



# Distributed Computing in a Pandemic: A Review of Technologies Available for Tackling COVID-19

Jamie J. Alnasir

Department of Computing, Imperial College London, 180 Queen's Gate, London, SW7, UK  
2AZ.j.alnasir@imperial.ac.uk

## KEYWORD

SARS-CoV-2;  
COVID-19;  
distributed;  
HPC;  
supercomputing;  
grid; cloud;  
cluster

## ABSTRACT

*The current COVID-19 global pandemic caused by the SARS-CoV-2 betacoronavirus has resulted in over a million deaths and is having a grave socio-economic impact, hence there is an urgency to find solutions to key research challenges. Some important areas of focus are: vaccine development, designing or repurposing existing pharmacological agents for treatment by identifying druggable targets, predicting and diagnosing the disease, and tracking and reducing the spread. Much of this COVID-19 research depends on distributed computing.*

*In this article, I review distributed architectures — various types of clusters, grids and clouds — that can be leveraged to perform these tasks at scale, at high-throughput, with a high degree of parallelism, and which can also be used to work collaboratively. High-performance computing (HPC) clusters, which aggregate their compute nodes using high-bandwidth networking and support a high-degree of inter-process communication, are ubiquitous across scientific research — they will be used to carry out much of this work. Several bigdata processing tasks used in reducing the spread of SARS-CoV-2 require high-throughput approaches, and a variety of tools, which Hadoop and Spark offer, even using commodity hardware.*

*Extremely large-scale COVID-19 research has also utilised some of the world's fastest supercomputers, such as IBM's SUMMIT — for ensemble docking high-throughput screening against SARS-CoV-2 targets for drug-repurposing, and high-throughput gene analysis — and Sentinel, an XPE-Cray based system used to explore natural products. Likewise, RSC's TORNADO has been employed in aptamer design. Grid computing has facilitated the formation of the world's first Exascale grid computer. This has accelerated COVID-19 research in molecular dynamics simulations of SARS-CoV-2 spike protein interactions*

Jamie J. Alnasir

Distributed Computing in a Pandemic: A Review of  
Technologies available for Tackling COVID-19



*through massively-parallel computation and was performed with over 1 million volunteer computing devices using the Folding@home platform. Grids and clouds both can also be used for international collaboration by enabling access to important datasets and providing services that allow researchers to focus on research rather than on time-consuming data-management tasks.*

## 1. Introduction

A novel betacoronavirus named SARS-CoV-2 (Severe Acute Respiratory Syndrome coronavirus 2) is the cause of the clinical disease COVID-19 — its spread is responsible for the current coronavirus pandemic and the resulting global catastrophe (Lake, 2020). The initial outbreak of the disease was first detected in December 2019 in Wuhan (Hubei province, China) manifesting as cases of pneumonia, initially of unknown aetiology. On the 10th of January 2020, Zhang et al. released the initial genome of the virus (Zhang, 2020). Shortly after, it was identified — by deep sequencing analysis of lower respiratory tract samples — as a novel betacoronavirus and provisionally named 2019 novel coronavirus (2019-nCoV) (Lu et al., 2020; Huang et al., 2020a). By the 30th of January 2020, the WHO (World Health Organisation) declared the outbreak a Public Health Emergency of International Concern (WHO, 2020), and a global pandemic on the 11th of March (Organization et al., 2020). At this present time of writing (September 2021) there now are over 233 million reported cases of COVID-19 globally and more than 4,779,000 deaths have occurred as a result of the disease (Dong et al., 2020). In addition to the casualties, the pandemic is also having a grave socio-economic impact; it is a global crisis to which researchers will typically apply a variety computational techniques and technologies to several key areas of focus (Zhang et al., 2020; Nicola et al., 2020). These include, but are not limited to, vaccine development, designing or repurposing existing pharmacological agents for treatment by identifying druggable targets, predicting and diagnosing the disease, e.g. clinical decision support, and tracking and reducing the spread (Ferretti et al., 2020; Kissler et al., 2020; Perez and Abadi, 2020). Many of the tasks involved can leverage a variety of distributed computing approaches which can be applied at scale, at high-throughput, and with a high degree of parallelism — they often also need to be performed collaboratively.

The classification of SARS-CoV-2 as a betacoronavirus, and the release of its genome earlier on January 2020, has enabled research to focus on specific strategies. It is known from the previous 2003 SARS outbreak that ACE2 (Angiotensin Converting Enzyme) is the main entry point the virus targets to infect its host (Li et al., 2003; Kuba et al., 2005). To this end, for drug repurposing or development, COVID-19 research is focused on modelling the interaction between the coronavirus spike protein (S-protein) and ACE2, and in understanding the structure of the S-protein as an epitope for vaccine development. Other important targets are the virus’s proteome and the Papain-like and Main proteases — *PL-pro* and *ML-pro*, respectively (Hilgenfeld, 2014). Given the urgency to reduce mortality, significant efforts are being made to re-purpose medicines that are appropriate and already approved. Whilst a WHO scientific briefing refers to this practice — *off-label prescribing* — in the clinical setting, much of the initial work to predict potential drug candidates will be carried out computationally via *in-silico* screening (Kalil, 2020; WHO, 2020). Furthermore, the scale of the pandemic and the global production of bigdata, particularly whilst vaccines are still being developed, will rely on bigdata analytics to model the spread of the disease, and inform government policy and reduce the death rate.



This review paper explores the variety of distributed and parallel computing technologies which are suitable for the research effort to tackle COVID-19, and where examples exist, work carried out using them. This review will not cover machine-learning or deep-learning methods, which although they employ highly-parallel GPU computing, are not necessarily distributed - whilst they are highly relevant to several COVID-19 research areas, it is separate area in its own right.

The objectives of this research are multiple. Firstly, to shed light on and provide a better understanding of the computational work being carried out in research for the COVID-19 pandemic. In particular, the research tasks performed, what distributed computing architectures have been used, how they have been configured to enable the analyses, and the tools and datasets used. Secondly, given the extent and time urgency of the pandemic, to gauge the scale, throughput and parallelism achieved in such work. Thirdly, to examine how existing large-scale computational work, done before the pandemic, can be applied to COVID-19 research. Lastly, to understand the ways in which different distributed platforms can be applied to aspects of the research, for example the use of commodity Hadoop Spark clusters for bigdata processing, and what different distributed computing resources are available for the tasks at hand.

