

# GigaScience

## Hot-starting software containers for bioinformatics analyses

--Manuscript Draft--

<b>Manuscript Number:</b>	GIGA-D-17-00326	
<b>Full Title:</b>	Hot-starting software containers for bioinformatics analyses	
<b>Article Type:</b>	Technical Note	
<b>Funding Information:</b>	National Institutes of Health (U54HL127624)	Not applicable
	AMEDD Advanced Medical Technology Initiative	Not applicable
<b>Abstract:</b>	<p>Background: Using software containers has become standard practice to reproducibly deploy and execute biomedical workflows on the cloud. However, some applications which contain time-consuming initial steps will produce unnecessary costs for repeated executions.</p> <p>Findings: We demonstrate that hot-starting, from containers that have been frozen after the application has already begun execution, reduces the costs of cloud computing by avoiding repetitive initialization steps. In particular, we use an open source tool called Checkpoint and Restore in Userspace (CRIU) to save the state of the containers as checkpoint files after executing the initialization steps. We then migrate the checkpoint files to other virtual machines on the cloud and restored containers from checkpoint files. As a proof-of-concept example, we create a hot-start container for the STAR aligner and deploy this container to align RNA sequencing data. In addition, we compare the performance of the alignment step with and without checkpoints on cloud platforms using local and network disks.</p> <p>Conclusions: We demonstrate that hot-starting Docker containers from snapshots taken after repetitive initialization steps are completed, significantly reduces the costs of executing the STAR aligner on all experimental platforms, including Amazon Web Services (AWS), Microsoft Azure and local virtual machines. Our method can be employed in other bioinformatics applications in which a checkpoint can be inserted after a repetitive initialization phase.</p>	
<b>Corresponding Author:</b>	Ka Yee Yeung, Ph.D. University of Washington Tacoma UNITED STATES	
<b>Corresponding Author Secondary Information:</b>		
<b>Corresponding Author's Institution:</b>	University of Washington Tacoma	
<b>Corresponding Author's Secondary Institution:</b>		
<b>First Author:</b>	Pai Zhang	
<b>First Author Secondary Information:</b>		
<b>Order of Authors:</b>	Pai Zhang	
	Ling-Hong Hung, Ph.D.	
	Wes Lloyd, Ph.D.	
	Ka Yee Yeung, Ph.D.	
<b>Order of Authors Secondary Information:</b>		
<b>Opposed Reviewers:</b>		
<b>Additional Information:</b>		
<b>Question</b>	<b>Response</b>	

<p>Are you submitting this manuscript to a special series or article collection?</p>	<p>No</p>
<p><b>Experimental design and statistics</b></p> <p>Full details of the experimental design and statistical methods used should be given in the Methods section, as detailed in our <a href="#">Minimum Standards Reporting Checklist</a>. Information essential to interpreting the data presented should be made available in the figure legends.</p> <p>Have you included all the information requested in your manuscript?</p>	<p>Yes</p>
<p><b>Resources</b></p> <p>A description of all resources used, including antibodies, cell lines, animals and software tools, with enough information to allow them to be uniquely identified, should be included in the Methods section. Authors are strongly encouraged to cite <a href="#">Research Resource Identifiers</a> (RRIDs) for antibodies, model organisms and tools, where possible.</p> <p>Have you included the information requested as detailed in our <a href="#">Minimum Standards Reporting Checklist</a>?</p>	<p>Yes</p>
<p><b>Availability of data and materials</b></p> <p>All datasets and code on which the conclusions of the paper rely must be either included in your submission or deposited in <a href="#">publicly available repositories</a> (where available and ethically appropriate), referencing such data using a unique identifier in the references and in the “Availability of Data and Materials” section of your manuscript.</p> <p>Have you have met the above requirement as detailed in our <a href="#">Minimum Standards Reporting Checklist</a>?</p>	<p>Yes</p>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

## Hot-starting software containers for bioinformatics analyses

Pai Zhang

Institute of Technology, University of Washington, Tacoma, Washington 98402, USA.

