Performance Evaluation and Tuning of BioPig for Genomic Analysis

Lizhen Shi Zhong Wang Weikuan Yu Xiandong Meng

Presenter: Lizhen Shi

Outline

- Background
  - DNA Sequence Analysis.
  - BioPig Toolkit.
  - K-mer Counting
- Initial Setup
- Tuning BioPig on YARN
- Related Work and Conclusion

DNA Sequence Analysis

- **DNA**: the hereditary material in humans and almost all other organisms
- **DNA sequence:**

- **DNA sequence analysis**: the process of subjecting a DNA to any of a wide range of analytical methods to understand its features, functions, structure, or evolution.

BioPig

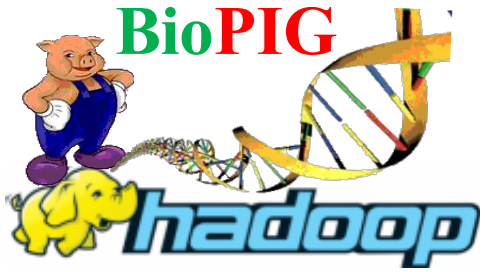

- **Hadoop-based** toolkit for **large-scale** sequence analysis.
- Built on **Hadoop** framework and **Pig** Latin data flow language.
- Developed by Lawrence Berkeley National Lab.
- The second version (2.0) is built on **Hadoop 2.7** and pig 0.15.

K-mer Counting

- **K-mer**: a short DNA sequence consisting of a fixed number (K) of bases.

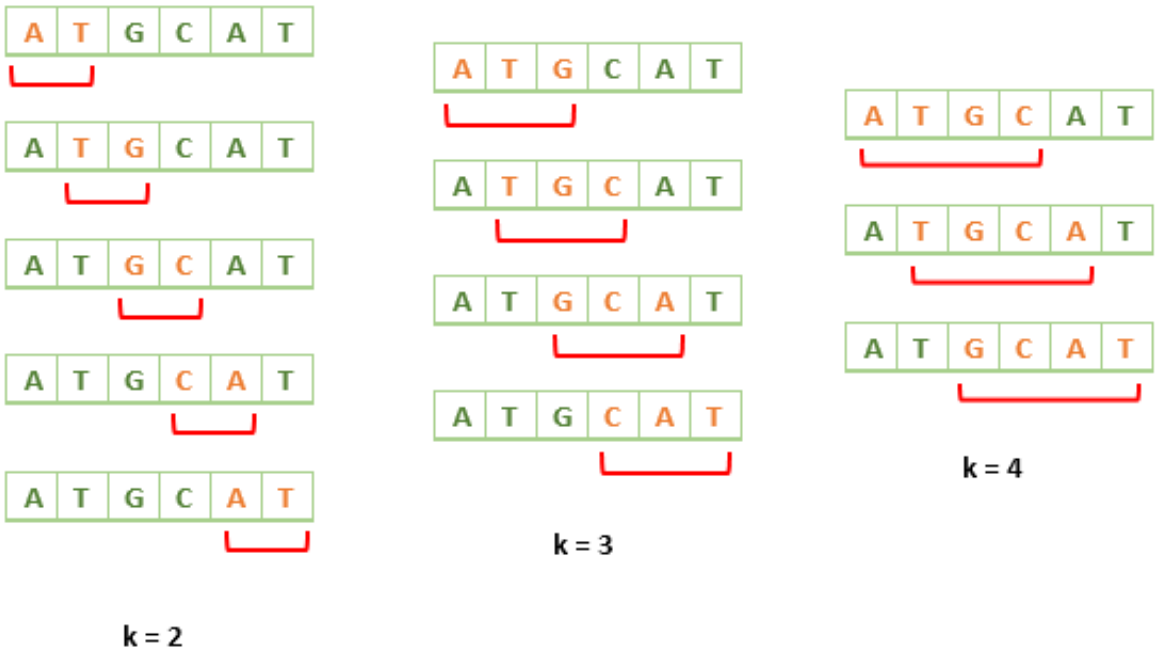

- **K-mer Counting**: core module of the BioPig toolkit and a prerequisite step of many bioinformatics applications.