This research has identified the ways in which different types of distributed computing architectures, many of which are state of the art, have been employed for making important contributions to COVID-19 research. Some of the examples covered are: the identification of new pharmacological hits for the S-protein:ACE2 interface, prospective natural product compounds identified through extensive pharmacophore analysis, an unprecedented 0.1s of MD (molecular dynamics) simulation data, aptamer leads to target the RBD (Receptor Binding Domain), and the analysis of over 580 million geo-tagged tweets. The detailed technical discussion of configuration, compute and storage resources, software, and datasets used, provides guidance for implementing such computational workflows. Overall, the paper acts as a road map, providing various routes for using distributed computing for meeting the challenges of such research. It is highly likely that the lessons learned herein are applicable to a variety of other similar research scenarios.

This paper is organised as follows. Firstly, we introduce cluster computing and discuss high-performance computing, particularly the application of high-throughput ensemble docking in identifying pharmacological targets for COVID-19. Next, we examine how some of the world's supercomputers have been applied during the current pandemic, specifically in drug repurposing, high-throughput gene analysis, exploring natural products that could be used to develop COVID-19 leads and in aptamer design. We then discuss how Hadoop and Spark can be applied to COVID-19 bigdata analytics. The wide-ranging applications of grid computing are covered — from forming the world's first exascale distributed computer using large-scale parallel processing, through to docking experiments utilising grid infrastructure, as well as international collaboration. Finally, cloud computing is discussed and a list of distributed computing resources available for COVID-19 researchers is provided.

The distributed computing architectures that are suitable for application to COVID-19 research exist in several different topologies — they can be fundamentally categorised as *clusters*, *grids* and *clouds* (Hussain et al., 2013), and will be covered in the next sections.

## 2. Cluster Computing

Cluster computing, unlike client-server or n-tier architectures — where the focus is on delineating resources group together compute nodes to improve performance through concurrency (Coulouris

et al., 2005). The increasing amount of COVID-19 research to be completed on such systems, coupled with its urgency, will necessitate further performance increases for large-scale projects. They can be achieved by *vertical scaling* —increasing the number of CPU cores in individual compute nodes of the system — or *horizontal scaling* — increasing the number of compute nodes in the system, hence some of the distributed systems employed in the research we review here exhibit both of these characteristics, often at very large-scales.

## 2.1 High-performance Computing with MPI

High-performance computing (HPC) is a key enabling technology for scientific and industrial research (EPSRC, 2016); HPC systems are ubiquitous across scientific and academic research institutions. Most of the computational research projects investigating the structure, function and genome of SARS-CoV-2 will be performed on HPC, executed via computational workflows and pipelines (Alnasir, 2021). This will be predominantly in-house, but in some cases will be via access to external HPC resources, e.g. in collaboration between institutions.

By employing high-bandwidth networking interconnects, HPC clusters facilitate a high degree of inter-process communication and extremely large scalability, for instance the Message Passing Interface (MPI) framework. Software implemented using MPI can exploit for many of the computationally complex problems in COVID-19 research, such as ensemble docking and mathematical modeling (Hill et al., 2000). The use of MPI for distributing complex scientific computation is well-established and many HPC systems are dependent on MPI libraries such OpenMPI and MVAPICH. Consequently, there have been further developments and refinement of these libraries over the last two decades — mainly in reducing latency and memory requirements (Shipman et al., 2006). OpenMPI and MPICH have had their most recent releases in 2020. Recently, new MPI implementations are coming to the fore, such as in LinaLC, a docking program employing strategies such as mixed multi-threading schemes to achieve further performance gains at an extremely large scale, that can be applied to COVID-19 research which we will discuss in the next section.

### 2.1.1 Ensemble Docking

A key task in identifying potential pharmacological agents to target SARS-CoV-2 is molecular docking — *in-silico* simulation of the electrostatic interactions between a ligand and its target — is used to score ligands according to their affinity to the target (Morris and Lim-Wilby, 2008; Meng et al., 2011). The complex computational process is extensively used in drug development and repurposing and is often time-consuming and expensive (Moses et al., 2005; Rawlins, 2004). The protein and enzyme targets that are docked against are not static, but are constantly moving in ways which are dependent on several factors such as temperature, electrostatic attractions and repulsions with nearby molecules, solvation (interaction with the solvent environment) etc. These factors cause atoms in the molecules, within the constraints of the types of bonds the bind them, to adopt spatial arrangements — termed conformations — that correspond to local energy minima on the energy surface. Molecular Dynamics (MD) uses *in-silico* computation to simulate this process, the outcome of which is typically clusters («ensembles») of the most probable conformations for docking, i.e. ensemble docking (Amaro et al., 2018).

In the past, popular tools such as AutoDock Vina — widely-used for performing both molecular docking and virtual screening — were being used primarily on single high-end workstations. Consequently, their parallelism was optimised for multithreading on multicore systems. However,

further gains in such tools have been made by developing or re-implementing existing code for the fine-grained parallelism offered by MPI, and at the same time, leveraging the scale at which HPC systems can operate. In previous work, Zhang et al. have further modified the AutoDock Vina source to implement a mixed MPI and multi-threaded parallel version called VinaLC. They have demonstrated this works efficiently at a very large-scale — 15K CPUs — with an overhead of only 3.94%. Using the DUD dataset (Database of Useful Docking Decoys), they performed 17 million flexible compound docking calculations which were completed on 15,408 CPUs within 24 h. with 70% of the targets in the DUD data set recovered using VinaLC. Projects such as this can be repurposed and applied to identifying potential leads for binding to the SARS-CoV-2 S-protein or the S-protein:Human ACE2 interface, either through the repurposing or the identification of putative ligands (Zhang et al., 2013). Furthermore, given the urgency in finding solutions to the current COVID-19 pandemic — where high-throughput performance gains and extreme scalability are required — these features can be achieved by re-implementing tools in similar ways to which VinaLC has been optimised from the AutoDock codebase.

## 2.2 Supercomputers and COVID-19

### 2.2.1 Drug Repurposing

In recent COVID-19 focused research, Smith et al. have utilised IBM’s SUMMIT supercomputer — the world’s fastest between November 2018 and June 2020 — to perform ensemble docking virtual high-throughput screening against both the SARS-CoV-2 S-protein and the S-protein:Human ACE2 interface (Smith and Smith, 2020; Kerner, S.M., 2018).