Email: [paizhang@uw.edu](mailto:paizhang@uw.edu)

Ling-Hong Hung

Institute of Technology, University of Washington, Tacoma, Washington 98402, USA.

Email: [lhung@uw.edu](mailto:lhung@uw.edu)

Wes Lloyd

Institute of Technology, University of Washington, Tacoma, Washington 98402, USA.

Email: [wllloyd@uw.edu](mailto:wllloyd@uw.edu)

Ka Yee Yeung

Institute of Technology, University of Washington, Tacoma, Washington 98402, USA.

Email: [kayee@uw.edu](mailto:kayee@uw.edu)

Correspondence should be addressed to K.Y.Y. ([kayee@uw.edu](mailto:kayee@uw.edu))

1  
2  
3  
4 **ABSTRACT**  
5

6 **Background:** Using software containers has become standard practice to reproducibly deploy  
7 and execute biomedical workflows on the cloud. However, some applications which contain  
8 time-consuming initial steps will produce unnecessary costs for repeated executions.  
9

10  
11 **Findings:** We demonstrate that hot-starting, from containers that have been frozen after the  
12 application has already begun execution, reduces the costs of cloud computing by avoiding  
13 repetitive initialization steps. In particular, we use an open source tool called Checkpoint and  
14 Restore in Userspace (CRIU) to save the state of the containers as checkpoint files after  
15 executing the initialization steps. We then migrate the checkpoint files to other virtual machines  
16 on the cloud and restored containers from checkpoint files. As a proof-of-concept example, we  
17 create a hot-start container for the STAR aligner and deploy this container to align RNA  
18 sequencing data. In addition, we compare the performance of the alignment step with and  
19 without checkpoints on cloud platforms using local and network disks.  
20  
21  
22

23  
24 **Conclusions:** We demonstrate that hot-starting Docker containers from snapshots taken after  
25 repetitive initialization steps are completed, *significantly* reduces the costs of executing the  
26 STAR aligner on all experimental platforms, including Amazon Web Services (AWS), Microsoft  
27 Azure and local virtual machines. Our method can be employed in other bioinformatics  
28 applications in which a checkpoint can be inserted after a repetitive initialization phase.  
29  
30  
31  
32  
33  
34  
35  
36

37 **Keywords:** software container, reproducibility of research, cloud computing  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

## FINDINGS

### Background

With the availability of high-throughput next generation sequencing technologies and the subsequent explosion of big biomedical data, the processing of biomedical big data has become a major challenge. Cloud computing plays an important role in addressing this challenge by offering massive scalable computing and storage, data sharing and on-demand access to resources and applications [1, 2]. The National Institutes of Health is launching a Data Commons Pilot Phase to provide access and storage of biomedical data and bioinformatics tools on the cloud (<https://commonfund.nih.gov/bd2k>). Additionally, software containers have become increasingly popular for deploying bioinformatics workflows on the cloud. Docker (<https://www.docker.com/>), an open source project, has become the *de facto* standard for container software. Docker packages executables with all the necessary software dependencies ensuring that the same software environment is replicated regardless of the host hardware and operating system. Thus, containerization enhances the reproducibility of bioinformatics workflows. In the context of cloud computing, the utility of containers comes from the ease in which a virtual cloud cluster can be rapidly provisioned with all of the necessary dependencies for a complicated workflow by simply downloading a set of containers, each of which take a few seconds to spin up. Recently, Vivian et al. processed over 20,000 RNA sequencing (RNA-seq) samples from the Cancer Genome Atlas (TCGA) using Docker containers on the cloud [3]. Tatlow *et al.* used software containers to study the performance and cost profiles of different cloud-based configurations in processing RNA-seq data from public cancer compendia [4].

When containers are deployed, applications are launched *de novo* each time the container is spun up. This means that any initial preparatory steps are repeated each time the container is used. For applications such as the alignment of reads, these initial steps can be quite substantive as an entire reference genome is read in and indices are generated. In an automated large-scale deployment, these steps are replicated many times. It would be far more efficient if one could “*checkpoint*” and save containers in states where the application has already completed the initialization steps so as to avoid unnecessary repetitions. One could then “*hot-start*” workflows from these checkpoints. This is analogous to hot-start PCR where all the necessary reagents are pre-mixed awaiting only the addition of the template.

