Towards Analyzing Computational Costs of Spark for SARS-CoV-2 Sequence Comparisons on a Commercial Cloud

Alan L. Nunes
Alba Cristina M. A. de Melo
Cristina Boeres
Daniel de Oliveira
Lúcia M. A. Drummond

Instituto de Computação

UnB

CAPES Programa de Internacionalização CAPES/PrInt

CNPq aws

WSCAD 2021 - XXII Symposium on High Performance Computer Systems

(Brazilian Workshop in Conjunction SBAC PAD)

Towards Analyzing Computational Costs of Spark for SARS-CoV-2 Sequence Comparisons on a Commercial Cloud

Alan L. Nunes
Alba Cristina M. A. de Melo
Cristina Boeres
Daniel de Oliveira
Lúcia M. A. Drummond

Instituto de Computação

UnB

CAPES Programa de Internacionalização CAPES/PrInt

CNPq aws

Motivation

Motivation

- Applications that generate a large amount of data.
    - *How to process/query this volume of data and extract useful knowledge from it in a timely manner?*

Motivation

- Relevant case: Covid-19
  - the comparison of SARS-CoV-2 biological sequences are crucial to understand the behavior of this disease

  - ≈ 5 millions SARS-CoV-2 sequences available for general use in public genomic databases - NCBI and GISAID

Motivation

- Traditional data management solutions, such as Relational Database Management Systems (RDBMS), do not scale for this volume of (commonly heterogeneous) data.
    - Big Data frameworks, e.g., **Apache Spark** has been widely adopted due to its high scalability.
- Cloud environments offers several advantages compared to dedicated infrastructures, like: rapid *resources provisioning and reduction of operational costs*.

Motivation

- Amazon Elastic Compute Cloud (EC2) has distinct virtual machines (VMs) price markets:
  - **On-Demand VMs** – fixed financial cost per time unit of use; not revoked by the provider;
  - **Spot VMs** – lower prices (up to 90% cheaper); can be revoked by the cloud provider at any time.
    - nice option when either the application can handle failures or the application is executed within a framework that provides a mechanism for recovering in case of a **VM revocation**.
      - *Apache Spark offers fault tolerance mechanisms, in case of Worker/Executor failures*

