

iQCAR: Inter-Query Contention Analyzer

Prajakta Kalmegh, Shivnath Babu, Sudeepa Roy
Department of Computer Science, Duke University
308 Research Drive Duke Box 90129, Durham, NC 27708-0129
{pkalmegh,shivnath,sudeepa}@cs.duke.edu

ABSTRACT

There are many approaches in use today to either prevent or minimize the impact of inter-query interactions on a shared cluster. Preventive measures often provide query execution isolation at the resource allocation level to guarantee a predictable query performance. Despite these measures, performance issues due to concurrent executions of mixed workloads are a common problem in large scale data processing systems. As a result, answering questions like *who is causing my query to slowdown* is important to diagnose resource conflicts in a multi-tenant environment for accurate blame attribution. However, accurate analysis of resource contention is challenging owing to a complex cause-effect relationship between resource utilization and runtime of concurrent queries (see Figure 1). For example, when some tasks get delayed because of a high demand for a particular resource (e.g. if they are blocked on CPU), they hold on to other resources (e.g. memory) as well, thus causing contention for other concurrently running queries on the held resources. Based on our user-study experience, this process is non-trivial and tedious, and involves hours of manually debugging through a cycle of query interactions.

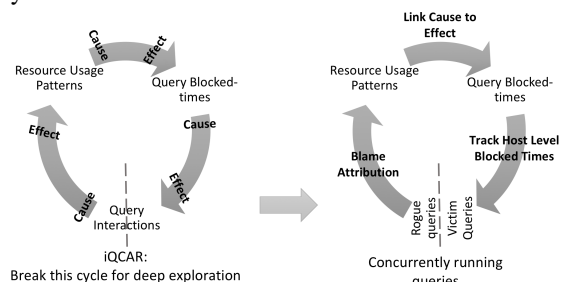


Figure 1: iQCAR aims to break the query interactions cycle into a multi-level blame distribution graph.

This work was supported in part by NSF awards IIS-1423124, IIS-1552538, and IIS-1703431, and NIH Award 1R01EB025021-01.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SoCC '18, October 11–13, 2018, Carlsbad, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6011-1/18/10.

<https://doi.org/10.1145/3267809.3275473>

In this research, we develop iQCAR- an inter Query Contention Analyzer tool that explores the resource interference cycle among concurrent queries and detects high-impacting queries systematically. Our contributions are as follows:

- **Blame Attribution:** We develop a metric called *Resource Acquire Time Penalty* (RATP) to capture the blocked-time distributions of a query for every resource. Using RATP of a task as a basis, we develop a metric to attribute blame to concurrent queries in causing contention to a victim query.
- **Explanations:** Our multi-level Directed Acyclic Graph (DAG), called iQC-Graph, enables distribution of blame values to concurrently running queries, their stages and tasks for different hosts and resources, and generates explanations for the slowdown of a query.
- **End-to-end system for deep exploration and blame analysis:** Using RATP and iQC-Graph, the web-based front-end of iQCAR [1, 2] helps administrators of cluster computing systems detect (i) hot resources, (ii) slow nodes, (iii) high impact *causing* queries, and (iv) high impact *receiving* victim queries.
- **Rules for Cluster Scheduler:** We use the top-*k* explanations in generating alternative query placement rules that can be applied by a cluster scheduler in an online execution.

KEYWORDS

Performance evaluation; contention analysis; blame attribution; resource bottleneck; cluster computing systems

ACM Reference Format:

Prajakta Kalmegh, Shivnath Babu, Sudeepa Roy, Department of Computer Science, Duke University, 308 Research Drive Duke Box 90129, Durham, NC 27708-0129, {pkalmegh,shivnath,sudeepa}@cs.duke.edu, . 2018. iQCAR: Inter-Query Contention Analyzer. In *ACM Symposium on Cloud Computing (SoCC '18)*, October 11–13, 2018, Carlsbad, CA, USA. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3267809.3275473>

REFERENCES

- [1] iQCAR: Inter-Query Contention Analyzer Tool. <https://www.cs.duke.edu/~pkalmegh/iqcar.html>.
- [2] P. Kalmegh, H. Lundberg, F. Xu, S. Babu, and S. Roy. iqcar: A demonstration of an inter-query contention analyzer for cluster computing frameworks. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD '18*, pages 1721–1724, New York, NY, USA, 2018. ACM.