### Our approach

1  
2  
3  
4 Our key idea is to save and restore memory states in software containers using the Checkpoint  
5 Restore in Userspace (CRIU) tool. CRIU freezes a running container and saves the checkpoint  
6 as a collection of files on disk ([https://criu.org/Main\\_Page](https://criu.org/Main_Page)). These files can subsequently be  
7 used to restore and resume the application from that checkpoint. CRIU was originally  
8 developed for Linux, but has recently become available for Docker (<https://criu.org/Docker>).  
9 While it is possible to stop Docker containers with native docker commands, this process does  
10 not preserve the memory state. Although re-starting from a ready-to-go state is an intuitive  
11 application of checkpointing, we have been unable to find any previous description of using  
12 checkpointing as a general method for improving the efficiency of container deployments.  
13  
14  
15  
16  
17  
18  
19  
20

21 We demonstrate that hot-starting from a saved container checkpoint can significantly reduce the  
22 execution time using the STAR aligner [5, 6] for RNA-seq data analyses. We choose STAR as  
23 a proof-of-concept example because it has an option to save an intermediate state. However,  
24 our idea of using checkpoints has broad applications in optimizing performance using software  
25 containers on the cloud *when deploying any bioinformatics task where a pause could be*  
26 *inserted to capture a re-usable state.*  
27  
28  
29  
30  
31

32 The STAR aligner [5, 6] consists of two steps. In the first step, genome indices using the  
33 reference genome as input are generated. In the second step, read sequences from a specific  
34 experiment sample are mapped to the reference genome assuming that the genome indices  
35 have already been generated. In particular, STAR has the option of keeping the indices in  
36 memory after they have been generated to avoid repeating the first step when multiple files are  
37 to be aligned to the same reference genome. We used the CRIU tool to create checkpoints after  
38 the first step of generating genome indices. Instead of launching a new container and starting  
39 STAR from scratch, we restore the container state using CRIU and resume running STAR after  
40 it has loaded the indices. **Figure 1** shows an overview of our approach with and without using  
41 checkpoints.  
42  
43  
44  
45  
46  
47  
48  
49  
50

## 51 **Testing**

52 To test the checkpointing methodology, we used RNA-seq data generated by Himes *et al.* which  
53 measure the gene expression changes in human airway smooth muscle cells in response to  
54 asthma medications [7]. We compared the time required to align the sequences with a normal  
55 container where STAR starts from scratch, and the time required when hot-starting from a  
56 container checkpoint where STAR has already generated indices. We performed empirical  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 studies on cloud instances including the Amazon Web Services (AWS) and Microsoft Azure,  
5 using both local and networked disks. On AWS, we compared the performance with data on the  
6 local host and Amazon Elastic Block Store (EBS). On Microsoft Azure, we compared the  
7 performance with data on local host and Azure File Storage. Please refer to the Online  
8 Methods for details of our experimental setup. Our empirical results are shown in **Figure 2**.

9  
10  
11  
12  
13  
14 **Figure 2** shows that the STAR aligner with checkpoint reduces the execution time compared to  
15 STAR without checkpoint. On AWS, we observed 1.89x speedup with data stored on the local  
16 disk and 1.42x speedup with data on network disk (Amazon EBS). On Microsoft Azure, we  
17 achieved a 1.34x speedup with data stored on the local disk, and 3.57x speedup with data on  
18 Azure File Storage. With respect to execution time, we show that hot-starting from checkpoint  
19 containers save 2 minutes on fast local disks and Amazon EBS disks. The saving is almost 20  
20 minutes when using Azure network storage where the disk caching scheme appears to be much  
21 less favorable to STAR's index generation process.