SUMMIT, launched by ORNL (Oak Ridge National Laboratory) and based at its Oak Ridge Leadership Computing Facility, comprises 4,608 compute nodes, each with two IBM POWER9 CPUs (containing nine cores each), and six Nvidia Tesla Volta GPUs for a total of 9,216 CPUs and 27,648 GPUs (Vazhkudai et al., 2018). Nodes each have 600 GB of memory, addressable by all CPUs and GPUs, with an additional 800 GB of non-volatile RAM that can be used as a burst buffer or as extended memory. SUMMIT implements a heterogeneous computing model — in each node the two POWER9 CPUs and Nvidia Volta GPUs are connected using Nvidia’s high-speed NVLink. The interconnect between the nodes consist of 200 Gb/s Mellanox EDR Infiniband for both storage and inter-process messaging and supports embedded in-network acceleration for MPI and SHMEM/PGAS.

For a source of ligands, they used the SWEETLEAD dataset which is a highly-curated *in-silico* database of 9,127 chemical structures representing approved drugs, chemical isolates from traditional medicinal herbs, and regulated chemicals, including their stereoisomers (Novick et al., 2013). The work involved three phases of computation: structural modelling, molecular dynamics simulations (ensemble building), and *in-silico* docking. Since the 3D structure of the SARS-CoV-2 S-protein was not yet available during the initial phase of this research, the first phase (structural modelling) was required and a 3D model was built with SWISSMODEL (Schwede et al., 2003) using the sequences for the COVID-19 S-protein (NCBI Ref. Seq: YP\_009724390.1) and the crystal structure of SARS-CoV S-protein as a template to generate the model of the SARS-CoV-2 S-protein:ACE2 complex. In the second phase, molecular dynamics simulations were carried out using GROMACS (compiled on ORNL SUMMIT and run with CHARMM36 force-field (Ossyra et al., 2019; Abraham et al., 2015)) to generate an ensemble of highest probability, lowest energy conformations of the complex which were selected via clustering of the conformations. In the final *in-silico* docking phase, AutoDock Vina was run in parallel using an MPI wrapper.



This work has identified 47 hits for the S-protein:ACE2 interface, with 21 of these having US FDA regulatory approval and 30 hits for the S-protein alone, with 3 of the top hits having regulatory approval.

### 2.2.2 High-throughput and Gene Analysis

Another research project by Garvin et al., that has also been undertaken using SUMMIT, focused on the role of bradykinin and the RAAS (Renin Angiotensin Aldosterone System) in severe, life-threatening COVID-19 symptoms by analysing 40,000 genes using sequencing data from 17,000 bronchoalveolar lavage (BAL) fluid samples (Garvin et al., 2020; Smith, T., 2020). RAAS regulates blood pressure and fluid volume through the hormones renin, angiotensin and aldosterone. Key enzymes in this system are ACE (Angiotensin Converting Enzyme), and ACE2 which work in antagonistic ways to maintain the levels of bradykinin, a nine-amino acid peptide that regulates the permeability of the veins and arterioles in the vascular system. Bradykinin induces hypotension (lowering of blood pressure) by stimulating the dilation of arterioles and the constriction of veins, resulting in leakage of fluid into capillary beds. It has been hypothesised that dysregulation of bradykinin signaling is responsible for the respiratory complications seen in COVID-19 — the *bradykinin storm* (Roche and Roche, 2020).

This work involved massive-scale, gene-by-gene RNA-Seq analysis of SARS-CoV2 patient samples with those of the control samples, using a modified transcriptome. The modified transcriptome was created to allow the researchers to quantify the expression of SARS-CoV2 genes and compare them with the expression of human genes. To create the modified transcriptome, reads from the SARS-CoV2 reference genome were appended to transcripts from the latest human transcriptome, thereby allowing the mapping of reads to the SARS-CoV2 genes. The SUMMIT supercomputer enabled the exhaustive gene-wise tests (all the permutations of all the genes) to be performed at a massive scale in order to test for differential expression, with the Benjamini-Hochberg method applied to the resulting p-values to correct for multiple comparisons.

Their analysis appears to confirm dysregulation of RAAS, as they found decreased expression of ACE together with increased expression of ACE2, renin, angiotensin, key RAAS receptors, and both bradykinin receptors. They also observed increased expression of kininogen and a number of kallikrein enzymes that are kininogen activating — the activated form, kinins, are polypeptides that are involved in vasodilation, inflammatory regulation, and blood coagulation. As they point out, atypical expression levels for genes encoding these enzymes and hormones are predicted to elevate bradykinin levels in multiple tissues and organ systems, and explain many of the symptoms observed in COVID-19.

### 2.2.3 Exploring Natural Products for Treatment

So far, we have discussed some examples of the use of *in-silico* docking and screening that have utilised HPC to identify existing medicines that could potentially be re-purposed for treating COVID-19. However, another strategy — one that is used to develop new therapeutics — explores the chemistry of natural products, i.e. chemical compounds produced by living organisms. To this end, in another research project that also performs *in-silico* docking and screening using a supercomputer, Sentinel, Baudry et al. focus on natural products (Byler et al., 2020). They point out that, natural products, owing to the long periods of natural selection they are subjected to, perform highly selective functions. Their work, therefore, aims to identify pharmacophores (spatial arrangement of chemical functional groups that interact with a specific receptor or target molecular structure) that can be used to develop rationally designed novel medicines to treat COVID-19. In addition to simulating the interaction with the



S-protein RBD, they also included those with the SARS-2 proteome (the sum of the protein products transcribed from its genome), specifically the main protease and the papain-like protease enzymes. These are also important targets as they are highly conserved in viruses and are part of the replication apparatus.

Sentinel, the cluster used for this research, is a Cray XC50, a 48-node, single-cabinet supercomputer featuring a massively parallel multiprocessor architecture and is based in the Microsoft Azure public cloud data centre. It has 1,920 physical Intel Skylake cores operating at 2.4GHz with Hyperthreading (HT) / Simultaneous Multi- Tasking (SMT) enabled, therefore providing 3,840 hyperthreaded CPU cores. Each node has 192 GB RAM and they are connected by an Aries interconnect<sup>28</sup> in a Dragonfly topology. A Cray HPE ClusterStor-based parallel file system is used, providing 612 TB of shared storage that is mounted on every node.