Outline

- Background
  - DNA sequence analysis.
  - BioPig Toolkit.
  - K-mer Counting
- Initial Setup
- Tuning BioPig on YARN
- Related work and Conclusion

Challenges in Hadoop Tuning

- Large number of parameters need to be specified.
- Parameters determine the runtime of the job.
- The job execution time is sensitive to these parameters.
- The set of near-optimal values is dependent on the **hardware resource, input data size** and **application characteristics**.

Characteristics of K-mer Counting

Characteristics of Intermediate Data for Common Hadoop Applications

Job Name	Input size (TB)	Int. data size (TB)	Int./Input
LogProc	1.10	1.10	100%
NdayModel	3.54	3.54	100%
BehaviorModel	3.60	9.47	263%
ClickAttribution	6.80	8.20	121%
SegmentExploder	14.10	25.20	179%
LogRead	1.10	1.10	100%
LogCount	1.10	0.04	4%

Characteristics of Intermediate Data for k-mer Counting (k=20)

Input size (GB)	Int. data size (GB)	Int./Input
1	13.5	1350%
10	135.4	1354%
20	270.9	1355%
40	541.5	1354%
60	830.0	1383%
100	1381.0	1381%

Initial Setup

- System Configuration:
  - 15 EC2 c3.8x large type instances(RAM60GB, 108 Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz)
  - **SSD**, **10Gbit/s Ethernet** network.
- Input data:
  - Cow Rumen metagenomic dataset I250
- Workloads:
  - 1GB,10GB,20GB,40GB,60GB,100GB

Baseline Configuration

Configuration parameters	Original Value	Initial Value
yarn.nodemanager.resource.memory-mb	8192	58296
yarn.nodemanager.resource.cpu-vcores	8	30
yarn.scheduler.minimum-allocation-mb	1024	2048
yarn.scheduler.maximum-allocation-mb	8192	58296
yarn.scheduler.minimum-allocation-vcores	1	1
yarn.app.mapreduce.am.resource.mb	1536	2048
yarn.app.mapreduce.am.command-opts	-Xmx1024m	-Xmx1624m
mapreduce.map.memory.mb	1024	2048
mapreduce.reduce.memory.mb	1024	2048
mapreduce.map.java.opts	-Xmx200m	-Xmx1624m
mapreduce.reduce.java.opts	-Xmx200m	-Xmx1624m
io.sort.mb	100	1000
io.sort.factor	10	100
mapreduce.reduce.shuffle.parallelcopies	5	20
dfs.block.size	64M	64M
mapreduce.map.output.compress	FALSE	FALSE
mapreduce.map.output.compress.codec	DefaultCodec	DefaultCodec
mapreduce.output.fileoutputformat.compress	FALSE	FALSE
mapreduce.output.fileoutputformat.compress.codec	DefaultCodec	DefaultCodec
mapreduce.job.reduce.slowstart.completedmaps	0.05	0.05
yarn.nodemanager.pmem-check-enabled	TRUE	TRUE

