Interference Analysis of Co-located Container Workloads: A Perspective from Hardware Performance Counters

Wen-Yan Chen[1], Ke-Jiang Ye[1,*], Cheng-Zhi Lu[1] , Dong-Dai Zhou[2]， Cheng-Zhong Xu[3]

[1] Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen 518055, China

[2] College of Information Science and Technology, Northeast Normal University, Changchun 130117, China

[3] Faculty of Science and Technology, University of Macau, Macau 999078, China

Chen WY, Ye KJ, Lu CZ et al. Interference analysis of co-located container workloads: A perspective from hardware performance counters. JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY 35(2): 412–417 Mar. 2020. DOI 10.1007/s11390-020-9707-y

Problem & Contribution

Research Problem

- Interference is an important issue needs to be solved in containers, as the container isolation is much worse than VM isolation. Due to the increasing population of container and complexity of workloads, it's very necessary to study this problem.

- Understanding the system performance and characteristics is very necessary to solve this problem and avoid resource preemption caused by uneven resource allocation strategies.

Main Contribution

- We analyzed the characteristics of different workload patterns from the hardware event level, and then speculated the causes of performance degradation when the workloads are co-located.

- According to the analysis results, we provided some valuable suggestions of recommended and unrecommended co-location workload patterns for datacenter administrators.

Experiments

- We select seven representative workloads including data analysis workloads and online service workloads from BigDataBench and CloudSuite benchmarks (Table 1).

Table 1. Representative Workloads used in This Paper

Workload Name	Domain	Application Type
WordCount (WC)	SE, SN, EC	Offline Analysis
CollaborativeFiltering (CF)	SE, SN, EC	Offline Analysis
Kmeans (K)	SE, SN, EC	Offline Analysis
ConnectedComponent (CC)	SE, SN	Graph Analysis
PageRank (PR)	SE	Graph Analysis
Data Caching (DC)	SE, SN	Online Service
Media Streaming (MS)	MS	Online Service

Note: *SE: search engine, SN: social network, EC: electronic commerce, MS: media streaming.*

- We select valuable information by pruning metrics from nearly 200 metrics (Table 2).

Table 2. Core Metrics used in This Paper

Metric	Description	Event Type
IPC	Instructions per cycle	Hardware event
BPM	Branch prediction misses	Hardware event
CS	Context switches	Software event
L1DCLM	L1 data cache load misses	Hardware cache event
LLCM	Last level cache misses	Hardware cache event
dTLBLM	dTLB load misses	Hardware cache event

- We co-locate workloads together to observe the performance interference phenomenon (Table 3).

Table 3. Matrix Table of Workload Co-location Patterns

	CC	WC	PR	CF	K	DC	MS
CC		1	1	1	1	1	1
WC	1		1	1	1	1	1
PR	1	1		1	1	1	1
CF	1	1	1		1	1	1
K	1	1	1	1		1	1

Note: CC, WC, PR, CF, K, DC and MS represent the abbreviations of the workload names. The matrix table represents co-location patterns of workloads, 1 means two workloads are co-located, such as CC-WC, CC-PR.

Results

- The experiment shows that:
  - **a certain typical hardware event has little impact on the running time of workloads**
  - **some types of workloads can be co-located to improve the resource utilization**

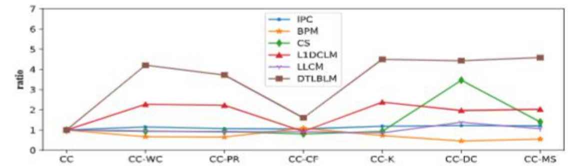

Fig.1. Metrics changes of different workload patterns of CC-X

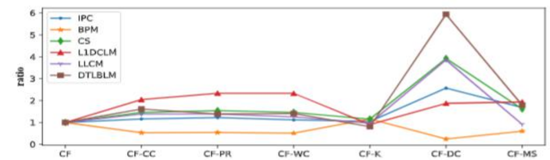

Fig.2. Metrics changes of different workload patterns of CF-X.

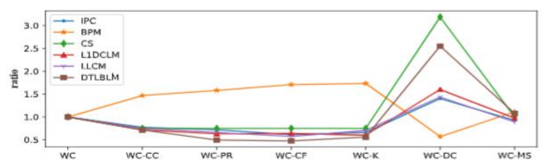

Fig.3. Metrics changes of different workload patterns of WC-X.

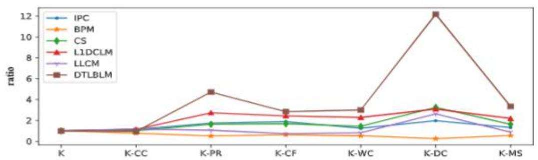

Fig.4. Metrics changes of different workload patterns of K-X.

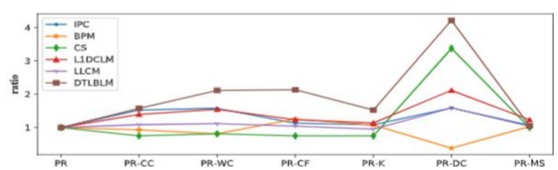

Fig.5. Metrics changes of different workload patterns of PR-X.

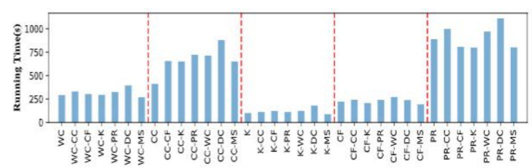

Fig.7. Running time of each Data Analysis Workload (The x-axis labels such as WC-CC represents the running time of WC when it is co-located with CC, others are the same to WC-CC).

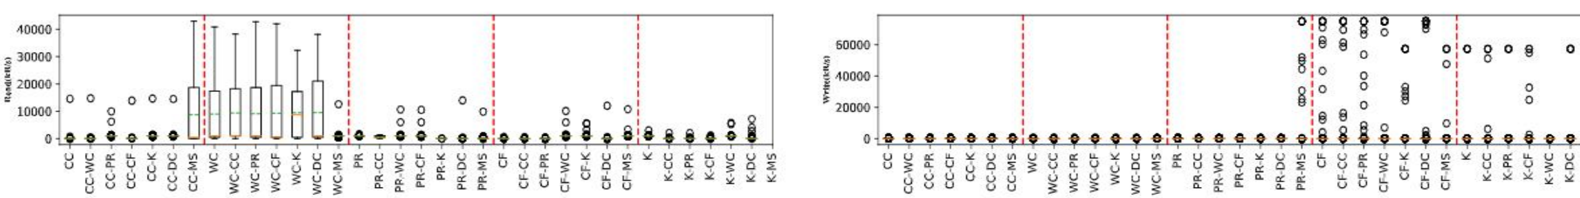

(a) (b)

Fig.6. The box-plots of Reading/Writing rate of each Data Analysis workload when it is co-located with other workloads. (a) Reading rate. (b) Writing rate. (The orange line represents median value, and the green line represents the average value. The circles simply point outside of the wide [(Q1-1.5 IQR), (Q3+1.5 IQR)] margin below. Q1: First Quartile, IQR: Interquartile Range.)

Conclusion

- This paper studies the **interference issue** in different co-location patterns by using diverse workloads (Offline Data Analysis and Online Service) in containers.

- Through the analysis of co-located patterns with hardware level information, we provide recommended and unrecommended workload patterns, and provide **valuable deployment suggestions** for datacenter administrators in order to optimize system performance.

- In the future, we plan to use the co-location pattern knowledge to **guide the workload deployment and scheduling** in the large-scale container-based cloud environment.