The development of medicines from natural products is challenging for several reasons: supply of the plant and marine organisms, seasonal variation in the organism, extinction of organism sources, and natural products often occur as mixtures of structurally related compounds, even after fractionation, only some of which are active. Contamination, stability, solubility of the compounds, culturing source microorganisms, and cases where synergistic activities require two constituents to be present to display full activity can also present difficulties (Li and Vederas, 2009). Baudry et al., therefore, performed their screening using a curated database of 423,706 natural products, COCONUT (COLleCtion of Open NatUral producTs) (Sorokina, M.; Steinbeck, C., 2020). COCONUT has been compiled from 117 existing natural product databases for which citations in literature since 2000 exist.

Using molecular dynamics simulation coordinate files for the structures of the S-protein, main protease and the papain-like protease enzymes — generated with GROMACS (Abraham et al., 2015) and made available by Oak Ridge National Laboratory — they generated an ensemble using the ten most populated confirmation for each, for *in-silico* docking. As AutoDock was used, its codebase was compiled optimised for Sentinel, with some optimisations for Skylake CPU and memory set in the Makefile compiler options.

They performed pharmacophore analysis of the top 500 unique natural product conformations for each target (S-protein, PL-pro, M-pro). Filtering was applied to the list of putative ligands such that there were no duplicate instances of the same compound, they only occurred in the set for a single target, and were deemed to be drug-like using the MOE (Molecular Operating Environment) descriptor (Vilar et al., 2008) from the COCONUT dataset. This resulted in 232, 204, and 164 compounds for the S-protein, PL-pro, M-pro, respectively. Of these, the top 100 natural products were superimposed onto their respective predicted binding locations on their binding proteins and those that correctly bind to the correct region (i.e. active site) were subjected to for pharmacophoric analysis. For the S-protein, two clusters of 24 and 73 compounds were found to bind to either side of a loop that interacts with ACE2. For PL-pro, the papain-like protease, again two clusters of 40 and 60 compounds were found to bind to either side of a beta-sheet. Finally, for ML-pro, the main protease, five clusters of binding compounds were found, one cluster in in close proximity to the proteases catalytic site.

The common pharmacophores partaking in these interactions were assessed from the relevant clusters, resulting in a greater understanding of the structure-activity relationship of compounds likely to be inhibitory to the SARS-CoV-2 S-protein, PL-pro, and ML-pro proteases. As a result, several natural product leads have been suggested which could undergo further testing and development, and the pharmacophore knowledge could be used to refine existing leads and guide rational drug design for medicines to treat COVID-19.

## 2.2.4 Aptamer Design - the Good Hope Net project

In another response to the COVID-10 pandemic, a collaboration — the Good Hope Net project — was established to develop treatments targeted at Coronavirus. It comprises 21 scientific and research institutions from 8 countries, currently Russia, Finland, Italy, China, Taiwan, Japan, USA and Canada (The Good Hope Net, 2020) and focuses a multi-disciplinary, geographically distributed team of researchers on exploring novel therapeutics for COVID-19. As the development of antibodies is typically laborious, the Good Hope Net's approach concentrates on the development of aptamers — short single-stranded ssDNA or ssRNA that can bind to the proteins and peptides present in SARs-CoV2 — that can be used to synthesise therapeutics. Aptamers offer several advantages over antibodies — they are smaller and hence can be more easily formulated for pulmonary delivery (inhalation via the lungs), and they have high thermal stability facilitating their transportation and storage at room temperature. Furthermore, aptamers may be used in adjunctive therapy with antibodies in order to reduce the incidence of adverse drug reactions (ADRs) (Sun et al., 2021). To this end, the collaboration employs the RSC MVS-10P TORNADO supercomputer based at the Joint Supercomputing Center of the Russian Academy of Sciences (JSCC RAS) in Moscow to perform in-silico docking and MD calculations using X-ray crystallographic models of the SARS-CoV-2 S-protein:ACE2 complex. The MVS-10P TORNADO is built using Xeon E5-2690 8C 2.9GHz CPUs (581 x 2 socket servers) providing 28,704 cores and 16,640 GB (16.6 TB) of RAM with the nodes connected using Infiniband FDR fabric (Savin et al., 2019). RSC Tornado features a 4 tier storage system for Very hot (lowest latency), hot, warm and cold data storage (greatest latency) for compute and Lustre distributed file system.

This work has resulted in 256 putative oligonucleotide aptamer ligands with the most promising aptamer selected for further refinement. The resulting candidate was further improved through sequential matagenesis — the mutation of individual bases at various points in the sequence, in order to improve the binding energy — thereby ensuring the most energetically favourable binding between the aptamer and the SARS-CoV-2 S-protein is achieved. Further study of the binding complex has been carried out using spectral fluorophotometry with fluorescence polarisation (to confirm co-localisation and binding), as well as X-ray crystallographic techniques (which feed back into the *in-silico* simulations) (HPC Wire, 2020).

## 2.3 Hadoop and Spark

Apache Hadoop is an open-source software «ecosystem» comprising a collection of interrelated, interacting projects, distributed platform and software framework that is typically installed on a Linux compute cluster, notwithstanding that it can be installed on a single standalone machine (usually only for the purposes of study or prototyping). Hadoop is increasingly used for bigdata processing (Messerschmitt et al., 2005; Joshua et al., 2013). Bigdata — which we will discuss in more detail with respect to the COVID-19 pandemic — is characterised as data possessing large volume, velocity, variety, value and veracity — known as the v's of bigdata (Laney, 2001; Borgman, 2015). A significant portion of bigdata generated during the COVID-19 pandemic will be semi-structured data from a variety of sources. MapReduce is a formalism for programatically accessing distributed data across Hadoop clusters which store and process data as sets of key-value pairs (i.e. tuples) on which Map and Reduce operations are carried out (Fish et al., 2015). This makes MapReduce particularly useful for processing this semi-structured data and building workflows.

Apache Spark, often viewed as as the successor to Hadoop is a distributed computing framework in its own right, which can be used standalone or can utilise the Hadoop platform's distributed file system





(HDFS), and a resource scheduler — typically Apache YARN. Spark, therefore, can also run MapReduce programs (written in Python, Java, or Scala) (Shanahan and Dai, 2015). It has been designed to overcome the constraints of Hadoop’s acyclic data flow model, through the introduction of a distributed data structure — the Resilient Distributed Dataset (RDD) — which facilitates the re-usability of intermediate data between operations, in-memory caching, and execution optimisation (known as lazy evaluation) for significant performance gains over Hadoop (Zaharia et al., 2012).

