Characterizing Performance Variation of Genomic Data Analysis Workflows on the Public Cloud

David Perez, Ling-Hung Hong, Sonia Xu, KYY1
Ka Yee Yeung, Wes Lloyd
daperez@uw.edu, wlloyd@uw.edu

August 17-24, 2020

School of Engineering and Technology
University Of Washington, Tacoma

CBDCOM 2020: IEEE International Conference on Cloud and Big Data Computing

1

Outline

- Background
- Research Questions
- Use Case
- Experimental Implementation
- Experimental Results
- Summary
- Conclusions

Slide 1

KYY1 I'd suggest to add: Correspondence Email: wlloyd@uw.edu
Ka Yee Yeung, 7/27/2020

CPU Heterogeneity

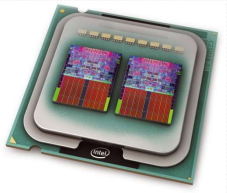

- Public cloud providers offer distinct **VM types** to simplify resource allocation to users
- **VM types:**
  - Have distinct configurations: (e.g. # of virtual CPUs (vCPUs), memory/storage capacity, and network bandwidth)

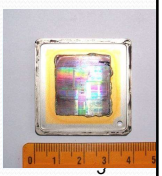

Resource Contention

- Resource Contention is when there is a competition over shared resources on a shared server

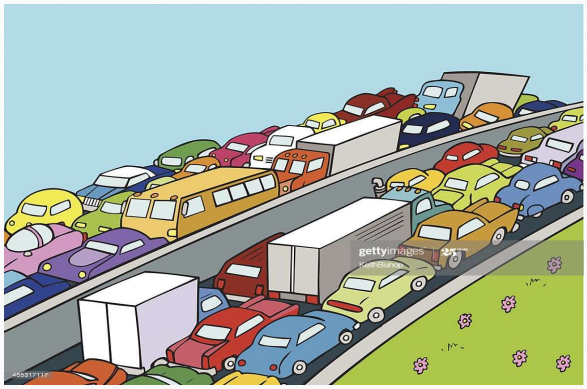

Provisioning Variation

- Provisioning variation is the random nature of VM placement across physical servers that occurs when cloud providers load balance VM launch requests.
- Where these VMs are hosted on public clouds is abstracted and not easily inferable in real time.

Outline

- Background
- Research Questions
- Use Case
- Experimental Implementation
- Experimental Results
- Summary
- Conclusions

Research Questions

KYY3

RQ1: What is the performance variation of running genomics data analytical tasks on the public cloud?

How much do factors such as provisioning variation, CPU heterogeneity, and resource contention contribute to performance variation?

RQ2: What relationships exist between Linux resource utilization metrics (CPU, memory, disk, and network) and workflow runtime?

Outline

- Background
- Research Questions
- Use Case
- Experimental Implementation
- Experimental Results
- Summary
- Conclusions

Slide 7

KYY3 I'd propose to use "animation" to show RQ1 first, before showing RQ2.
Ka Yee Yeung, 7/27/2020

Use Case: UMI RNA Sequencing Workflow (Xiong, Yuguang, et al)

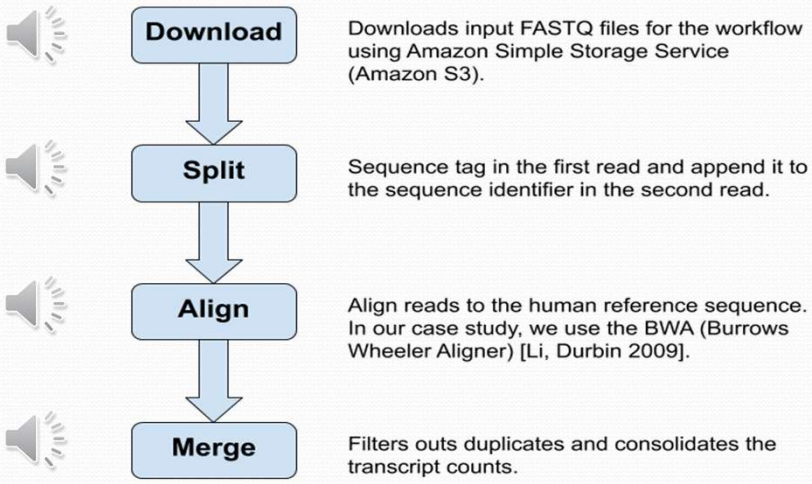

https://www.nature.com/articles/s41598-017-14892-x.pdf

Outline

- Background
- Research Questions
- Use Case
- Experimental Implementation
- Experimental Results
- Summary
- Conclusions

WJL7 DP1

Container Profiler KYY10

The Container Profiler measures and records resource utilization of any containerized task capturing over 50+ Linux system metrics to characterize CPU, memory, disk, and network utilization at the VM, container, and process levels. KYY9

These metrics are important as they can help identify what system resources your workflow is consuming the most.

