International Journal of Scientific Research in Computer Science, Engineering and Information Technology

ISSN : 2456-3307 OPEN ACCESS

Available Online at : www.ijsrcseit.com

doi : https://doi.org/10.32628/CSEIT2390660

Plausibility of BBR as CUBIC'S Replacement and proposed improvement to BBR using GENET

Dr. Kunwar Asif, Hemanjali Kadali, Sathwik Varma Mudduluri, Pawan Sai Krishna Reddy Kerelly

Department of Computer Science, New Jersey Institute of Technology, New Jersey, USA

ARTICLEINFO

Article History:

Accepted: 10 Dec 2023

Published: 26 Dec 2023

Publication Issue

Volume 9, Issue 6

November-December-2023

Page Number

303-309

ABSTRACT

Recent advancements in deep reinforcement learning (RL) have opened new avenues for enhancing network congestion control algorithms. Our research builds upon these developments, particularly focusing on the BBR (Bottleneck Bandwidth and Round-trip propagation time) congestion control algorithm. We propose integrating GENET's reinforcement learning framework, a novel training paradigm that has demonstrated success in various network adaptation algorithms, including adaptive video streaming, congestion control, and load balancing. GENET leverages curriculum learning to effectively train RL models by progressively introducing more challenging network environments. This method counters the common pitfalls in RL training, such as suboptimal performance in a wide range of environments and poor generalization in narrowly defined training scenarios. Our approach exploits the strengths of GENET in identifying and emphasizing network conditions where the current RL model underperforms compared to traditional rule-based baselines, thereby facilitating significant improvements. This research aims to demonstrate that applying GENET's methodology to the BBR congestion control algorithm can yield RL policies that surpass both regularly trained RL policies and conventional baselines, thereby advancing the efficiency and reliability of network congestion control.

Keywords: Bottleneck Bandwidth and Round-trip propagation time algorithm, congestion control algorithms, GENET's

I. INTRODUCTION

The ever-evolving landscape of Internet traffic management necessitates continual advancements in congestion control mechanisms. Among these, the BBR (Bottleneck Bandwidth and Round-trip propagation time) congestion control algorithm has emerged as a significant innovation. Developed by Google, BBR represents a paradigm shift from

Copyright © 2023 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution **4.0 International License (CC BY-NC 4.0)** which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

traditional congestion control algorithms, which typically rely on packet loss as a signal to gauge network congestion. Instead, BBR utilizes estimates of the network's bottleneck bandwidth and round-trip time to adjust its sending rate, aiming to optimize throughput while minimizing latency.

Despite its advancements, BBR's performance can be further enhanced, particularly in dynamic network environments where predicting conditions such as bandwidth and latency is challenging. This is where GENET (Generic Evolutionary Network), a deep reinforcement learning (RL) framework, becomes pivotal. GENET introduces a novel approach to training RL-based network adaptation algorithms. It employs curriculum learning, a technique that progressively exposes the model to more complex network scenarios. This method is particularly effective in overcoming common challenges in RL, such as suboptimal performance across diverse environments and poor generalization when trained in narrowly defined conditions.

In our research, we aim to explore the potential of improving BBR's congestion control algorithm by employing GENET's reinforcement learning capabilities. The core idea is to use RL to predict network characteristics that are crucial for BBR's algorithm to adjust the sending rate effectively. By accurately forecasting network conditions like bandwidth availability and latency, we anticipate that BBR can make more informed decisions, thereby enhancing its efficiency and adaptability.

This integration of GENET's RL techniques with BBR's congestion control algorithm is an innovative approach that has the potential to significantly advance the field of network traffic management. Through this research, we seek to not only improve upon the existing capabilities of BBR but also contribute to the broader understanding of how AI and machine learning can be effectively leveraged in networking to address complex, dynamic challenges.

II. LITERATURE REVIEW

The realm of congestion control in network communications has been a subject of extensive research, leading to the development of various algorithms and methodologies. This review focuses on two significant contributions in this field: the BBR congestion control algorithm and GENET's reinforcement learning framework. We explore their performance in different network conditions, examine previous attempts to improve congestion control algorithms, particularly through reinforcement learning, and critically assess the strengths and weaknesses of these approaches.

BBR Congestion Control Algorithm:

The BBR (Bottleneck Bandwidth and Round-trip propagation time) algorithm, developed by Google, marks a notable departure from traditional congestion control algorithms that primarily rely on packet loss as a congestion signal. BBR's unique approach involves estimating the network's bottleneck bandwidth and round-trip time to adjust its data transmission rate, aiming to optimize network throughput while minimizing latency. This method has shown significant improvements over classic algorithms like TCP Cubic, especially in environments with high bandwidth and low latency. However, BBR's performance can degrade in networks with highly variable bandwidth or in scenarios where packet loss is not directly indicative of congestion, leading to potential issues like underutilization of available bandwidth or unfairness when coexisting with loss-based algorithms.