## 2.4 COVID-19 Bigdata Analytics

As briefly mentioned earlier, bigdata refers to data sets that, by virtue of their massive size or complexity, cannot be processed or analysed by traditional data-processing methods, and, therefore, usually require the application of distributed, high-throughput computing. Bigdata analytics — the collection of computational methods that are applied for gaining valuable insight from bigdata — employs highly specialised platforms and software frameworks, such as Hadoop or Spark. In a paper that focused on AI for bigdata analytics in infectious diseases, which was written over a year before the current COVID-19 pandemic, Wong et al. point out that, in our current technological age, a variety of sources of epidemiological transmission data exist, such as sentinel reporting systems, disease centres, genome databases, transport systems, social media data, outbreak reports, and vaccinology related data (Wong et al., 2019). In the early stages of global vaccine roll out, compounded by the difficulty of scaling national testing efforts, this data is crucial for contact tracing, and for building models to understand and predict the spread of the disease (Sun et al., 2020).

Furthermore, given the current COVID-19 pandemic has rapidly reached a global scale, the amount of data produced and the variety of sources is even greater than before. Such data is, in most cases, semi-structured or unstructured and, therefore, requires pre-processing (Agbehadji et al., 2020). The size and rate in which this data is being produced during this pandemic, particularly in light of the urgency, necessitates bigdata analytics to realise the potential it has to aid in finding solutions to arrest the spread of the disease by, for example, breaking the chain of transmission (i.e. via track-and-trace systems), and informing government policy (Bragazzi et al., 2020).

The Apache Hadoop ecosystem has a several projects ideally suited to processing COVID-19 big data, and by virtue of them all utilising Hadoop’s cluster infrastructure and distributed file system, they gain from the scalability and fault-tolerance inherent in the framework. For example, for pre-processing bigdata — often referred to as cleaning dirty data — Pig is a high-level data-flow language that can compile scripts into sequences of MapReduce steps for execution on Hadoop (Olston et al., 2008). Apache Spark, owing to its in-memory caching and execution optimisations discussed earlier, offers at least two orders of magnitude faster execution than Hadoop alone and, though centred around MapReduce programming, is less constrained to it. Hive (Thusoo et al., 2009) is a data-warehousing framework which has an SQL type query language, HBase (George, 2011) a distributed scalable database, and Mahout (Lyubimov and Palumbo, 2016) can be used for machine-learning and clustering of data.

An example of how Hadoop can be applied to analytics of COVID-19 big data is shown in recent work by Huang et al. who have analysed 583,748,902 geotagged tweets for the purposes of reporting on human mobility — a causal factor in the spread of the disease (Huang et al., 2020b). In doing so they have demonstrated that bigdata captured from social media can be used for epidemiological purposes and can do so with less invasion of privacy that such data offers. They do point out, however, that a limitation to this approach is that only a small portion of the total twitter corpus is available via



the API. That said, an important outcome of this work is their proposed metric for capturing overall mobility during phases of pandemics — the MRI (Mobility-Responsiveness) Indicator which can be used as a proxy for human mobility.

Whilst Hadoop and Spark are frequently applied to data analytics, they have also been employed in bioinformatics — such as in processing Next-generation sequencing data, e.g. SNP genotyping, de novo assembly, read alignment, reviewed in (Taylor, 2010) and structural biology, e.g. *in-silico* molecular docking, structural alignment / clustering of protein-ligand complexes, and protein analysis reviewed in (Alnasir and Shanahan, 2020).

### 3. Grid Computing

Grids provide a medium for pooling resources and are constructed from a heterogeneous collection of geographically dispersed compute nodes connected in a mesh across the internet or corporate networks. With no centralised point of control, grids broker resources by using standard, open, discoverable protocols and interfaces to facilitate dynamic resource-sharing with interested parties (Foster et al., 2008). Particularly applicable to COVID-19 research are the extremely large scalability grid computing offers and the infrastructure for international collaboration they facilitate, which we will discuss in the following sections.

#### 3.1 Large-scale Parallel-processing Using Grids

The grid architecture allows for massive parallel computing capacity by the *horizontal* scaling of heterogeneous compute nodes, and the exploitation of underutilised resources through methods such as idle CPU-cycle scavenging (Bhavsar and Pradhan, 2009). Distributed, parallel processing using grids is ideally suited for batch tasks that can be executed remotely without any significant overhead.

An interesting paradigm of grid computing, that has now been applied to Molecular Dynamics research for COVID-19, leverages this scalability, particularly for applications in scientific computing, is known as volunteering distributed computing having evolved during the growth of the internet from the 2000s onwards. This involves allocating work to volunteering users on the internet (commonly referred to as the @home projects) with tasks typically executed while the user's machine is idle (Krieger and Vriend, 2002).

#### 3.2 The World's First Exascale Computer Assembled Using Grid Volunteer Computing

A recent project, that focused on simulating the conformations adopted by the SARS-CoV-2 S-protein, culminated in the creation of the first Exascale grid computer. This was achieved by enabling over a million citizen scientists to volunteer their computers to the Folding@home grid computing platform, which was first founded in 2000 to understand protein dynamics in function and dysfunction (Zimmerman et al., 2020; Beberg et al., 2009). The accomplishment of surmounting the Exascale barrier by this work is based on a conservative estimate that the peak performance of 1.01 exaFLOPS on the Folding@home platform was achieved at a point when 280,000 GPUs and 4.8 million CPU cores were performing simulations. The estimate counts the number of GPUs and CPUs that participated during a three-day window, and makes the conservative assumption about the computational



performance of each device. Namely, that each GPU/CPU participating has worse performance than a card released before 2015.

In addition to understanding how the structure of the SARS-CoV-2 S-protein dictates its function, simulating the ensemble of conformations that it adopts allows characterisation of its interactions. These interactions with the ACE2 target, the host system antibodies, as well as glycans on the virus surface, are key to understanding the behaviour of the virus. However, as pointed out in this work, datasets generated by MD simulations typically consist of only a few microseconds of simulation — at most millisecond timescales — for a single protein. An unprecedented scale of resources are therefore required to perform MD simulations for all of the SARS-CoV-2 proteins. The Folding@home grid platform has enabled this, generating a run of 0.1 s of simulation data that illuminates the movement and conformations adopted by these proteins over a biologically relevant time-scale.