Apache Spark's Cluster Mode Overview

Apache Spark's Cluster Mode Overview

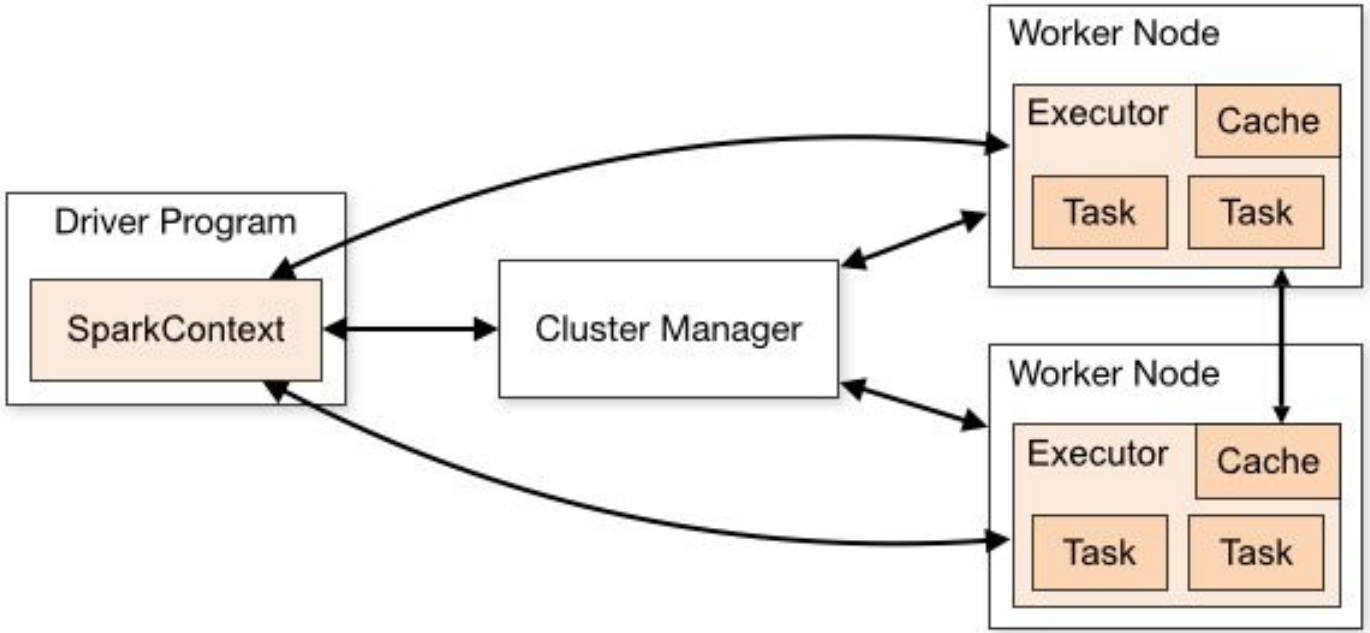

https://spark.apache.org/docs/latest/cluster-overview.html

Diff Sequences Spark's Algorithm Overview

Diff Sequences Spark's Algorithm Overview

N biological sequences to compare