GENET's Reinforcement Learning Framework: GENET (Generic Evolutionary Network) introduces a novel reinforcement learning framework specifically tailored for network adaptation algorithms. It utilizes curriculum learning, gradually exposing the learning model to increasingly complex network environments. This approach addresses common challenges in reinforcement learning, such as suboptimal performance across diverse environments and poor generalization in narrowly defined training conditions. GENET has demonstrated its effectiveness in various case studies, including adaptive video streaming, congestion control, and load balancing, outperforming both traditional rule-based and regularly trained RL algorithms. The strength of GENET lies in its ability to identify and focus on network conditions where the current RL model underperforms, thereby facilitating significant improvements.

Previous Attempts at Reinforcement Learning in

Congestion Control:
Reinforcement learning has been increasingly applied to improve congestion control algorithms. Prior research has explored RL for dynamically adjusting transmission rates based on network feedback, aiming to enhance throughput and reduce latency. While these efforts have shown promise, they often struggle with the inherent unpredictability and complexity of real-world network conditions. The challenge lies in training models that can generalize well across diverse environments without sacrificing performance in specific scenarios.

Critical Evaluation:

The integration of reinforcement learning into congestion control algorithms, like BBR, offers a promising avenue for advancements. However, this approach is not without its challenges. RL models require extensive training and are often sensitive to the variability in network conditions. GENET's curriculum learning approach mitigates some of these issues by focusing training on environments where the model underperforms, but the effectiveness of this method in real-world, dynamic network scenarios remains to be thoroughly evaluated. Additionally, while BBR significantly improves upon traditional congestion control algorithms in certain aspects, its performance in varied network conditions, particularly in competition with loss-based algorithms, highlights areas for further improvement.

In conclusion, the literature underscores the potential of using advanced reinforcement learning techniques, like GENET, to enhance congestion control algorithms, such as BBR. While these approaches demonstrate considerable strengths, they also reveal limitations that need to be addressed through continued research and development.

III.METHODOLOGY

This section outlines the methodology employed in our research to enhance the BBR congestion control algorithm using GENET's reinforcement learning framework. It includes a detailed description of the BBR algorithm, an overview of GENET, the proposed modifications to BBR, and the simulation setup for testing.

3.1 Detailed Description of the BBR Algorithm:

BBR (Bottleneck Bandwidth and Round-trip propagation time) is a congestion control algorithm developed by Google. Unlike traditional algorithms that rely on packet loss as a congestion indicator, BBR estimates the network's bottleneck bandwidth and round-trip time. It operates in two primary modes 'Startup' for bandwidth discovery and 'Drain' to eliminate excess queueing. BBR then uses these estimates to adjust its data sending rate, aiming to maximize network throughput while minimizing latency. The algorithm operates in a cycle of probing for bandwidth, draining any queue it has built, and then cruising at the estimated bandwidth.

3.2 Overview of GENET's Reinforcement

Learning Framework:

GENET (Generic Evolutionary Network) employs a novel reinforcement learning framework, built on curriculum learning, to adapt network algorithms effectively. Curriculum learning in GENET introduces increasingly complex network scenarios to the RL model progressively. This approach aims to overcome challenges like suboptimal performance and poor generalization in diverse network environments. GENET identifies network conditions where the RL model underperforms compared to rule-based baselines and emphasizes training in these scenarios, promoting substantial improvements in performance.

3.3 Proposed Modifications to BBR Using GENET:

Our approach involves integrating GENET's reinforcement learning techniques with the BBR algorithm. The objective is to use RL to predict network characteristics crucial for BBR's decision-making process, particularly in estimating bottleneck bandwidth and round-trip time. GENET's curriculum learning will be employed to train the model in various network conditions, focusing on scenarios where BBR's traditional approach might underperform. This integration aims to enhance BBR's adaptability and efficiency in diverse and dynamic network environments.

3.4 Simulation Setup and Parameters for Testing:

The evaluation of the modified BBR algorithm will be conducted through a series of simulations.

The simulation environment will be set up to mimic a range of network conditions, including varying levels of bandwidth, latency, and packet loss. Key parameters to be varied include link capacity, propagation delay, and queue size. The performance of the enhanced BBR algorithm will be compared against the standard BBR algorithm and other prevalent congestion control algorithms. Metrics for evaluation will include throughput, latency, packet loss, and fairness. This comprehensive testing will allow us to assess the effectiveness of incorporating GENET's RL techniques into the BBR congestion control algorithm.