### 3.3 International Collaboration Through Grids

Grids are an ideal infrastructure for hosting large-scale international collaboration. This was demonstrated by the Globus Toolkit produced by the Global Alliance, which became a de facto standard software for grids deployed in scientific and industrial applications. It was designed to facilitate global sharing of computational resources, databases and software tools securely across corporate and institutions (Ferreira et al., 2003). However, development of the toolkit ended in 2018 due to a lack of funding and the service remains under a freemium mode. Globus’s work in enabling worldwide collaboration continue through their current platform which now employs cloud computing to provide services — this is discussed further in section 4.

Some notable large-scale grids participating in COVID-19 research are: the World Community Grid launched by IBM (IBM, 2020), the WLCG (Worldwide LHC Computing Grid) at CERN (Sciaba et al., 2010), Berkeley Open Infrastructure for Network Computing (BOINC) (Anderson, 2004), the European Grid Infrastructure (EGI) (Gagliardi, 2004), the Open Science Grid (OSG) (Pordes et al., 2007) and previously Globus. Interestingly, grids can be constructed from other grids — for example, BOINC is part of IBM WCG, and CERN’s WLCG is based on two main grids, the EGI and OSG, which is based in the US.

### 3.4 COVID-19 Research on Genomics England Grid

Genomics England, founded in 2014 by the UK government and owned by the Department of Health & Social Care, has been tasked with delivering the 100,000 genomes project which aims to study the genomes of patients with cancer or rare diseases (Siva, 2015; James Gallagher, BBC, 2014). It was conceived at a time when several government and research institutions worldwide announced large-scale sequencing projects — akin to an *arms race* of sequencing for patient-centric precision medicine research. In establishing the project, the UK government and Illumina decided to secure sequencing services for the project from Illumina (Marx, 2015). Sequencing of the 100,000 genomes has resulted in 21 PB of data and involved 70,000 UK patients and family members, 13 genomic medicines centres across 85 recruiting NHS trusts, 1,500 NHS staff, and 2,500 researchers and trainees globally (Genomics England, 2014).

In 2018, after sequencing of the 100,000 genomes was completed, the UK government announced the significant expansion of the project — to sequence up to five million genomes over five years (Genomics England, 2018). At the time, the Network Attached Storage (NAS) held 21 PB of data



and had reached its node-scaling limit and so a solution that could scale to hundreds of Petabytes was needed — after consultation with Nephos Technologies, a more scalable storage system comprising a high-performance parallel file system from WekaIO, Mellanox® high-speed networking, and Quantum ActiveScale object storage was implemented (HPC wire, 2020). Genomics England’s Helix cluster, recently commissioned in 2020, has 60 compute nodes each with 36 cores (providing 2,160 cores) and approximately 768 GB RAM. It has a dedicated GPU node with 2x nVidia Tesla V100 GPUs installed (Genomics England, 2020).

GenOMICC (Genetics Of Mortality In Critical Care) is a collaborative project, first established in 2016, to understand and treat critical illness such as sepsis and emerging infections (e.g. SARS/MERS/Flu) is now also focusing on the COVID-19 pandemic. The collaboration involves Genomics England, ISARIC (The International Severe Acute Respiratory and Emerging Infection Consortium), InFACT (The International Federation of Acute Care Trialists), Asia-Pacific Extra-Corporeal Life Support Organisation (AP ELSO) and the Intensive Care Society. The aim is to recruit 15,000 participants for genome sequencing, who have experienced only mild symptoms, i.e. who have tested positive for COVID-19, but have not been hospitalised. The rationale is that in addition to co-morbidities, there are genetic factors that determine whether a patient will suffer mild or severe, potentially life-threatening illness — this would also explain why some young people, who are fit and healthy have suffered severely and others who are old and frail did not. Furthermore, since many people who have suffered severe illness from COVID-19 were elderly or from ethnic minorities, the aim is to recruit participants that are from these backgrounds who suffered from mild symptoms of COVID-19. To this end, the project will carry out GWAS (Genome Wide Association Studies) to identify associations between genetic regions (loci) and increased susceptibility to COVID-19 (Pairo-Castineira et al., 2020).

### 3.5 Other COVID-19 Research on Grids

During the previous 2002-4 SARS-CoV-1 outbreak, DiscoveryNet — a pilot designed and developed at Imperial College and funded by the UK e-Science Programme — enabled a collaboration between its team and researchers from SCBIT (Shanghai Centre for Bioinformation Technology) to analyse the evolution of virus strains from individuals of different countries (Au et al., 2004). This was made possible through its provision of computational workflow services, such as an XML based workflow language and the ability to couple workflow process to datasources, as part of an e-Science platform to facilitate the extraction of knowledge from data (KDD) (Rowe et al., 2003). It is coincidental that this grid technology in its infancy, and in its pilot phase, was used in a prior pandemic, especially since many of the services will be employed during the current one, in particular the support for computational workflows and the use of large datasets made available through grids and clouds.

In work that utilises various grid resources, including EGI and OSG, together with the European Open Science Cloud (EOSC) (Ayrís et al., 2016), Hassan et al. have performed an *in-silico* docking comparison between human COVID-19 patient antibody (B38) and RTA-PAP fusion protein (ricin a chain-pokeweed antiviral protein) against targets (S-protein RBD, Spike trimer, and membrane-protein) in SARS-CoV-2 (Hassan et al., 2020). RTA-PAP, plant-derived N-glycosidase ribosomal-inactivating proteins (RIPs), is a fusion of ricin a chain isolated from *Ricinus communis* — and pokeweed antiviral protein — isolated from *Phytolacca Americana*, which the same researchers had demonstrated to be anti-infective against Hepatitis B in prior work (Hassan et al., 2018). They also utilised a grid based service called WeNMR, which provides computational workflows for NMR (Nuclear Magnetic Resonance)/SAX (Small-angle X-ray scattering) via easy-to-use web interfaces (Wassenaar



et al., 2012), and the CoDockPP protein-protein software to perform the docking (Kong et al., 2019). They found favourable binding affinities (low binding energies) for the putative fusion protein RTA-PAP binding with both the SARS-CoV-2 S-protein trimer and membrane protein, which can be further explored for development as antivirals for use against COVID-19.