Baseline Performance

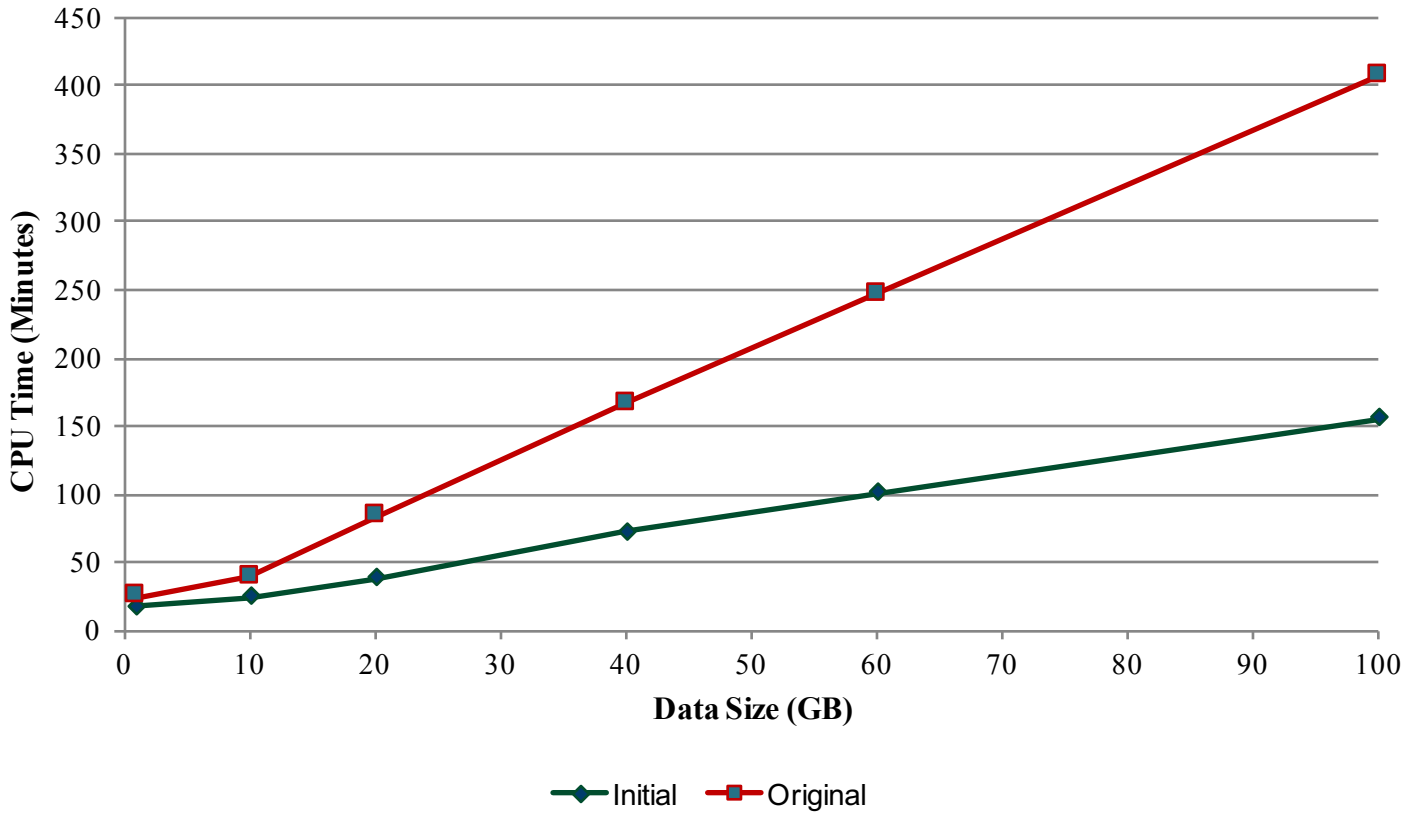

The overall execution time was reduced by **60%** using the initial configuration.

Outline

- Background
  - DNA sequence analysis.
  - Challenges involved in Bioinformatics.
  - MapReduce-based bioinformatics tools.
  - BioPig Toolkit.
- Initial Setup
- Tuning BioPig on YARN
- Related work and Conclusion

Performance Tuning Cycle

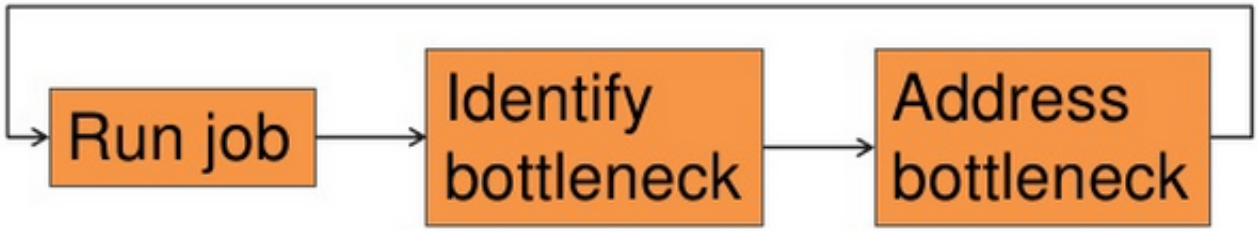

5 Iterations:

-Data Compression
-Block size
-Heap size
-JVM Garbage Collection(GC)
-Reducer start time