22  
23  
24  
25  
26  
27  
28  
29 In this article, we have presented a novel idea for optimising cloud deployments using  
30 checkpointing to save containers where the applications are already started. Using CRIU for  
31 Docker, we can save the container with a preloaded genome for STAR alignment and restore  
32 the container from these checkpoint files to any host. We have achieved successful migration  
33 of checkpointed containers to different virtual machine instances running on the Amazon and  
34 Azure cloud platforms while realizing up to a 3.57x speedup using our approach saving up to 20  
35 minutes for a single STAR alignment workflow on Azure with network disks. For STAR  
36 alignment, it is possible to use a checkpointed container to align multiple sequences at once by  
37 retaining the genomic indices in memory. Our approach yields a significant benefit with hot-  
38 starting when as few as one or two files are aligned. Additionally, multiple STAR alignment  
39 tasks can be computed in parallel using the same genome indices hosted by different  
40 processes. For automated schedulers such as Docker Compose  
41 (<https://docs.docker.com/compose/>), "hot-starting" reduces execution time every single time the  
42 STAR container is launched. While it is possible to design a workflow to perform all the  
43 alignments in a single container first, load-balancing optimizations would be better utilized by  
44 allowing the scheduler to distribute the computation over the cluster as shorter jobs.

45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57 Our hot-start strategy only requires that there is a convenient place for a pause, checkpoint and  
58 re-start. In the case of STAR, this is provided by a flag that allows the container to keep  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 genomic indices in shell memory between invocations of STAR. For other workflows, one could  
5 add a flag to pause the computation where the checkpoint is to be created, and a flag to resume  
6 the computation afterwards. With these straightforward modifications, any application could take  
7 advantage of checkpointing to avoid repetitive initialization. This is a novel and unexplored  
8 approach to optimising containerized workflows while reducing the costs of cloud computing.  
9  
10  
11  
12

## 13 **METHODS**

14  
15  
16 **CRIU.** CRIU (Checkpoint/Restore In Userspace) is a Linux software tool that freezes a running  
17 application and saves it as a collection of files to disk (<https://criu.org>). The application can later  
18 be restored on the same or on a different host. Docker currently integrates CRIU as an  
19 experimental checkpoint sub-command that saves the state of processes to a collection of files  
20 on disk. The checkpointing command has been used to migrate containers from the source host  
21 to a target host when the resources of the source are limited [8], for fault tolerance purposes [9],  
22 and to provide highly available and scalable micro-services [10].  
23  
24  
25  
26  
27  
28

### 29 **Cloud configurations tested**

30  
31 In our experiments, we deployed our containers on instances from two cloud platforms:  
32 Microsoft Azure and Amazon Web Services (AWS). Ubuntu 16.04 was the host operating  
33 system in all our tests. Testing was conducted using a standard DS13 v2 instance with 8 virtual  
34 CPUs and 56 Gb memory on an Azure and a m4.xlarge instance with 16 virtual CPUs and 64  
35 GB memory on AWS. As disk I/O is an important factor in the efficiency of CRIU restoration and  
36 the generation of indices without CRIU, instances were tested using both network based disks  
37 (EBS for AWS and Microsoft Azure File Storage for Azure) and locally attached disks.  
38  
39  
40  
41  
42  
43

### 44 **Creating hot-start containers**

45  
46 We installed CRIU on the host Ubuntu system. Docker Community Edition (Docker CE), which  
47 includes the experimental checkpointing tool, was then installed. The STAR binary was  
48 compiled from source (<https://github.com/alexdobin/STAR>) using Ubuntu 16.04 and g++ and  
49 then copied into a fresh Ubuntu 16.04 container to get rid of all the intermediate build files. The  
50 build code and Dockerfiles are available from <https://github.com/BioDepot/ubuntu-star>. To  
51 create the checkpoint, STAR was launched with the *genomeLoad* flag set to *LoadAndKeep*.  
52 This keeps the indices in shared memory after STAR exits. To trap the container in this state,  
53 we launched STAR using a parent shell script that did not exit, and checkpointed the container  
54 after STAR exited. This results in the generation of checkpoint files that store the state of the  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65