## 4. Cloud Computing

A consequence of the data-driven, integrative nature of bioinformatics and computational biology (Dudley et al., 2010), as well as advancements in high-throughput next-generation sequencing (Dai et al., 2012), is that cloud-services such as for instance Amazon AWS, Microsoft Azure, and Google Cloud, are increasingly being used in research (Schatz et al., 2010; Shanahan et al., 2014). These areas of research underpin the COVID-19 research effort and hence the use of cloud services will no doubt contribute significantly to the challenges faced.

Clouds provide pay-as-you-go access to computing resources via the internet through a service provider and with minimal human interaction between the user and service provider. Resources are accessed on-demand, generically as a service, without regard for physical location or specific low-level hardware and in some cases without software configuration (Smith and Nair, 2005). This has been made possible by the developments in virtualisation technologies such as Xen, and Windows Azure Hypervisor (WAH) (Younge et al., 2011; Barham et al., 2003). Services are purchased on-demand in a metered fashion, often to augment local resources and aid in completion of large or time-critical computing tasks. This offers small research labs access to infrastructure that they would not be able to afford to invest in for use on-premises, as well as services that would be time consuming and costly to develop (Navale and Bourne, 2018). Furthermore, there is variety in the types of resources provided in the form of different service models offered by cloud providers, such as Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) (Mell et al., 2011).

To a great extent, scientific and bioinformatics research projects utilise cloud services through IaaS (Infrastructure as a Service) and PaaS (Platform as a Service). In the IaaS approach, processing, storage and networking resources are acquired and leased from the service provider and configured by the end user to be utilised through the use of virtual disk images. These virtual disks are provided in proprietary formats, for instance the AMI (Amazon Machine Image) on AWS or VHD (Virtual Hard Disk) on Azure, serve as a bit-for-bit copy of the state of a particular VM (Shanahan et al., 2014). They are typically provisioned by the service provider with an installation of commonly used Operating Systems configured to run on the cloud service's Infrastructure, and service providers usually offer a selection of such images. This allows the end user to then install and precisely configure their own or third party software, save the state of the virtual machine, and deploy the images elsewhere.

In contrast, in the PaaS approach, the end user is not tasked with low level configuration of software and libraries which are instead provided to the user readily configured to enable rapid development and deployment to the cloud. For example AWS provides a PaaS for MapReduce called Elastic MapReduce which it describes as a «Managed framework for Hadoop that makes it easy, fast, and cost-effective to process vast amounts of data across dynamically scalable Amazon EC2 instances» (Amazon, 2016). In fact MapReduce is offered as a PaaS by all of the major cloud-service providers (Amazon AWS, Google Cloud and Microsoft Azure) (Gunaratne et al., 2010).

Cloud computing is market-driven and has emerged thanks to improvements in capabilities the internet which have enabled the transition of computational research from mainstay workstations and

HPC clusters into the cloud. Clouds offer readily provisionable resources which, unlike grids — where investment can be lost when scaled down — projects utilising cloud infrastructure do not suffer the same penalty. However, there is no up-front cost to establishing infrastructure in the case of clouds. One potential drawback is that whilst the ingress of data into clouds is often free, there is invariably a high cost associated with data egress which, depending on the size of the computational results, can make it more costly than other infrastructures in terms of extricating computational results (Navale and Bourne, 2018). These are salient factors with respect to the likely short-term duration of some of the pandemic research tasks that are being carried out.

## 4.1 International Collaboration through Clouds

As discussed in earlier (section 3.3), the Globus service has evolved from the Globus Alliance grid consortium’s work on the standardisation and provision of grid services. Currently, Globus is a cloud-enabled platform that facilitates collaborative research through the provision of services which focus primarily on data management. Globus is used extensively by the global research community, and at the time of writing, there are over 120,000 users across more than 1,500 institutions registered and connected by more than 30,000 global *endpoints*. Endpoints are logical file transfer locations (source or destination) that are registered with the service and represent a resource (e.g. a server, cluster, storage system, laptop, etc.) between which files can be securely transferred by authorised users. Globus has enabled the transfer of more than 800 PB of data to-date — presently more than 500 TB are transferred daily. Some of the services it provides are given in Table 1 below.

Another cloud project to enable research collaboration, currently still in development, is the European Open Science Cloud (EOSC), which was proposed in 2016 by the European Commission with a vision to enable Open Science (Ayrís et al., 2016). It aims to provide seamless cloud-services for storage, data management and analysis and facilitate re-use of research data by federating existing scientific infrastructures dispersed across EU member states. After an extensive consultation period with scientific and institutional stakeholders, the outcome is a road-map of project milestones, published in 2018, these are: I) Architecture, II) Data, III) Services, IV) Access and Interface, and V) Rules and Governance — and are anticipated to be completed by 2021.

*Table 1. Some important services Globus provides*

Feature	Description
<b>Identity management</b>	Authentication and authorization interactions are brokered between end-users, identity providers, resource servers (services), and clients
<b>File transfers</b>	Can be performed securely, either by request or automated via script
<b>File sharing</b>	Allows sharing between users, groups, and setting access permissions
<b>Workflow automation</b>	Automate workflow steps into pipelines
<b>Dataset assembly</b>	Researchers can develop and deposit datasets, and describe their attributes using domain-specific metadata
<b>Publication repository</b>	Curators review, approve and publish data
<b>Collaboration</b>	Collaborators can access shared files via Globus login — no local account is required — and then download
<b>Dataset discovery</b>	Peers and collaborators can search and discover datasets



## 5. Distributed Computing Resources provided freely to COVID-19 Researchers

In order to facilitate and accelerate COVID-19 research, a number of organisations are offering distributed computational resources freely to researchers (Miller, S., 2020). For instance, a number of research institutions that extensively use HPC — many of which host the world’s most powerful supercomputers — have joined together to form the COVID-19 High Performance Computing Consortium. Cloud providers such as Amazon AWS, Google, Microsoft and Rescale are also making their platforms available, generally through the use of computational credits, and are largely being offered to researchers working on COVID-19 diagnostic testing and vaccine research. Table 2 lists some of the computational resources on offer, and the specific eligibility requirements for accessing them.

## 6. Conclusions

There are a variety of distributed architectures that can be employed to perform efficient, large-scale, and highly-parallel computation requisite for several important areas of COVID-19 research. Some of the large-scale COVID-19 research projects we have discussed that utilise these technologies are summarised in Table 3 — these have focused on *in-silico* docking, MD simulation and gene-analysis.