IO Reduction from Data Compression

Data Size	Counter Group	Uncompressed			Compressed			Difference		
		Map	Reduce	Total	Map	Reduce	Total	Map	Reduce	Total
40GB	Number of bytes read(GB)	604	598	1,202	283	131	414	321	467	788
	Number of bytes written(GB)	1,200	598	1,798	550	131	681	649	467	1,117
60GB	Number of bytes read(GB)	929	917	1,846	442	167	609	487	751	1,237
	Number of bytes written(GB)	1,839	917	2,756	853	167	1,020	986	751	1,736

More than 50% IO operation was reduced!!!

The overall execution time was reduced by 6 minutes for the 40GB input and 8 minutes for the 60GB input

IO Reduction from Heap Size Tuning

Data Size	Counter Group	Spill			No spill			Difference		
		Map	Reduce	Total	Map	Reduce	Total	Map	Reduce	Total
40GB	Number of bytes read	283	131	414	0	125	125	283	6	289
	Number of bytes written	550	131	681	278	125	403	272	6	278
	Map output records	32	0	32	32	0	32	0	0	0
	Spilled Records	63	31	95	32	32	63	31	0	31
60GB	Number of bytes read	442	167	609	0	156	156	442	11	453
	Number of bytes written	853	167	1,020	431	156	587	422	11	433
	Map output records	49	0	49	49	0	49	0	0	0
	Spilled Records	97	48	145	48	48	97	48	0	48

More than 50% IO operation was reduced!!!

The overall execution time was reduced by 1 minutes for the 40GB input and 3 minutes for the 60GB input

Performance Gain from GC Tuning

Data Size	Counter Group	Before GC tuning			After GC tuning			Difference		
	GC Time (s)	Map	Reduce	Total	Map	Reduce	Total	Map	Reduce	Total
40GB	Job1	183,905	529	184,433	2,014	104	2,118	181,891	424	182,315
	Job2	18,807	5,736	24,543	7,720	1,510	9,230	11,087	4,226	15,313
60GB	Job1	25,283	390	25,672	7,135	141	7,277	18,147	249	18,396
	Job2	26,577	8,120	34,697	11,040	2,302	13,342	15,538	5,818	21,355

The overall execution time was reduced by 4 minutes for the 40GB input and 8 minutes for the 60GB input