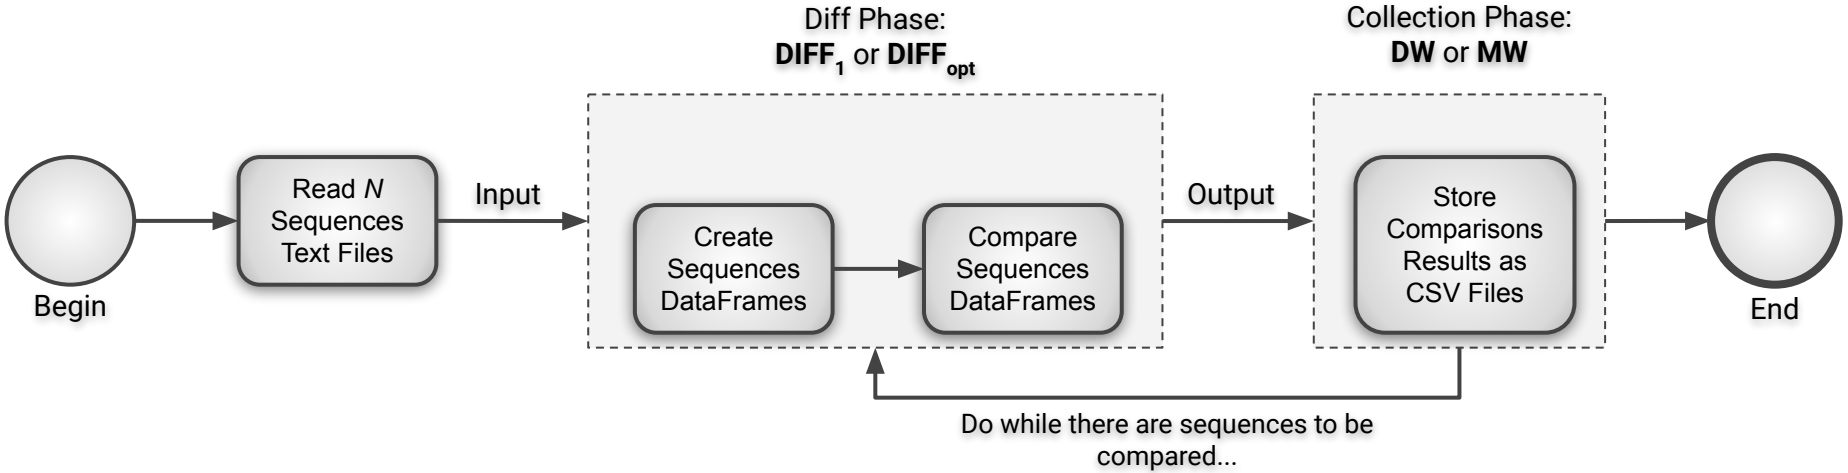

Diff Sequences Spark's Algorithm Overview

N biological sequences to compare

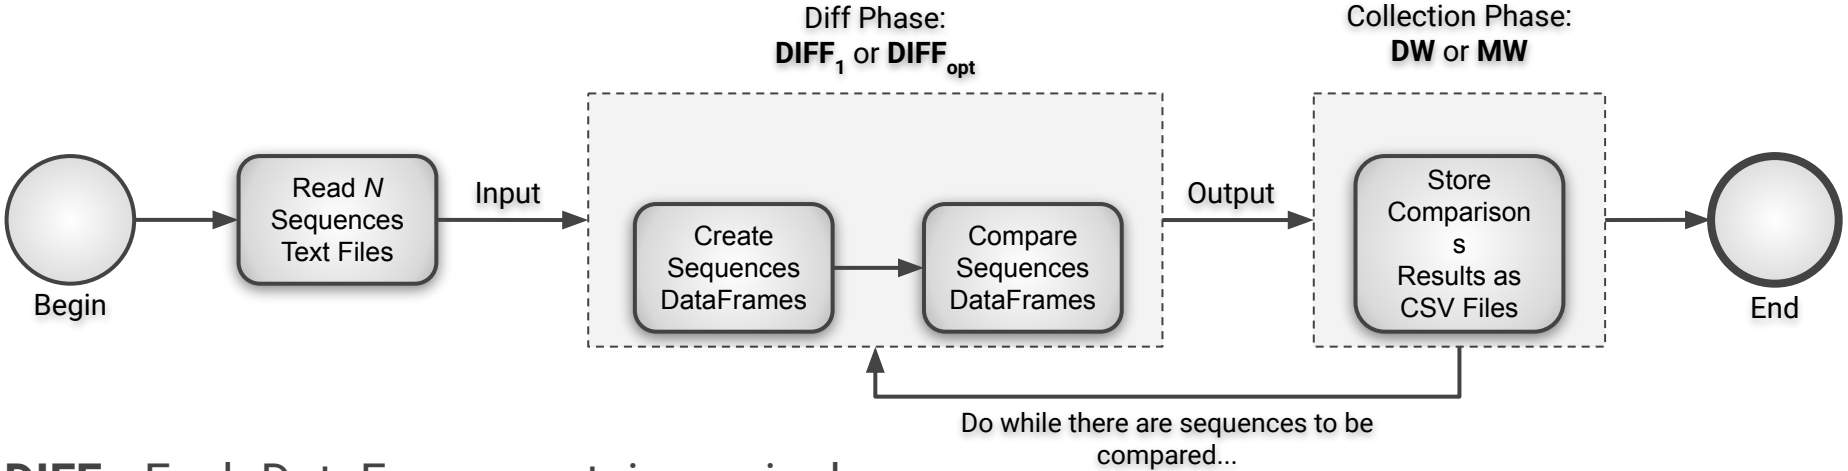

DIFF$_1$: Each DataFrame contains a single sequence.

DIFF$_{opt}$: Each DataFrame contains [1, max$_D$] sequence(s), where $1 \leq \max_D < N$.