IV. IMPLEMENTATION

The implementation phase of our research project is divided into two stages. The first stage, which has been completed, involved testing the BBR congestion control algorithm against older algorithms like Cubic and Reno. The second stage, yet to be initiated, will focus on integrating GENET's reinforcement learning framework into BBR. This section details the implementation process for both stages.

4.1 Testing BBR Against Traditional Congestion Control Algorithms:

For the initial stage of testing, we set up a controlled environment using two Virtual Machine (VM) instances on Google Cloud. This setup was chosen for its scalability and the ability to closely mimic real-world network conditions.

a. Setup and Tools:

We utilized iperf3, a widely-used network measurement tool, to generate traffic and measure throughput between the VM instances.

Linux traffic control (tc) was used to simulate network conditions with high latency and high packet loss. This allowed us to evaluate the performance of BBR in scenarios that closely resemble challenging real-world network environments.

b. Procedure:

The VM instances were configured to represent a client-server model. One VM acted as the iperf3 server, while the other operated as the client initiating traffic. We conducted a series of tests comparing BBR with traditional congestion control algorithms, namely

Cubic and Reno. These algorithms were chosen due to their widespread use and historical significance in the evolution of congestion control.
Each algorithm was tested under varying conditions of latency and packet loss to evaluate their performance in terms of throughput, latency, and fairness.

c. Observations:
Preliminary results indicated that BBR outperformed Cubic and Reno in environments with high bandwidth and low latency. However, in conditions of high latency and packet loss, the differences in performance were more nuanced and required deeper analysis.

4.2 Planned Integration of GENET into BBR:
The next phase involves the integration of GENET's reinforcement learning framework into the BBR algorithm. This stage aims to leverage GENET's curriculum learning approach to enhance BBR's adaptability and performance in diverse network conditions.

a. Integration Plan:
GENET's RL model will be trained using data collected from the VM instances under various network conditions. This will include environments where BBR's performance was suboptimal.

The training will focus on improving BBR's ability to predict and adapt to changes in bandwidth and latency, leveraging GENET's strength in identifying challenging network scenarios for targeted learning.

b. Expected Challenges:

Integrating an RL framework into an existing congestion control algorithm presents unique challenges. These include ensuring the stability of the algorithm in dynamic environments and the computational efficiency of the RL model.
Ensuring that the modified BBR algorithm remains fair when coexisting with other congestion control algorithms in the network will also be a key consideration.

c. Future Testing and Validation:
Once the integration is complete, a series of tests similar to the initial stage will be conducted. These tests will compare the performance of the modified BBR algorithm with the standard BBR and other algorithms under various network conditions.

Metrics for evaluation will remain consistent with the initial testing phase to provide a comparative analysis.

Through these stages of implementation, we aim to not only validate the superiority of BBR over traditional algorithms but also to pave the way for an advanced congestion control algorithm that leverages the latest advancements in reinforcement learning.