Impact of Reducer start time

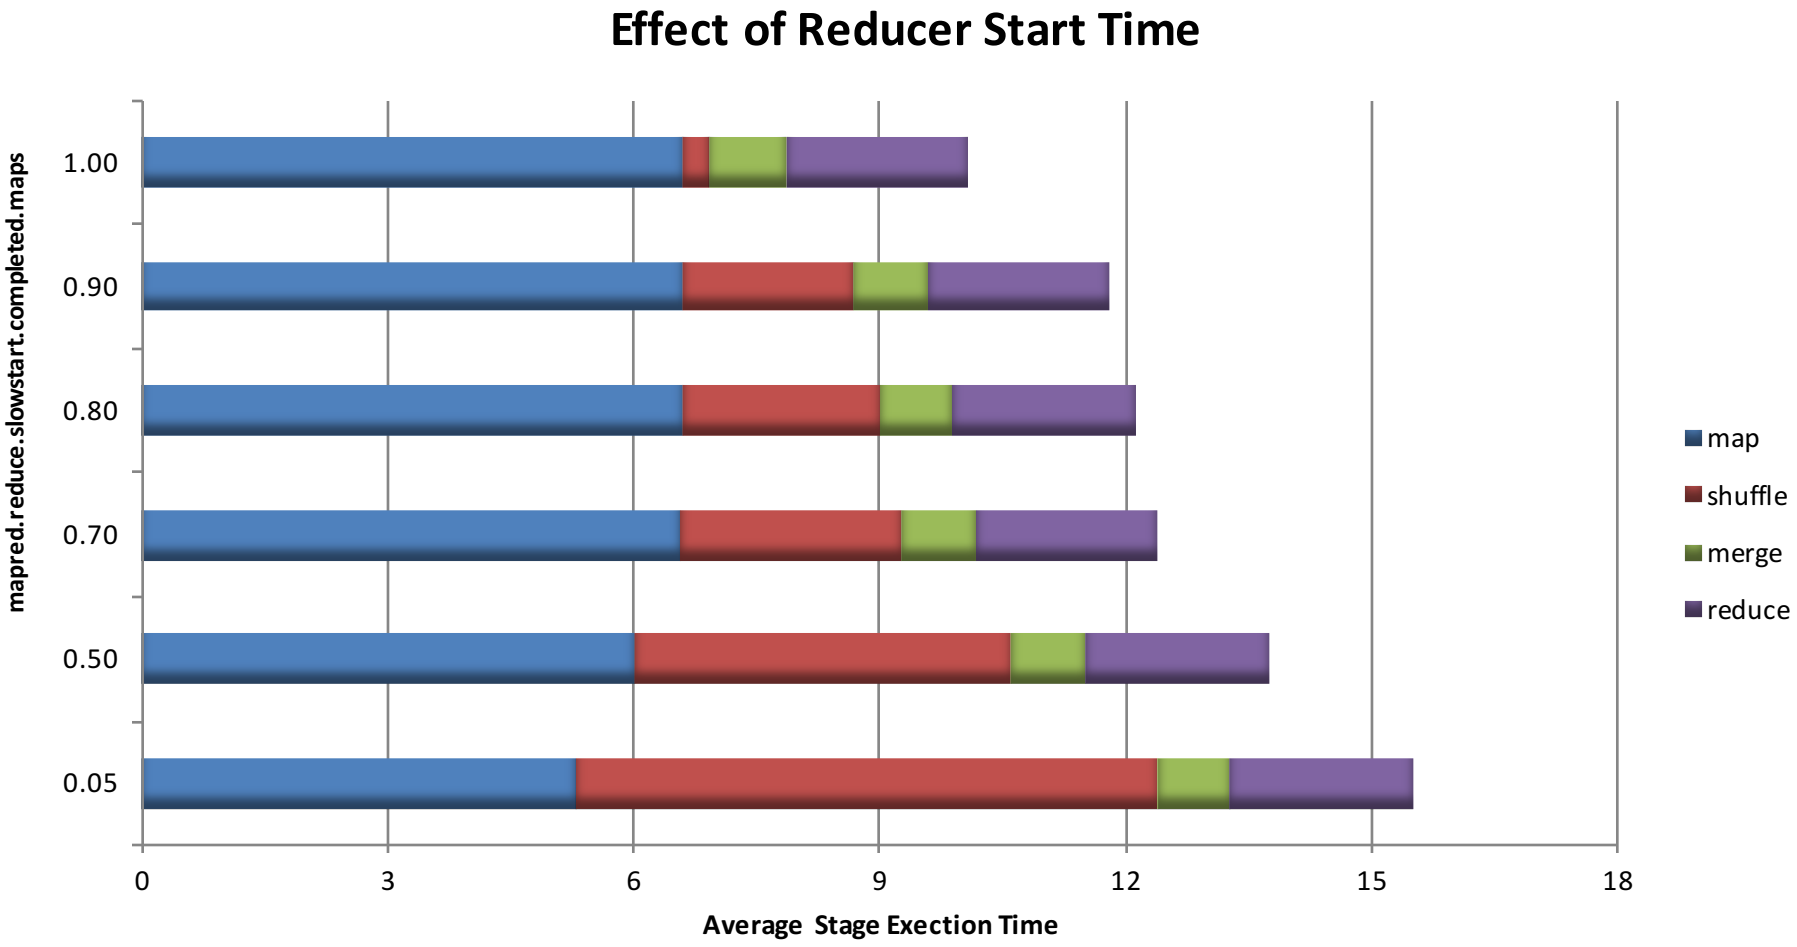

The overall execution time was reduced by 9 minutes for the 40GB input and 11 minutes for the 60GB input