Diff Sequences Spark's Algorithm Overview

N biological sequences to compare

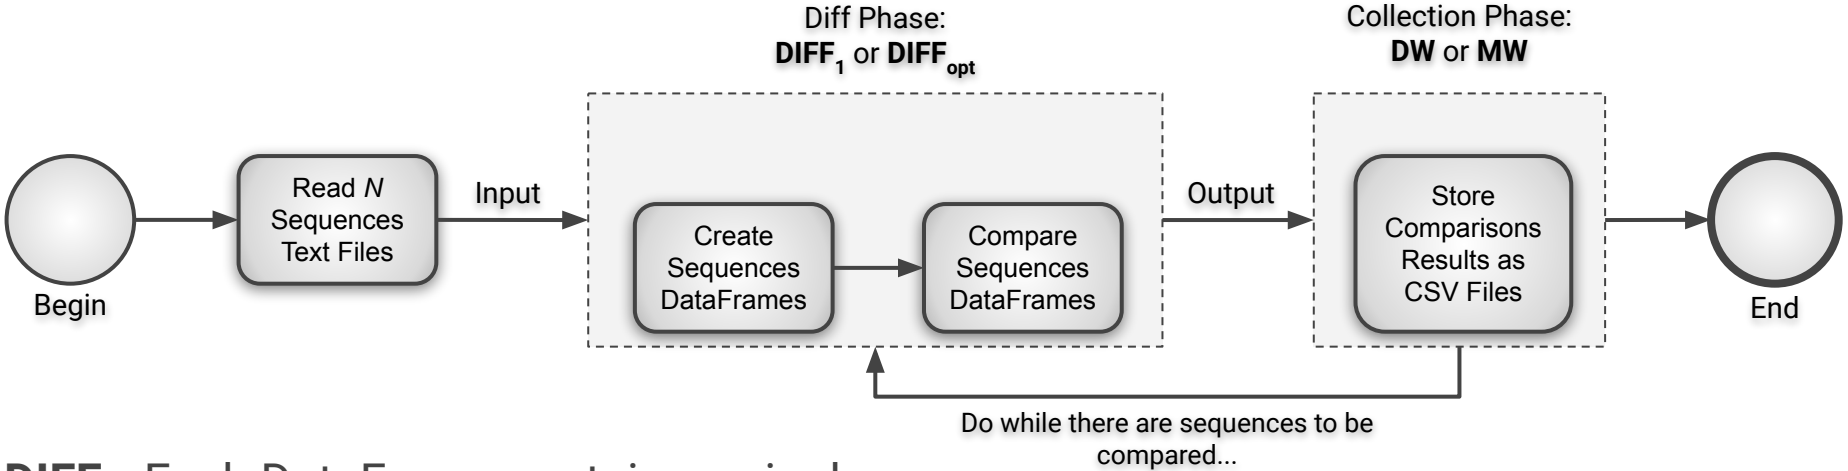

$DIFF_1$: Each DataFrame contains a single sequence.

$DIFF_{opt}$: Each DataFrame contains $[1, max_D]$ sequence(s), where $1 \leq max_D < N$.

DW: Each Executor writes its partial comparisons results to a local files
MW: One Executor merges and writes all partial results on a single file

Example

Example

Input:

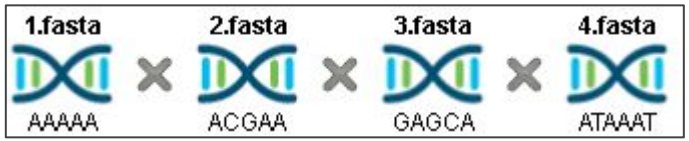

Example

Input:

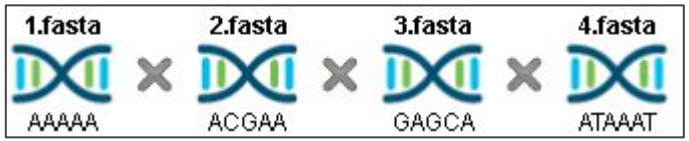

***DIFF*$_1$'s output (*N* = 4):**

Example

Input:

***DIFF*$_1$'s output (*N* = 4):**

Index	Seq_1.fasta	Seq_2.fasta
1	A	C
2	A	G

Index	Seq_1.fasta	Seq_3.fasta
0	A	G
2	A	G
3	A	C

Index	Seq_1.fasta	Seq_4.fasta
1	A	T

Index	Seq_2.fasta	Seq_3.fasta
0	A	G
1	C	A
3	A	C

Index	Seq_2.fasta	Seq_4.fasta
1	C	T
2	G	A

Index	Seq_3.fasta	Seq_4.fasta
0	G	A
1	A	T
2	G	A
3	C	A

Example

Input:

***DIFF$_{opt}$*'s output (*N* = 4, *max$_D$* = 3):**

Example

Input:

***DIFF$_{opt}$*'s output (*N* = 4, *max$_D$* = 3):**

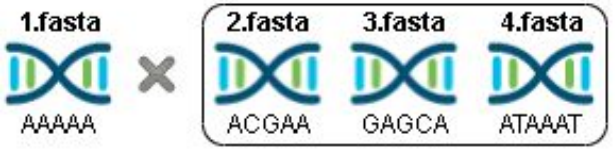

Index	Seq_1.fasta	Seq_2.fasta	Seq_3.fasta	Seq_4.fasta
0	A	=	G	=
1	A	C	=	T
2	A	G	G	=
3	A	=	C	=

Index	Seq_2.fasta	Seq_3.fasta	Seq_4.fasta
0	A	G	=
1	C	A	T
2	G	=	A
3	A	C	=

Index	Seq_3.fasta	Seq_4.fasta
0	G	A
1	A	T
2	G	A
3	C	A

Preliminary Experiments on Amazon EC2

Preliminary Experiments on Amazon EC2 – Execution Settings

Execution Settings

- **One** On-Demand Master Node (Driver Program)
- **8** Spot Worker Nodes
- Input (*N*): **{2, 4, 8, 16, 32, 64}** South America SARS-CoV-2 Sequences
- DataFrames' number of partitions: **{*Auto*, *Customized*}**
- Diff Phase: **{$DIFF_1$, $DIFF_{opt}$}**
- max_D: *N* - 1, when Diff Phase = $DIFF_{opt}$
- Collection Phase: **{*DW*, *MW*}**

Preliminary Experiments on Amazon EC2 – VM Selection

Family	Instance Name	Number of vCPUs	Memory (GiB)	Storage Type	Network Speed (Gbps)	Cost per Hour (USD)	
						On-Demand	Spot
General Purpose	t2.medium	2	4	EBS	0.3	0.0464	0.0139

Preliminary Experiments on EC2 – Average Times & Costs

Experiments Round	*N* SARS-CoV-2 Sequences	Average Diff Phase Time (Minutes)		Average Collection Phase Time (Minutes)		Average Execution Time (Minutes)		Average Execution Cost (USD)	
		$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$
FR_{PE} *Auto + DW*	2	0.0022	0.0053	0.2815	0.2893	0.4256	0.4362	0.0011	0.0011
	4	0.0045	0.0083	0.9854	0.5825	1.1537	0.7530	0.0030	0.0020
	8	0.0105	0.0196	2.8430	1.1275	3.0950	1.3454	0.0081	0.0035
	16	0.0267	0.0468	8.1126	1.9528	8.6050	2.2651	0.0226	0.0059
	32	0.0762	0.1850	29.6007	3.5592	30.9374	4.2317	0.0813	0.0111
	64	0.3822	0.5770	123.6268	5.6632	130.7519	9.2003	0.3434	0.0242

Preliminary Experiments on EC2 – Average Times & Costs

Experiments Round	*N* SARS-CoV-2 Sequences	Average Diff Phase Time (Minutes)		Average Collection Phase Time (Minutes)		Average Execution Time (Minutes)		Average Execution Cost (USD)	
		$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$
FR_{PE} *Auto + DW*	2	0.0022	0.0053	0.2815	0.2893	0.4256	0.4362	0.0011	0.0011
	4	0.0045	0.0083	0.9854	0.5825	1.1537	0.7530	0.0030	0.0020
	8	0.0105	0.0196	2.8430	1.1275	3.0950	1.3454	0.0081	0.0035
	16	0.0267	0.0468	8.1126	1.9528	8.6050	2.2651	0.0226	0.0059
	32	0.0762	0.1850	29.6007	3.5592	30.9374	4.2317	0.0813	0.0111
	64	0.3822	0.5770	123.6268	5.6632	130.7519	9.2003	0.3434	0.0242