1  
2  
3  
4 hot-start container. Due to different kernel versions being used, we created separate hot-start  
5 containers for AWS and Azure.  
6  
7

### 9 **Comparing hot-start containers and standard cold-start containers**

10 The paired-end fastq files were 9 Gb in size comprising 22,935,521 reads. Times were recorded  
11 for the generation of aligned BAM files using STAR in the standard container and using STAR  
12 with the hot-start container. Times include the time required to restore the hot-start container  
13 from the checkpointed files.  
14  
15  
16  
17

### 18 **AVAILABILITY AND REQUIREMENTS**

19 **Project name:** Hot-starting software container for STAR Alignment  
20

21 **Project homepage:** <https://github.com/paizhang/Hotstarting-For-STAR-Alignment>  
22

23 **DockerHub URL:** <https://hub.docker.com/r/biodepot/star-for-criu/>  
24

25 **Operating system:** Ubuntu 16.04  
26

27 **Programming language:** Shell  
28

29 **Other requirements:** Docker API version 1.25 or higher is required, Linux kernel v3.11 or  
30 higher is required.  
31

32 **License:** MIT License.  
33  
34

### 35 **AVAILABILITY OF SUPPORTING DATA**

36 The fastq files used in our tests were generated by Himes et al. and are publicly available from  
37 GEO with accession number GSE52778  
38 (<https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE52778>).  
39  
40  
41  
42

### 43 **ADDITIONAL FILES**

44 **Additional File 1:** User manual for “Hot-starting software container for STAR Alignment”  
45  
46  
47

### 48 **ABBREVIATIONS:**

49  
50 AWS : Amazon Web Services;  
51

52 CRIU: Checkpoint/Restore In Userspace;  
53

54 EBS: Elastic Block Store;  
55

56 NIH: National Institutes of Health;  
57

58 STAR: Spliced Transcripts Alignment to a Reference;  
59

60 TCGA : The Cancer Genome Atlas.  
61  
62  
63  
64  
65

1  
2  
3  
4  
5  
6  
7 **Conflict of interests**

8 The authors declare that they have no competing interests.  
9

10  
11 **AUTHOR CONTRIBUTIONS**

12 P.Z. implemented the Docker containers and conducted the empirical experiments. P.Z., L.H.H.  
13 and K.Y.Y. drafted the manuscript. K.Y.Y. and L.H.H. designed the case study. W.L. provided  
14 cloud computing expertise. K.Y.Y. coordinated the empirical study. All authors edited the  
15 manuscript.  
16  
17  
18  
19  
20  
21  
22

23 **ACKNOWLEDGEMENTS**

24 L.H.H. and K.Y.Y. are supported by NIH grant U54HL127624 and the AMEDD Advanced  
25 Medical Technology Initiative. We would like to acknowledge support from the AWS Cloud  
26 Credits for Research (to Lloyd and Yeung) and the Microsoft Azure for Research programs (to  
27 Hung and Lloyd) for providing cloud computing resources. We would like to acknowledge the  
28 Student High Performance Computing Club and the eScience Institute at University of  
29 Washington for both technical assistance and computing resources to Pai Zhang.  
30  
31  
32  
33  
34  
35  
36

37 **REFERENCES**

- 38  
39 1. Calabrese B and Cannataro M. Cloud Computing in Bioinformatics: current solutions and  
40 challenges. PeerJ Preprints. 2016;4:e2261v1.  
41  
42 2. Shanahan HP, Owen AM and Harrison AP. Bioinformatics on the cloud computing  
43 platform Azure. PLoS One. 2014;9 7:e102642. doi:10.1371/journal.pone.0102642.  
44  
45 3. Vivian J, Rao A, Nothhaft FA, Ketchum C, Armstrong J, Novak A, et al. Rapid and efficient  
46 analysis of 20,000 RNA-seq samples with Toil. bioRxiv. 2016.  
47  
48 4. Tatlow PJ and Piccolo SR. A cloud-based workflow to quantify transcript-expression  
49 levels in public cancer compendia. Scientific reports. 2016;6:39259.  
50 doi:10.1038/srep39259.  
51  
52 5. Dobin A, Davis CA, Schlesinger F, Drenkow J, Zaleski C, Jha S, et al. STAR: ultrafast  
53 universal RNA-seq aligner. Bioinformatics. 2013;29 1:15-21.  
54 doi:10.1093/bioinformatics/bts635.  
55  
56 6. Dobin A and Gingeras TR. Mapping RNA-seq Reads with STAR. Current protocols in  
57 bioinformatics. 2015;51:11 4 1-9. doi:10.1002/0471250953.bi1114s51.  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

