent reduces the need for retransmission under packet loss. Comparing with packet duplication methods, Eloquent is able to reduce stall by 31.6% percent with similar redundancy rate. The right graph with 100ms RTT shows a similar result, when Eloquent reduces stall ratio by 60.0% compared with TCP and 31.3% compared with duplication with less size.

To learn about how Eloquent performs under different loss patterns, we fixed RTT to be 400ms and modified the loss parameters in Fig. 6. The left graph shows results with  $p = 0.5$  and  $q = 0.5$ . This network is more likely to enter lossy state than default. Since loss is more frequent, the stall ratios of all methods increased. But Eloquent’s improvements is more significant since packet loss are more frequent. On the right side we set  $p$  to 0.9 and  $q$  to 0.8. In this case, every lossy period lasts longer. As shown in the right hand side, although we still outperforms the baselines, our improvement decreased due the limitation mentioned in §3.3. Since both loss time (L) and RTT are big, many packets could not contain all unacked tokens, and this damages Eloquent’s performance. This potentially could be improved by reusing packet meta-data (font type, size...) to put more tokens in each packet.

### 4.3 Real Trace Experiment

We utilized packet arrival traces collected from the emulation platform in §2.2 to conduct end-to-end simulation under real-world loss patterns. We calculated whether the  $i$ th token packet was lost




Figure 7: Eloquent reduces stall ratio and late tokens greatly in Real Trace Driven Experiment

by TCP sequence number and used this data to determine packet loss in the simulation sender channel.

As shown in Fig 7, under 400ms RTT, Eloquent is able to reduce stall ratio by 27.4% compared with TCP and 18.0% compared with packet duplication with less size. Moreover, it reduces the proportion of late tokens (delay since generation over 400ms) by 40.4% over TCP, providing more real-time responses with no modification to computation.

## 5 FUTURE DIRECTIONS

### 5.1 Practical implementation

This paper only intends to demonstrate the promising improvement of a new transmission scheme over unstable network. More work is still needed before implementing Eloquent in practice.

To start with, since Eloquent does not require all sent bytes to receive at the receiver, Eloquent should be implemented within one of the transport protocols based on UDP. For reference, Eloquent could be implemented in QUIC for the stream of generated text as its loss recovery method, or on top of RTP similar to audio/video streaming. Different from the default TCP method that maintains the unacked buffer inside the operating system, Eloquent maintains the unacked buffer inside application space.

### 5.2 Quality of Experience for Token Streaming

There could be further effort on establishing the quality of experience (QoE) model for token streaming. Prior results in video QoE could provide insight on new system optimization. For example, in video streaming it is established that lower quality frames are better than having stalls [14, 15]. But in LLM Chatbot services people might be willing to suffer from rendering stall in order to receive more sophisticated answers. Moreover, users might have preference for a single long stall over frequent stalls. If this holds, then maintaining a token buffer at client side could help bring better experience. Building a better QoE model for LLM inference potentially improve our tradeoff curve for inference systems.

### 5.3 QoE-aware LLM Inference

Prior works including Andes [25] and Google’s Bard [9] maintains a buffer at client side to gradually release text at human reading speed. This damages the real-timeliness of response but increases smoothness while alleviating the pressure on system scheduling. Eloquent could be combined with these systems to selectively enter

redundancy mode when the buffer is close to empty since a potential packet loss would cause more damage during that time.

Moreover, we could also combine different inference systems. For example, VLLM optimizes throughput by increasing the batch size but may incur generation stall [22]. While Sarathi-Serve has smoother generation with potentially lower throughput [13]. With a user-side token buffer, we could optimize throughput while the buffer has many tokens and switch to methods like Sarathi-Serve to prevent generation stall when buffer is shallow.

## 5.4 Exposing user state in LLM serving

Eloquent targets the problem of token streaming after the tokens are generated by the data center. There could also be works that sends signals to data center to improve token generation. For example, when users switches tab, a signal could be transmitted back to the data center to deprioritize smoothness since users are not actively reading the content. On the otherhand, if the network is detected to be lossy, the client could try to switch to buffering mode and notify the user that the content will be available later.

## 6 CONCLUSION

In conclusion, we identify the transmission problem in the token streaming pipeline for LLM Chatbots and reveal the degraded performance of current application under unstable network with a measurement study. To improve the novel token streaming problem, we propose Eloquent, a novel transmission scheme that proactively puts unacked tokens into packets of newly generated tokens. Through simulation, Eloquent reduces stall ratio by 71.0% over TCP and 31.6% over baselines. In the end, we point out future research directions optimizing LLM inference with components outside data centers in consideration to spark more discussion.