Contribution Analysis

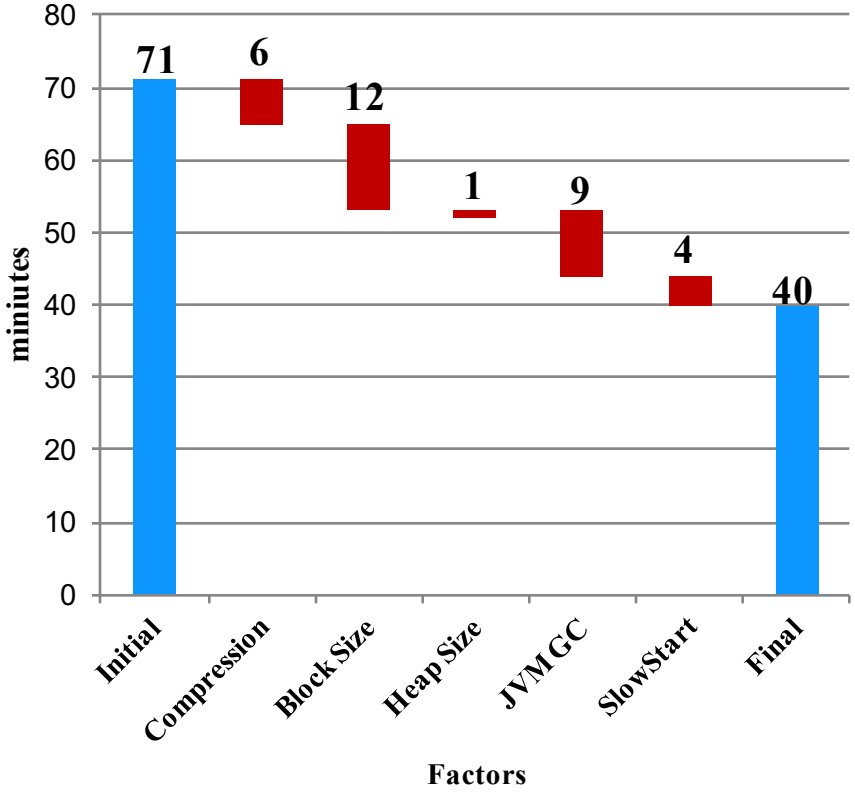

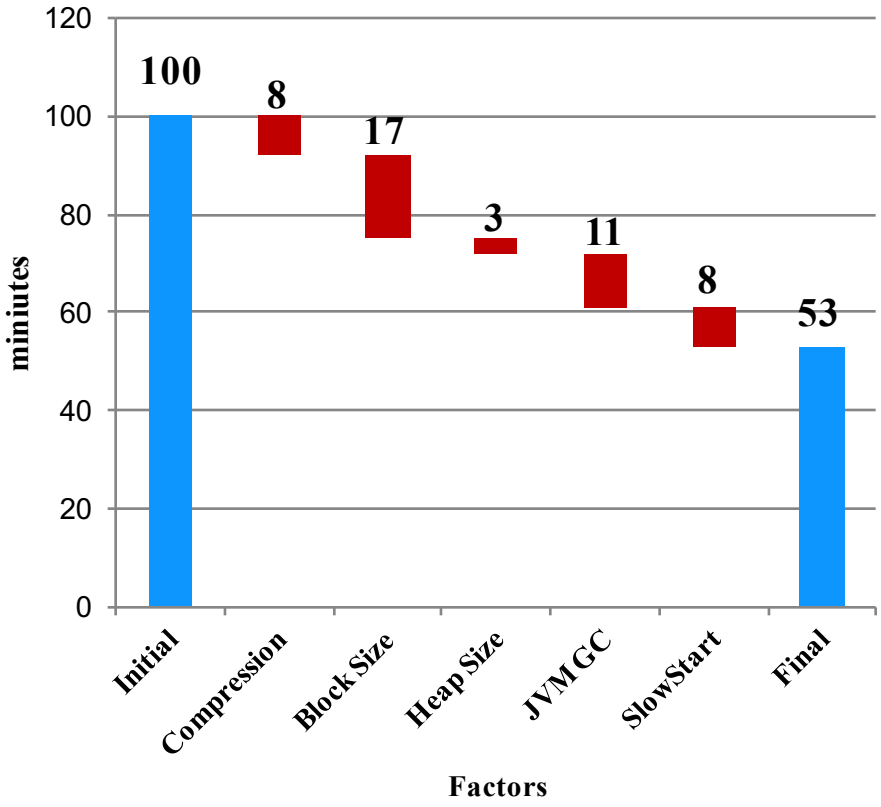

The overall execution time was reduced by **44%** for the 40GB input and **47%** for the 60GB input.