7. Himes BE, Jiang X, Wagner P, Hu R, Wang Q, Klanderma B, et al. RNA-Seq transcriptome profiling identifies CRISPLD2 as a glucocorticoid responsive gene that modulates cytokine function in airway smooth muscle cells. *PLoS One*. 2014;9 6:e99625. doi:10.1371/journal.pone.0099625.
8. Al-Dhuraibi Y, Paraiso F, Djarallah N and Merle P. Autonomic Vertical Elasticity of Docker Containers with ELASTICDOCKER. In: *IEEE 10th International Conference on Cloud Computing* Honolulu, CA, USA, 25-30 June 2017 2017, IEEE.
9. Ismail BI, Goortani EM, Karim MBA, Tat WM, Setapa S, Luke JY, et al. Evaluation of Docker as Edge Computing Platform. In: *IEEE Confernece on Open Systems (ICOS)* Bandar Melaka, Malaysia 24-26 Aug. 2015 2015, IEEE.
10. Chen Y. Checkpoint and Restoration of Micro-service in Docker Containers. In: *Third International Conference on Mechatronics and Industrial Informatics* 2015.

1  
2  
3  
4 **FIGURE CAPTION**  
5  
6

7 **Figure 1.** An overview of our approach with and without checkpoints. The left panel shows the  
8 two steps of the STAR aligner [5, 6]. Note that the generation of the genome indices in the first  
9 step only requires the reference genome as input, thus the data from the experimental sample is  
10 only used in the second (mapping) step of STAR. The right panel shows our approach using  
11 the Checkpoint Restore in Userspace (CRIU) tool that freezes a running container and saves  
12 the checkpoint as a collection of files on disk after the genome indices are generated using the  
13 reference genome. Our “hot-start” containers use these saved files to restore the application  
14 and map the reads from the experimental sample data to the reference.  
15  
16  
17  
18  
19  
20  
21


22 **Figure 2.** STAR alignment running time comparison with checkpoint and without checkpoint.  
23 We performed our empirical experiments on two cloud platforms: Amazon Web Services (AWS)  
24 and Microsoft Azure. Both the Azure File Storage and the Amazon Elastic Block Store (EBS)  
25 represent network disks. We observe that our “hot-start” containers (orange and grey bars)  
26 provide a major reduction in execution time, especially on local disks.  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