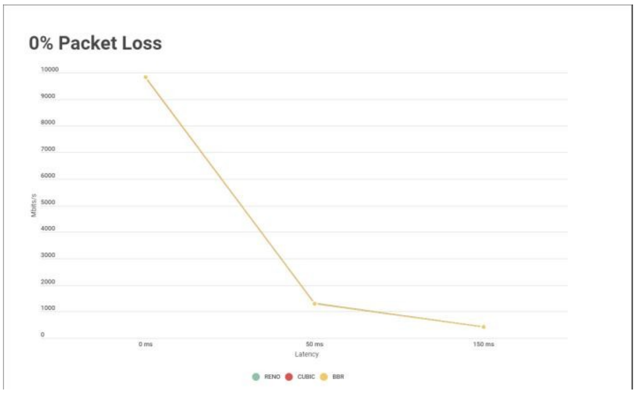

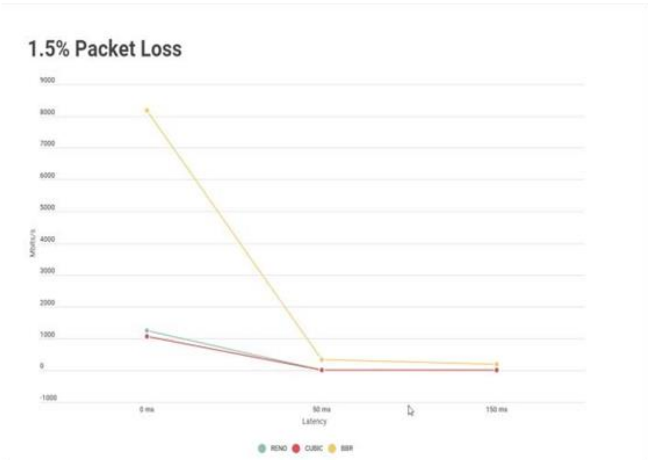

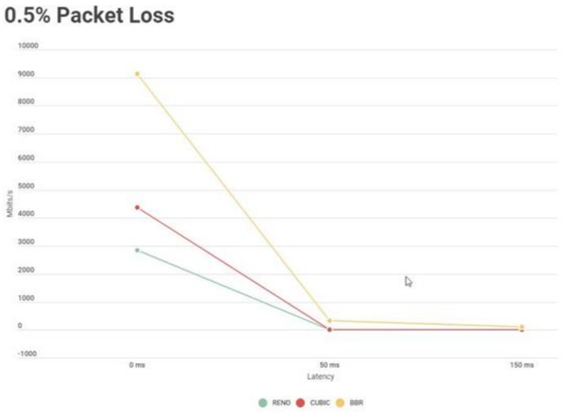

Our research project embarked on an in- depth analysis and testing of the BBR congestion control algorithm, comparing its performance against traditional algorithms such as Reno and Cubic, particularly under conditions of high packet loss. The findings from our tests have provided valuable insights into the capabilities and advantages of BBR in challenging network environments.

Through the implementation of controlled experiments using VM instances on Google Cloud, along with tools like `iperf3` and Linux traffic control, we observed that BBR consistently outperformed Reno and Cubic in scenarios characterized by high packet loss. This result aligns with BBR's design philosophy, which focuses on leveraging the bottleneck bandwidth and round-trip propagation time for congestion control, as opposed to the traditional approach of using packet loss as a primary indicator of congestion.

The superiority of BBR in these conditions suggests its potential effectiveness in modern, high- throughput networks, where managing congestion efficiently is crucial for maintaining optimal performance. However, while our tests have validated BBR's effectiveness over older algorithms in certain scenarios, the quest for an ideal congestion control mechanism that can adapt dynamically to a wide range of network conditions is ongoing.

In this context, we believe that integrating GENET's reinforcement learning framework into BBR holds significant promise. GENET, with its novel approach to training RL-based network adaptation algorithms, could potentially enhance BBR's adaptability and efficiency in diverse network environments. The curriculum learning aspect of GENET, which focuses on training models in progressively challenging scenarios, is particularly suited for addressing the complex dynamics of real- world network conditions.

However, it is important to note that this integration and its potential benefits are yet to be fully explored and tested. The implementation of GENET within BBR presents its own set of challenges, including ensuring stability and computational efficiency, and requires thorough research and testing. The future integration and testing of GENET with BBR will be a critical step in our ongoing efforts to improve congestion control algorithms.

V. CONCLUSION

In conclusion, our research reaffirms BBR's effectiveness over traditional algorithms like Reno and Cubic in high packet loss situations, and it sets the stage for further advancements in congestion control through the application of advanced reinforcement learning techniques. We anticipate that further exploration and development in this area will contribute significantly to the evolution of network congestion management, offering more robust, efficient, and adaptable solutions for the ever-changing landscape of network communications.

VI. REFERENCES

[1]. Cardwell, N., Cheng, Y., Gunn, C. S., Yeganeh, S. H., & Jacobson, V. (2016). BBR: Congestion-Based Congestion Control. Communications of the ACM, 60(2), 58-66.

[2]. Zhengxu Xia, Yajie Zhou, Francis Y. Yan, Junchen Jiang. (2022). GENET: Evolving Reinforcement Learning Algorithms for Adaptive Network Control.
[3]. Jacobson, V. (1988). Congestion avoidance and control. In Proceedings of SIGCOMM '88, Symposium on Communications Architectures and Protocols, 314-329.
[4]. Floyd, S., & Henderson, T. (1999). The NewReno Modification to TCP's Fast Recovery Algorithm. RFC 2582.
[5]. Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction (2nd ed.). MIT Press.
[6]. Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. Nature, 518, 529-533.
[7]. Ha, S., Rhee, I., & Xu, L. (2008). CUBIC: A New TCP-Friendly High-Speed TCP Variant. ACM SIGOPS Operating Systems Review, 42(5), 64-74.

Cite this article as :

Dr. Kunwar Asif, Hemanjali Kadali, Sathwik Varma Mudduluri, Pawan Sai Krishna Reddy Kerelly, "Plausibility of BBR as CUBIC'S Replacement and proposed improvement to BBR using GENET", International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT), ISSN : 2456-3307, Volume 9, Issue 6, pp.303-309, November-December-2023. Available at doi :
https://doi.org/10.32628/CSEIT2390660
Journal URL : https://ijsrcseit.com/CSEIT2390660