Poster Abstract: Configuration Tuning for Distributed IoT Message Systems Using Deep Reinforcement Learning

Zhuangwei Kang*
Vanderbilt University
Nashville, Tennessee
zhuangwei.kang@vanderbilt.edu

Yogesh D. Barve
Vanderbilt University
Nashville, Tennessee
yogesh.d.barve@vanderbilt.edu

Shunxing Bao
Vanderbilt University
Nashville, Tennessee
shunxing.bao@vanderbilt.edu

Abhishek Dubey
Vanderbilt University
Nashville, Tennessee
abhishek.dubey@vanderbilt.edu

Aniruddha Gokhale
Vanderbilt University
Nashville, Tennessee
a.gokhale@vanderbilt.edu

ABSTRACT

Distributed messaging systems (DMSs) are often equipped with a large number of configurable parameters that enable users to define application run-time behaviors and information dissemination rules. However, the resulting high-dimensional configuration space makes it difficult for users to determine the best configuration that can maximize application QoS under a variety of operational conditions. This poster introduces a novel, automatic knob tuning framework called DMSConfig. DMSConfig explores the configuration space by interacting with a data-driven environment prediction model(a DMS simulator), which eliminates the prohibitive cost of conducting online interactions with the production environment. DMSConfig employs the deep deterministic policy gradient (DDPG) method and a custom reward mechanism to learn and make configuration decisions based on predicted DMS states and performance. Our initial experimental results, conducted on a single-broker Kafka cluster, show that DMSConfig significantly outperforms the default configuration and has better adaptability to CPU and bandwidth-limited environments. We also confirm that DMSConfig produces fewer violations of latency constraints than three prevalent parameter tuning tools.

CCS CONCEPTS

• **Software and its engineering → Software configuration management and version control systems**; • **Computing methodologies → Policy iteration**.

KEYWORDS

Publish/Subscribe Middleware, System Configuration, Policy-based RL Algorithm

ACM Reference Format:
Zhuangwei Kang, Yogesh D. Barve, Shunxing Bao, Abhishek Dubey, and Aniruddha Gokhale. 2021. Poster Abstract: Configuration Tuning for Distributed IoT Message Systems Using Deep Reinforcement Learning. In *International Conference on Internet-of-Things Design and Implementation (IoTDI '21), May 18–21, 2021, Charlottesvle, VA, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3450268.3453517

IoTDI '21, May 18–21, 2021, Charlottesvle, VA, USA

ACM ISBN 978-1-4503-8354-7/21/05.
https://doi.org/10.1145/3450268.3453517

1 INTRODUCTION

A variety of IoT application domains, such as smart cities and smart grids, employ DMS as the middleware for data transmission through which messages can be produced, disseminated and consumed asynchronously. To ensure flexibility in a wide range of deployment scenarios, system topologies and runtime specifications, industrial-strength DMSs provide users with a set of continuous and discrete configurable parameters that have different data types (e.g., numeric, boolean, categorical) and value ranges, which together result a hybrid, multidimensional configuration space. These parameters control application runtime behaviors and resource allocation strategies, resulting in different variations of application performance measured across different metrics, such as throughput, latency, CPU utilization, etc.

Making prudent configuration decisions is challenging because the scale of the searching space boosts exponentially as the quantity of tunable parameters increases. It also requires significant domain knowledge and in-depth understanding of the impact of each parameter on application performance as well as their unseen interactions, which is difficult even for experts, not to mention common users. The default configurations provided by the software vendors are usually suboptimal, and utilizing naïve exhaustive search methods to find appropriate configurations are laborious, time-consuming, non-scalable, and likely be suboptimal due to the continuous nature of the configuration space (most parameters are numeric). Hence, this project proposes a Deep Reinforcement Learning (DRL)-based configuration recommendation system, called DMSConfig. We aim to optimize the publisher-side throughput of DMS applications while meeting latency constraints, which satisfies the demands of practical IoT streaming applications, such as online smart grid analytics, that usually have stringent requirements on both throughput and response time. This project uses Kafka, a popular event-processing framework used at the data analytics layer of IoT systems, as an example to validate our approach, but DMSConfig can be adapted to other DMSs since its system components are fully decoupled.