# Traditional STAR Alignment

Reference genome




Generating genome indices



Input reads from sample



Mapping reads to genome




Output files




# Hot-start STAR alignment

Checkpoint and save



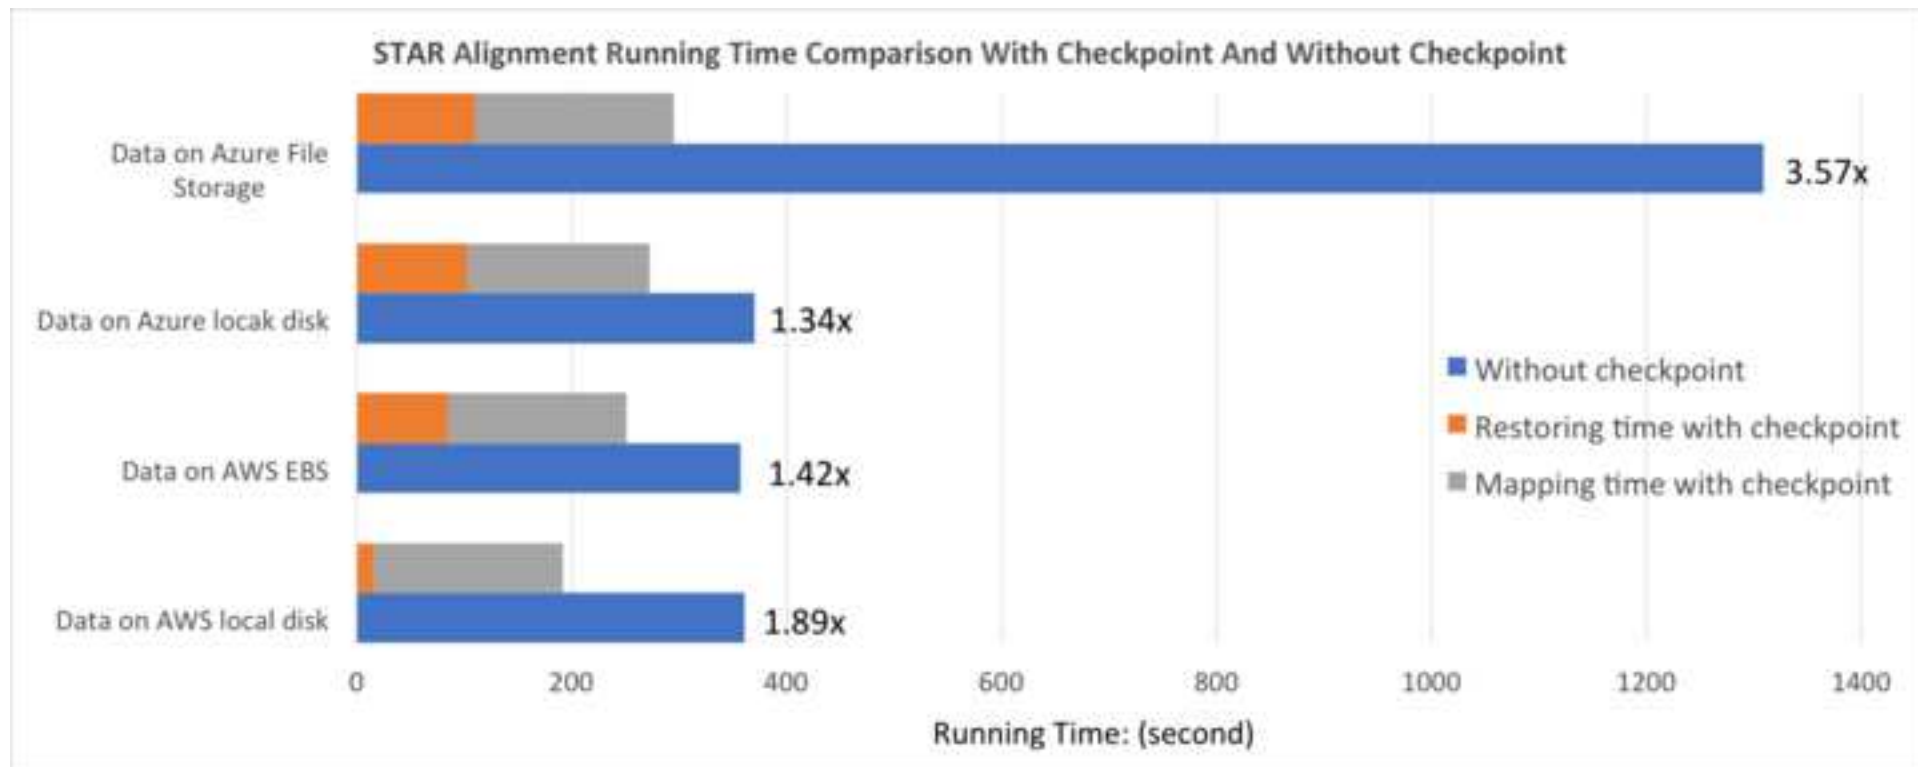
Input reads from sample

Mapping reads to genome



Output files







Click here to access/download  
**Supplementary Material**  
User Manual Hot-start Containers.pdf