WJL6

Controlling provisioning variation with AWS EC2 Placement Groups

- **Standard Placement:** No strategy – standard VM launch
- **Spread Placement:** Instances placed on distinct servers located on different server racks.
- **Cluster Placement: I**nstances placed packed together inside an Availability Zone

AWS. 2020. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html Last accessed July, 2020.

Slide 11

WJL7 probably should add a bullet to say that we used the container profiler for our measurements
Wes J. Lloyd, 7/22/2020

DP1 Can add graphic images from the gigascience container profiler paper
David Perez, 7/24/2020

KYY9 A typical bioinformatics audience may not know what Linux metrics you are talking about. I'd suggest to show some examples (those that you will show in later results).
Ka Yee Yeung, 7/27/2020

KYY10 I think the focus on this slide should be what metrics we are collecting and WHY. The structure of the code is not as important.
Ka Yee Yeung, 7/27/2020

Slide 12

WJL6 May consider cutting out partition placement since we don't use it in any experiments
Wes J. Lloyd, 7/22/2020

Experimental Setup

Using AWS EC2, we provisioned 30 x ec2 c5.2xlarge instances, 10 VMs for each placement strategy:

WJL4

	Total VMS	Standard	Cluster	Spread
8124M	16	4	4	8
8275CL	14	6	6	2

WJL18
WJL19

c5.2xlarge Heterogeneous CPU comparison

	Intel Xeon(R) Platinum 8124M CPU @ 3.00 GHZ	Intel Xeon(R) Platinum 8275CL @ 3.00 GHZ
EC2 Instance Type	c5.2xlarge	c5.2xlarge
Family/microns/yr	Skylake/14nm/2017	Cascade Lake/14nm/2019
Virtual CPU cores/host	72	96
Physical CPU cores/host	36	48
Burst clock MHz (Single/all)	3400/3500	3600/3900
L1 Cache (Per core)	64K (½ data, ½ instruction)	64k (½ data, ½ instruction)
L2 Cache (Per core)	1024K	1024K
L3 Cache (Per core)	1375K	1525K
Total Occurrences:	53%	47%
Standard Placement	13%	20%
Cluster Placement	13%	20%
Spread Placement	27%	7%

https://en.wikipedia.org/wiki/List_of_Intel_Skylake-based_Xeon_microprocessors#Xeon_Platinum_8124M
https://en.wikipedia.org/wiki/List_of_Intel_Cascade_Lake-based_Xeon_microprocessors#Xeon_Platinum_8275CL

Slide 13

WJL4 Suggest removing c5a.2xlarge and saving this for the ParBio presentation.

Can save and extend these slides for ParBio.
Wes J. Lloyd, 7/22/2020

Slide 14

WJL18 this slide reports % frequency of instances, I suggest adding the numbers too because it provides scale to the work. How many VMs for each
Wes J. Lloyd, 7/22/2020

WJL19 Suggest adding an animation frame the highlights each CPU type as you talk about it
Wes J. Lloyd, 7/22/2020

Outline

- Background
- Research Questions
- Use Case
- Experimental Implementation
- Experimental Results
- Summary
- Conclusions

WJL19

RQ-1: Performance Variation

WJL5

- What is the performance variation of running genomics data analytical tasks on the public cloud?

Slide 16

WJL5 Please select a clearer (sans-serif) font for the table.

This is any font without the extending features called "serifs" at the end of strokes.

Times new roman is a serif font
Arial is a sans-serif font
Wes J. Lloyd, 7/22/2020

WJL18 this slide reports % frequency of instances, I suggest adding the numbers too because it provides scale to the work. How many VMs for each
Wes J. Lloyd, 7/22/2020

WJL19 Suggest adding an animation frame the highlights each CPU type as you talk about it
Wes J. Lloyd, 7/22/2020

Performance Variation: Standard Placement

CPU runtime variation - c5.2xlarge, Standard placement:

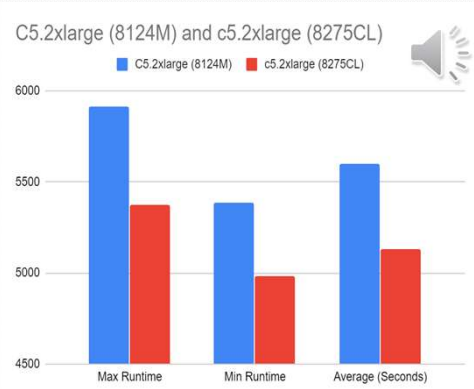

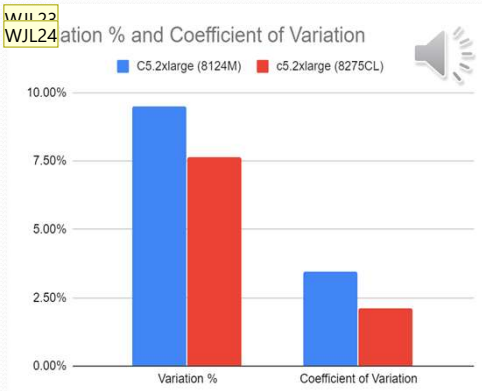

rformance Variation: Spread Placement

CPU runtime variation - c5.2xlarge, Spread placement:

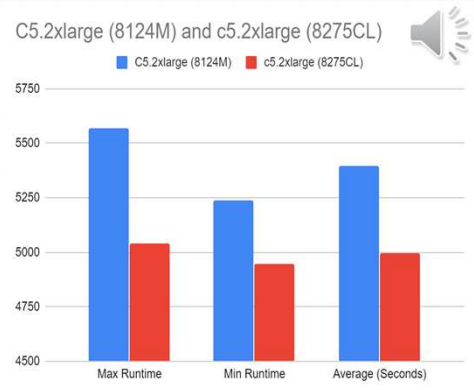

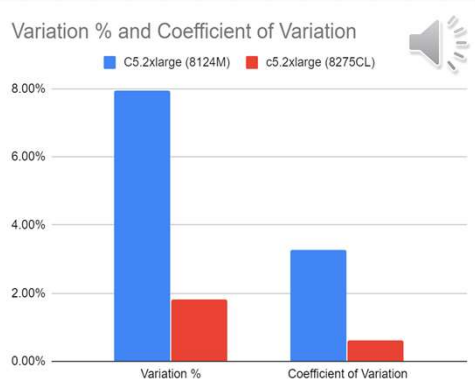

Slide 17

WJL9 needs a sans-serif font
Wes J. Lloyd, 7/22/2020

WJL14 probably needs to be replaced with graphs or at the very least small tables

We should not throw a slide full of numbers at the audience. At the very least, can animate a frame to appear around the sections of the table you'll talk about.
Wes J. Lloyd, 7/22/2020

WJL16 The presentation of the results should flow more. It could convery the experimental steps/process and the results. Can break this slide down into separate slides for standard, cluster, spread. Can add highlight (BOLD/color) for min and max. Can add a couple observation bullet points from each set of runs.
Wes J. Lloyd, 7/22/2020

WJL23 IS THIS the whole workflow or just alignment?
Wes J. Lloyd, 7/27/2020

WJL24 WHERE ARE THESE NUMBERS FROM THEY DON'T MATCH THE PAPER
Wes J. Lloyd, 7/27/2020

Slide 18

WJL9 needs a sans-serif font
Wes J. Lloyd, 7/22/2020

WJL14 probably needs to be replaced with graphs or at the very least small tables

We should not throw a slide full of numbers at the audience. At the very least, can animate a frame to appear around the sections of the table you'll talk about.
Wes J. Lloyd, 7/22/2020

WJL16 The presentation of the results should flow more. It could convery the experimental steps/process and the results. Can break this slide down into separate slides for standard, cluster, spread. Can add highlight (BOLD/color) for min and max. Can add a couple observation bullet points from each set of runs.
Wes J. Lloyd, 7/22/2020

WJL25 WHERE ARE THESE NUMBERS FROM?
THEY DON'T MATCH THE PAPER
Wes J. Lloyd, 7/27/2020

WJL26 They look like they are just the alignment
Wes J. Lloyd, 7/27/2020

Performance Variation: Cluster Placement

CPU runtime variation - c5.2xlarge, Cluster placement:

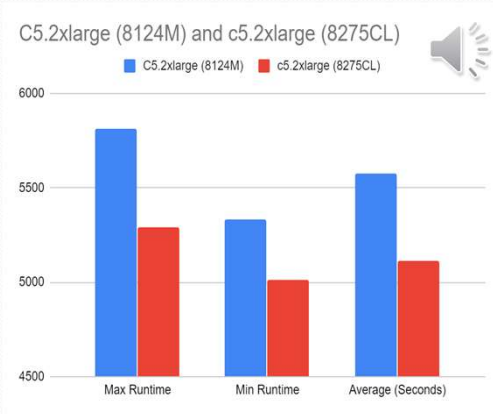

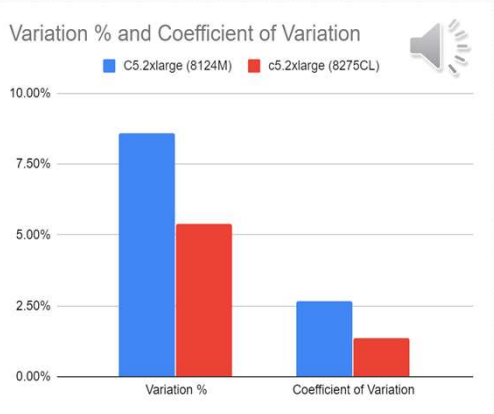

RQ-2: Inferring performance from resource utilization metrics

What relationships exist between Linux resource utilization metrics (CPU, memory, disk, and network) and workflow runtime?

Slide 19

WJL9 needs a sans-serif font
Wes J. Lloyd, 7/22/2020

WJL14 probably needs to be replaced with graphs or at the very least small tables

We should not throw a slide full of numbers at the audience. At the very least, can animate a frame to appear around the sections of the table you'll talk about.
Wes J. Lloyd, 7/22/2020

WJL16 The presentation of the results should flow more. It could convery the experimental steps/process and the results. Can break this slide down into separate slides for standard, cluster, spread. Can add highlight (BOLD/color) for min and max. Can add a couple observation bullet points from each set of runs.
Wes J. Lloyd, 7/22/2020

WJL10

RQ-2: Inferring performance from resource utilization metrics

Resource utilization heatmap using collected data from the Container Profiler with clustered rows.

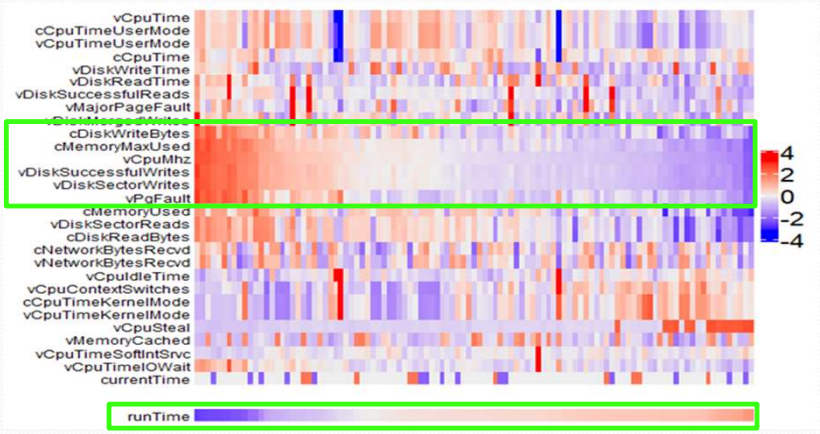

Outline

- Background
- Research Questions
- Use Case
- Experimental Implementation
- Experimental Results
- Summary
- Conclusions

Slide 21

WJL10 I would create a frame and animate it to highlight certain parts of the figure that you will talk about.

A frame can be created as a rectangle with no fill that has a thick outer line of point 3 or 4. Can use a pleasing green color for the highlight frame.

Use the appear/disappear animation to add and remove it to different areas of the figure that you'll refer to in the talk.

Wes J. Lloyd, 7/22/2020

WJL15

Summary

- RQ-1 Performance variation:

  Performance variance of long running compute-bound tasks on were found to be as high as 18.9% and as low as 12.5% using the same instance type (c5.2xlarge).

- RQ-2 Metric relationships with performance:

  A subset of metrics gathered by the Container profiler have been shown to exhibit a strong inverse relationship with runtime.

Outline

- Background
- Research Questions
- Use Case
- Experimental Implementation
- Experimental Results
- Summary
- Conclusions

Slide 23

WJL15 Add a QUESTIONS slide at the end.

> Can give your name, email, and well as my name and email so that people email us any questions.
> Wes J. Lloyd, 7/22/2020

Conclusions

From RQ-1 we determined when running our genomics data analysis workflow that:

- Spread is fastest and most consistent, with the fastest possible runtime.
- Standard is the slowest, least consistent, with the worst possible runtimes.
- Cluster is middle of the pack.

From RQ-2 we determined when running our genomics data analysis workflow that:

- cDiskWriteBytes, cMemoryMaxUsed, vCpuMhz, vDiskSuccessfulWrites, vDiskSectorWrites, vPgFaults have an inverse relationship to runtime.
- For future work we can use these metrics as candidates for categorizing whether a VM is slow, typical or fast.

THANK YOU FOR WATCHING

- Questions or Comments?
- Please Email:
- daperez@uw.edu or wlloyd@uw.edu