Preliminary Experiments on EC2 – Average Times & Costs

Experiments Round	*N* SARS-CoV-2 Sequences	Average Diff Phase Time (Minutes)		Average Collection Phase Time (Minutes)		Average Execution Time (Minutes)		Average Execution Cost (USD)	
		$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$
SR_{PE} *Custom + DW*	2	0.0112	0.0141	0.2579	0.2881	0.4139	0.4731	0.0011	0.0012
	4	0.0200	0.0218	0.8406	0.4648	1.0336	0.6453	0.0027	0.0017
	8	0.0458	0.0419	2.6194	0.9115	2.9148	1.1462	0.0077	0.0030
	16	0.1237	0.1021	6.5475	1.7639	7.2007	2.1544	0.0189	0.0057
	32	0.4009	0.1576	18.9651	2.8872	20.9013	3.5074	0.0549	0.0092
	64	1.4076	0.5242	65.6623	4.9424	72.4692	6.5623	0.1904	0.0172

Preliminary Experiments on EC2 – Average Times & Costs

Experiments Round	*N* SARS-CoV-2 Sequences	Average Diff Phase Time (Minutes)		Average Collection Phase Time (Minutes)		Average Execution Time (Minutes)		Average Execution Cost (USD)	
		$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$
SR_{PE} *Custom + DW*	2	0.0112	0.0141	0.2579	0.2881	0.4139	0.4731	0.0011	0.0012
	4	0.0200	0.0218	0.8406	0.4648	1.0336	0.6453	0.0027	0.0017
	8	0.0458	0.0419	2.6194	0.9115	2.9148	1.1462	0.0077	0.0030
	16	0.1237	0.1021	6.5475	1.7639	7.2007	2.1544	0.0189	0.0057
	32	0.4009	0.1576	18.9651	2.8872	20.9013	3.5074	0.0549	0.0092
	64	1.4076	0.5242	65.6623	4.9424	72.4692	6.5623	0.1904	0.0172

Preliminary Experiments on EC2 – Average Times & Costs

Experiments Round	*N* SARS-CoV-2 Sequences	Average Diff Phase Time (Minutes)		Average Collection Phase Time (Minutes)		Average Execution Time (Minutes)		Average Execution Cost (USD)	
		$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$
TR_{PE} *Custom + MW*	2	0.0112	0.0127	0.1836	0.1746	0.3596	0.3291	0.0009	0.0009
	4	0.0204	0.0215	0.5877	0.3152	0.7770	0.4955	0.0020	0.0013
	8	0.0478	0.0455	1.7232	0.5287	2.0287	0.7694	0.0053	0.0020
	16	0.1358	0.0883	5.0213	0.9464	5.7202	1.3004	0.0150	0.0034
	32	0.3989	0.2186	16.1036	1.8633	18.0104	2.5403	0.0473	0.0067
	64	1.4035	0.7150	61.7254	3.4916	68.4508	5.2757	0.1798	0.0139

Preliminary Experiments on EC2 – Average Times & Costs

Experiments Round	*N* SARS-CoV-2 Sequences	Average Diff Phase Time (Minutes)		Average Collection Phase Time (Minutes)		Average Execution Time (Minutes)		Average Execution Cost (USD)	
		$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$	$DIFF_1$	$DIFF_{opt}$
TR_{PE} *Custom + MW*	2	0.0112	0.0127	0.1836	0.1746	0.3596	0.3291	0.0009	0.0009
	4	0.0204	0.0215	0.5877	0.3152	0.7770	0.4955	0.0020	0.0013
	8	0.0478	0.0455	1.7232	0.5287	2.0287	0.7694	0.0053	0.0020
	16	0.1358	0.0883	5.0213	0.9464	5.7202	1.3004	0.0150	0.0034
	32	0.3989	0.2186	16.1036	1.8633	18.0104	2.5403	0.0473	0.0067
	64	1.4035	0.7150	61.7254	3.4916	68.4508	5.2757	0.1798	0.0139