## ACKNOWLEDGEMENT

We thank all the anonymous reviewers, Xu Zhang, Qizheng Zhang, and Zhengxu Xia for their constructive feedback. This project is supported by NSF CNS-2146496, CNS-1901466, CNS-2313190.

## REFERENCES

- [1] Multipath TCP lets Siri seamlessly switch between Wi-Fi and 3G/LTE. <https://arstechnica.com/gadgets/2013/09/multipath-tcp-lets-siri-seamlessly-switch-between-wi-fi-and-3g/le/>, 2019.
- [2] 5 Big Wireless Challenges: Signal Loss, Movement, Reach, Densification and Multipath Fading. <https://www.belden.com/blogs/smart-building/5-big-wireless-challenges-signal-loss-movement-reach-densification-and-multipath-fading/>, 2020.
- [3] API Reference - OpenAI API. <https://platform.openai.com/docs/api-reference/chat/create>, 2023.
- [4] Bring Gen AI & LLMs to Your Data . <https://www.snowflake.com/blog/generative-ai-llms-summit-2023/>, 2023.
- [5] Global Large Language Model(LLM) Market Research Report 2023 . <https://reports.valuates.com/market-reports/QYRE-Auto-30B13652/global-large-language-model-llm>, 2023.
- [6] LLMs are becoming increasingly popular in the enterprise world. <https://uptrain.medium.com/llms-are-becoming-increasingly-popular-in-the-enterprise-world-as-businesses-recognize-their-2ac2a61771bd>, 2023.
- [7] Anthropic Claude. <https://www.anthropic.com/>, 2024.
- [8] ChatGPT. <https://chat.openai.com/>, 2024.
- [9] Google Bard. <https://bard.google.com/chat>, 2024.
- [10] TCPDUMP documentation. <https://www.tcpdump.org/>, 2024.
- [11] Typing Speed Test. <https://www.livechat.com/typing-speed-test/#/>, 2024.
- [12] Unofficial OpenAI Status. <https://openai-status.llm-utils.org/>, 2024.
- [13] Amey Agrawal, Nitin Kedia, Ashish Panwar, Jayashree Mohan, Nipun Kwatra, Bhargav S. Gulavani, Alexey Tumanov, and Ramachandran Ramjee. Taming throughput-latency tradeoff in llm inference with sarathi-serve, 2024.
- [14] Nabajeet Barman and Maria G. Martini. Qoe modeling for http adaptive video streaming—a survey and open challenges. *IEEE Access*, 7:30831–30859, 2019.
- [15] Yihua Cheng, Ziyi Zhang, Hanchen Li, Anton Arapin, Yue Zhang, Qizheng Zhang, Yuhua Liu, Xu Zhang, Francis Y. Yan, Amrita Mazumdar, Nick Feamster, and Junchen Jiang. Grace: Loss-resilient real-time video through neural codecs, 2023.
- [16] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness, 2022.
- [17] Sadjad Fouladi, John Emmons, Emre Orbay, Catherine Wu, Riad S Wahby, and Keith Winstein. Salsify: {Low-Latency} network video through tighter integration between a video codec and a transport protocol. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 267–282, 2018.
- [18] Junxian Huang, Feng Qian, Alexandre Gerber, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. A close examination of performance and power characteristics of 4g lte networks. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services, MobiSys '12*, page 225–238, New York, NY, USA, 2012. Association for Computing Machinery.
- [19] J. Iyengar and M. Thomson. QUIC: A UDP-Based Multiplexed and Secure Transport. <https://www.rfc-editor.org/rfc/rfc9000.html>, 2021. RFC 9000.
- [20] Bart Jansen, Timothy Goodwin, Varun Gupta, Fernando Kuipers, and Gil Zussman. Performance evaluation of webrtc-based video conferencing. *SIGMETRICS Perform. Eval. Rev.*, 45(3):56–68, mar 2018.
- [21] Syed A. Khayam and Hayder Radha. Markov-based modeling of wireless local area networks. In *Proceedings of the 6th ACM International Workshop on Modeling Analysis and Simulation of Wireless and Mobile Systems, MSWIM '03*, page 100–107, New York, NY, USA, 2003. Association for Computing Machinery.
- [22] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention, 2023.
- [23] Kilho Lee, Daehyeok Kim, and Insik Shin. REboost : Improving throughput in wireless networks using redundancy elimination. *IEEE Communications Letters*, 21(1):160–163, 2017.
- [24] Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding, 2023.