Final Tuned Parameters

Configuration parameters	Default value	Baseline Value	Tuned Value
yarn.nodemanager.resource.memory-mb	8192	58296	58296
yarn.nodemanager.resource.cpu-vcores	8	30	30
yarn.scheduler.minimum-allocation-mb	1024	2048	2048
yarn.scheduler.maximum-allocation-mb	8192	58296	58296
yarn.scheduler.minimum-allocation-vcores	1	1	1
yarn.app.mapreduce.am.resource.mb	1536	2048	2048
yarn.app.mapreduce.am.command-opts	-Xmx1024m	-Xmx1624m	Xmx2048m
mapreduce.map.memory.mb	1024	2048	2048
mapreduce.reduce.memory.mb	1024	2048	2048
mapreduce.map.java.opts	-Xmx200m	-Xmx1624m	Xmx2048m①
mapreduce.reduce.java.opts	-Xmx200m	-Xmx1624m	Xmx2048m①
io.sort.mb	100	1000	1100
io.sort.factor	10	100	100
mapreduce.reduce.shuffle.parallelcopies	5	20	20
dfs.block.size	64M	64M	32M
mapreduce.map.output.compress	FALSE	FALSE	TRUE
mapreduce.map.output.compress.codec	DefaultCodec	DefaultCodec	Snappy
mapreduce.output.fileoutputformat.compress	FALSE	FALSE	TRUE
mapreduce.output.fileoutputformat.compress.codec	DefaultCodec	DefaultCodec	Snappy
mapreduce.job.reduce.slowstart.completedmaps	0.05	0.05	0.9
yarn.nodemanager.pmem-check-enabled	TRUE	TRUE	FALSE

① :-Xms2048m -Xmx2048m -XX:PermSize=128M -XX:ParallelGCThreads=4 -verbose:gc -XX:+PrintGC -XX:+PrintGCDetails -XX:+PrintGCTimeStamps

Performance Trends(4x speedup)