Main Experiments on Amazon EC2

Main Experiments on Amazon EC2 – Execution Settings

Execution Settings

- **1** On-Demand Master Node (Driver Program).
- **8** Spot Worker Nodes.
- Input (*N*): **540** South America SARS-CoV-2 Sequences.
- DataFrames' number of partitions: ***Customized***.
- Diff Phase: $\boldsymbol{DIFF_{opt}}$
- max_D: ***63***
- Collection Phase: ***MW***

Main Experiments on Amazon EC2 – VM Selection

Selection Criteria

- Optimization Families: **Memory** & **Storage**.
- Availability Zone: us-east-1a (Northern Virginia).
- Processor Architecture: 64 bits (x86_64).
- Availability Price Markets: On-Demand & Spot.
- Maximum On-Demand Price (Hour): 0.5 USD.

Main Experiments on Amazon EC2 – VM Selection

Family	Instance Name	Number of vCPUs	Memory (GiB)	Storage Type	Network Speed (Gbps)	Cost per Hour (USD)	
						On-Demand	Spot
Memory	r5.xlarge	4	32	EBS	Up to 10	0.2520	0.1374
	r5dn.xlarge	4	32	1x 150 NVMe SSD	Up to 25	0.3340	0.1232
	z1d.xlarge	4	32	1x 150 NVMe SSD	Up to 10	0.3720	0.1116
Storage	i3en.xlarge	4	32	1x 2500 NVMe SSD	Up to 25	0.4520	0.1356
	h1.2xlarge	8	32	1x 2000 HDD	Up to 10	0.4680	0.1404
	d3.xlarge	4	32	3x 2000 HDD	Up to 15	0.4990	0.1497

Main Experiments on Amazon EC2 – Average Times & Costs

Instance Name	Execution Time (Minutes)		Execution Cost (USD)		Percentage Change	
	Average	Standard Deviation	Average	Standard Deviation	Time	Cost
r5.xlarge	174.6113	3.0408	3.9322	0.0685	0%	0%
r5dn.xlarge	174.3515	1.0521	3.8346	0.0231	- 0.1487%	- 2.4820%
z1d.xlarge	135.1479	1.9064	2.8489	0.0402	- 22.6007%	- 27.5494%
i3en.xlarge	169.0511	5.7611	4.3300	0.1476	- 3.1843%	+ 10.1164%
h1.2xlarge	201.6883	1.1435	5.3488	0.0303	+ 15.5070%	+ 36.0256%
d3.xlarge	165.1419	0.7194	4.6697	0.0203	- 5.4231%	+ 18.7554%

Baseline: r5.xlarge instances results (cheapest On-Demand instances).

Main Experiments on Amazon EC2 – Average Times & Costs

Instance Name	Execution Time (Minutes)		Execution Cost (USD)			
	Average	Standard Deviation	Average	Standard Deviation		
r5.xlarge	174.6113	3.0408	3.9322	0.0685	0%	0%
r5dn.xlarge	174.3515	1.0521	3.8346	0.0231	- 0.1487%	- 2.4820%
z1d.xlarge	135.1479	1.9064	2.8489	0.0402	- 22.6007%	- 27.5494%
i3en.xlarge	169.0511	5.7611	4.3300	0.1476	- 3.1843%	+ 10.1164%
h1.2xlarge	201.6883	1.1435	5.3488	0.0303	+ 15.5070%	+ 36.0256%
d3.xlarge	165.1419	0.7194	4.6697	0.0203	- 5.4231%	+ 18.7554%

Memory Optimized

Baseline: r5.xlarge instances results (cheapest On-Demand instances).