- [25] Jiachen Liu, Zhiyu Wu, Jae-Won Chung, Fan Lai, Myungjin Lee, and Mosharaf Chowdhury. Andes: Defining and enhancing quality-of-experience in llm-based text streaming services, 2024.
- [26] David JC MacKay. Fountain codes. *IEE Proceedings-Communications*, 152(6):1062–1068, 2005.
- [27] David JC MacKay and Radford M Neal. Near Shannon limit performance of low density parity check codes. *Electronics letters*, 33(6):457–458, 1997.
- [28] Kyle MacMillan, Tarun Mangla, James Saxon, and Nick Feamster. Measuring the performance and network utilization of popular video conferencing applications. In *Proceedings of the 21st ACM Internet Measurement Conference, IMC '21*, page 229–244, New York, NY, USA, 2021. Association for Computing Machinery.
- [29] Yunzhe Ni, Feng Qian, Taide Liu, Yihua Cheng, Zhiyao Ma, Jing Wang, Zhongfeng Wang, Gang Huang, Xuanzhe Liu, and Chenren Xu. POLYCORN: Data-driven cross-layer multipath networking for high-speed railway through composable schedulerlets. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 1325–1340, Boston, MA, April 2023. USENIX Association.
- [30] Jon Postel. Transmission control protocol. <https://www.rfc-editor.org/rfc/rfc793.txt>, September 1981. RFC 793.
- [31] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee. Diagnosing wireless packet losses in 802.11: Separating collision from weak signal. In *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*, pages 735–743, 2008.
- [32] Keith Rayner, Elizabeth R. Schotter, Michael E. J. Masson, Mary C. Potter, and Rebecca Treiman. So much to read, so little time: How do we read, and can speed reading help? *Psychological Science in the Public Interest*, 17(1):4–34, 2016. PMID: 26769745.
- [33] Michael Rudow, Francis Y Yan, Abhishek Kumar, Ganesh Ananthanarayanan, Martin Ellis, and KV Rashmi. Tambur: Efficient loss recovery for videoconferencing via streaming codes. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 953–971, 2023.
- [34] Samuel H. Russ and Sasan Haghani. Behavior of 802.11g traffic at high sustained bit rates in the home. In *2009 Digest of Technical Papers International Conference on Consumer Electronics*, pages 1–2, 2009.
- [35] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A transport protocol for real-time applications. Internet RFC 3550, July 2003. RFC 3550.
- [36] Ying Sheng, Lianmin Zheng, Binhang Yuan, Zhuohan Li, Max Ryabinin, Daniel Y. Fu, Zhiqiang Xie, Beidi Chen, Clark Barrett, Joseph E. Gonzalez, Percy Liang, Christopher Ré, Ion Stoica, and Ce Zhang. Flexgen: High-throughput generative inference of large language models with a single gpu, 2023.
- [37] Stephen B. Wicker and Vijay K. Bhargava. *Reed-Solomon codes and their applications*. John Wiley & Sons, 1999.
- [38] Keith Winstein and Hari Balakrishnan. Mosh: An interactive remote shell for mobile clients. In *2012 USENIX Annual Technical Conference (USENIX ATC 12)*, pages 177–182, Boston, MA, June 2012. USENIX Association.
- [39] Bingyang Wu, Yinmin Zhong, Zili Zhang, Gang Huang, Xuanzhe Liu, and Xin Jin. Fast distributed inference serving for large language models, 2023.
- [40] Chenren Xu, Jing Wang, Zhiyao Ma, Yihua Cheng, Yunzhe Ni, Wangyang Li, Feng Qian, and Yuanjie Li. A first look at disconnection-centric tcp performance on high-speed railways. *IEEE Journal on Selected Areas in Communications*, 38(12):2723–2733, 2020.
- [41] Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. Orca: A distributed serving system for Transformer-Based generative models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 521–538, Carlsbad, CA, July 2022. USENIX Association.
- [42] Gina Yuan, Matthew Sotoudeh, David K. Zhang, Michael Welzl, David Mazières, and Keith Winstein. Sidekick: In-Network assistance for secure End-to-End transport protocols. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, pages 1813–1830, Santa Clara, CA, April 2024. USENIX Association.