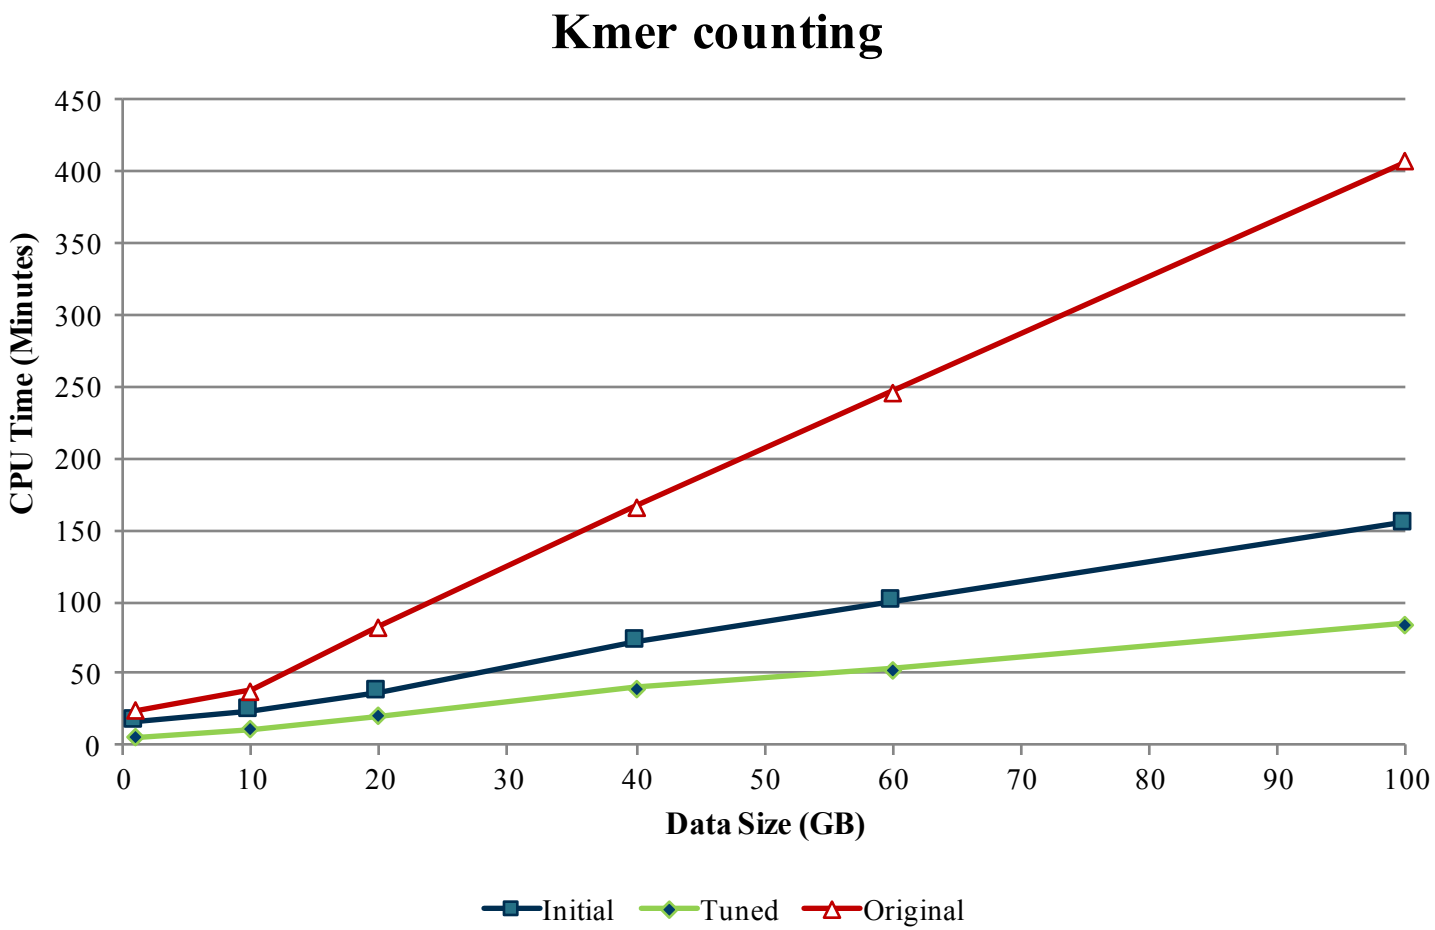

The initial performance was improved by **60%** Compared to original setting; the tuned performance Was improved by **50%** compared to initial performance

Outline

- Background
  - DNA sequence analysis.
  - Challenges involved in Bioinformatics.
  - MapReduce-based bioinformatics tools.
  - BioPig Toolkit.
- Initial Setup
- Tuning BioPig on YARN
- Related Work and Conclusion

Related Work

- **ICPE '12**: Apache hadoop performance-tuning methodologies and best practices
- **HPDC '14:** MRONLINE: MapReduce online performance tuning
- **ICAC '12:** AROMA: automated resource allocation and configuration of mapreduce environment in the cloud
- **CIDR '11:** Starfish: A Self-tuning System for Big Data Analytic Environment in the Cloud

Conclusion

- Tuned Hadoop parameters by 5 steps based on k-mer counting characteristics.
- Result showed these tunings achieved 4x speedup compared to default setting.
  - Baseline: 60% performance improvement compared to default
  - Tuned: 50% compared to baseline configuration.

Sponsors

NSF

FLORIDA STATE UNIVERSITY
VIRES
ARTES
MORES
1851

JGI
JOINT GENOME INSTITUTE
DEPARTMENT OF ENERGY

BERKELEY LAB
Lawrence Berkeley National Laboratory

Thank You!

Any questions, comments or suggestions are welcome !