Main Experiments on Amazon EC2 – Average Times & Costs

Revocation Scenarios:

- ***RS*$_1$**: 2 Workers revoked after 30 minutes of execution.
- ***RS*$_2$**: 2 Workers revoked after 60 minutes of execution.
- ***RS*$_3$**: 2 Workers revoked after 120 minutes of execution.

Main Experiments on Amazon EC2 – Average Times & Costs

Instance Name	Revocation Scenario	Execution Time (Minutes)		Execution Cost (USD)		Percentage Change	
		Average	Standard Deviation	Average	Standard Deviation	Time	Cost
z1d.xlarge	None	135.1479	1.9064	2.8489	0.0402	0%	0%
	RS_1	140.8041	1.5435	2.5560	0.0268	+ 4.1851%	- 10.2811%
	RS_2	139.2890	1.6629	2.6413	0.0289	+ 3.0641%	- 7.2870%
	RS_3	132.9266	5.0956	2.7540	0.0885	- 1.6436%	- 3.3311%
i3en.xlarge	None	169.0511	5.7611	4.3300	0.1476	0%	0%
	RS_1	183.7072	2.0837	3.3008	0.0362	+ 8.6696%	- 23.7690%
	RS_2	177.4600	1.4847	3.3039	0.0258	+ 4.9741%	- 23.6974%
	RS_3	175.9235	2.9039	3.5004	0.0504	+ 4.0652%	- 19.1593%

Baseline: z1d.xlarge (memory) & i3en.xlarge (storage) instances results without revocations.

Main Experiments on Amazon EC2 – Average Times & Costs

Instance Name	Revocation Scenario	Execution Time (Minutes)		Execution Cost (USD)		Percentage Change	
		Average	Standard Deviation	Average	Standard Deviation	Time	Cost
z1d.xlarge	None	135.1479	1.9064	2.8489	0.0402	0%	0%
	RS_1	140.8041	1.5435	2.5560	0.0268	+ 4.1851%	- 10.2811%
	RS_2	139.2890	1.6629	2.6413	0.0289	+ 3.0641%	- 7.2870%
	RS_3	132.9266	5.0956	2.7540	0.0885	- 1.6436%	- 3.3311%
i3en.xlarge	None	169.0511	5.7611	4.3300	0.1476	0%	0%
	RS_1	183.7072	2.0837	3.3008	0.0362	+ 8.6696%	- 23.7690%
	RS_2	177.4600	1.4847	3.3039	0.0258	+ 4.9741%	- 23.6974%
	RS_3	175.9235	2.9039	3.5004	0.0504	+ 4.0652%	- 19.1593%

Baseline: z1d.xlarge (memory) & i3en.xlarge (storage) instances results without revocations.

Conclusions and Future Directions

Conclusions and Future Directions

- When using memory and storage optimized spot instances **without any revocations**, **reductions of up to 22.60% of the average execution time and up to 62.22% of the average monetary cost** for the study case of **540 SARS-CoV-2 all-to-all sequence comparisons**.

- I**n the worst tested spot revocation scenario**, **reduction of the monetary cost in 10.28% and 23.77%** respectively for the z1d.xlarge and i3en.xlarge instances, while their respective **execution times slightly increased to 4.19% and 8.67%**, benefiting from the low overhead fault tolerance Spark framework.

- The experiments also showed **cost-benefit of running it on the memory optimized instances**, most outstanding being z1d.xlarge.

Future Steps

- **Concerning the data input**: variations in size and number of biological sequences;
- **Concerning the cluster**: variations in number of worker nodes and of executors per worker;
- **Concerning the resources metrics** (e.g., IOPS, bandwidth): better analysis of the performance and of the scaling bounds for Spark applications in the cloud.

Thank You!