High-performance computing (HPC) clusters are ubiquitous across scientific research institute and aggregate their compute nodes using high-bandwidth networking interconnects. Employing communications protocols, such as Message Passing Interface (MPI), they enable software to achieve a high degree of inter-process communication. Hadoop and Spark facilitate high-throughput processing suited for the bigdata tasks in COVID- 19 research. Even when Hadoop/Spark clusters are built using commodity hardware, their ecosystem of related software projects can make use of the fault-tolerant, scalable Hadoop framework i.e. HDFS distributed file system — features that are usually found in more expensive HPC systems. Although not widely adopted, nor a common use, Hadoop and Spark have also been employed for applications in bioinformatics (e.g. processing sequencing data) and structural biology (e.g. performing docking, clustering of protein-ligand conformations).

### Key points

- HPC is commonly used in research institutions. However, access to the world’s supercomputers allows for the largest scale projects to be completed quicker, which is particularly important given the time urgency of COVID-19 research.
- Bigdata generated during the pandemic - which can be used for epidemiological modeling and critical track and trace systems - can be processed using platforms such as Spark and Hadoop.
- Grid computing platforms offer unprecedented computing power through volunteer computing, enabling large-scale analysis during the pandemic that hitherto has not been achieved at this scale.
- Both grids and clouds can also be used for international research collaboration by providing services, frameworks and APIs, but differ in their geographical distribution and funding models.

COVID-19 research has utilised some of the world’s fastest supercomputers, such as IBM’s SUMMIT — to perform ensemble docking virtual high-throughput screening against SARS-CoV-2 targets



Table 2. Free HPC and Cloud-computing resources for COVID-19 researchers

Provider / Initiative	Offering	Eligibility
(COVID-19 HPC Consortium, 2020) COVID-19 High Performance Computing Consortium	<p>Access to global supercomputers at the institutions taking part in the consortium, such as:</p> <ul style="list-style-type: none"> <li>• Oak Ridge Summit</li> <li>• Argonne Theta</li> <li>• Lawrence Berkeley National Laboratory Cori</li> <li>• and many more</li> </ul> <p>Also, other resources contributed by members, such as:</p> <ul style="list-style-type: none"> <li>• IBM, HP, Dell, Intel, nVidia</li> <li>• Amazon, Google</li> <li>• National infrastructures (UK, Sweden, Japan, Korea, etc)</li> <li>• and many others</li> </ul>	<p>Requests need to demonstrate:</p> <ul style="list-style-type: none"> <li>• Potential near-term benefits for COVID-19 response</li> <li>• Feasibility of the technical approach</li> <li>• Need for HPC</li> <li>• HPC knowledge and experience of the proposing team</li> <li>• Estimated computing resource requirements</li> </ul>
(Amazon AWS, 2020b) Tech against COVID: Rescale partnership with Google Cloud and Microsoft Azure	HPC resources through the Rescale platform	Any researcher, engineer, or scientist can apply who is targeting their work to combat COVID-19 in developing test kits and vaccines.
(Amazon AWS, 2020a) AWS Diagnostic Development Initiative	AWS in-kind credits and technical support.	<p>Accredited research institutions or private entities:</p> <ul style="list-style-type: none"> <li>• a using AWS to support research-oriented workloads for the development of point-of-care diagnostics</li> <li>• other COVID-19 infectious disease diagnostic projects considered</li> </ul>
(Lifebit, 2020) Lifebit	Premium license for Lifebit CloudOS	Exact eligibility criteria not published, but is advertised for researchers developing diagnostics, treatments, and vaccines for COVID-19. Contact lifebit with details of project.

for drug-repurposing, and high-throughput gene analysis — RSC’s TORNADO for the design of aptamers, and Sentinel, an XPE-Cray based system used to explore natural products. During the present COVID-19 pandemic, researchers working on important COVID-19 problems, who have relevant experience, now have expedited and unprecedented access to supercomputers and other powerful resources through the COVID-19 High Performance Computing Consortium. Grid computing has also come to the fore during the pandemic by enabling the formation of an Exascale grid computer allowing massively-parallel computation to be performed through volunteer computing using the Folding@home platform.

Jamie J. Alnasir

Distributed Computing in a Pandemic: A Review of Technologies available for Tackling COVID-19



ADCAIJ: Advances in Distributed Computing and Artificial Intelligence Journal  
 Regular Issue, Vol. 11 N. 1 (2022), 19-43  
 eISSN: 2255-2863 - <https://adcaij.usal.es>  
 Ediciones Universidad de Salamanca - CC BY-NC-ND



Table 3. Comparison of COVID-19 research exploiting large-scale distributed computing

Ref. / Name	Platform	Scale	Research task	Tools	Outcome
Smith and Smith, 2020	IBM Summit supercomputer	up to 4,608 nodes, 9,216 CPUs, 27,648 GPUs	<i>in-silico</i> ensemble docking & screening of existing medicines for repurposing	GROMACS, CHARMM32, AutoDock Vina	Identified 47 hits for the S- protein:ACE2 interface, with 21 of these having US FDA regulatory approval. 30 hits for the S-protein alone, with 3 of the top hits having regulatory approval.
Garvin et al., 2020	IBM Summit supercomputer	up to 4,608 nodes, 9,216 CPUs, 27,648 GPUs	large-scale gene analysis	AutoDock	Observed atypical expression levels for genes in RAAS pointing to bradykinin dysregulation and storm hypothesis.
Byler et al., 2020	Cray Sentinel supercomputer	up to 48 nodes, 1,920 physical cores 3,840 HT/SMT cores	<i>in-silico</i> docking	AutoDock	Pharmacophore analysis of natural product compounds likely to be inhibitory to the SARS-CoV-2 Sprotein, PL-pro, and ML-pro proteases.
Zimmerman et al., 202)	Folding@ home grid	4.8 million CPU cores ~280,000 GPUs	MD simulations	GROMACS, CHARMM36, AMBER03	Generated an unprecedented 0.1 s of MD simulation data.
Hassan et al., 2020	EGI, OSD grids & EOSC cloud	<i>unspecified</i>	<i>in-silico</i> docking	weNMR, CoDockPP	Demonstrated high in-silico binding affinities of fusion protein RTA-PAP putative ligand with both the SARS-CoV-2 S-protein trimer and membrane protein.
Pauro-Castineira et al., 2020	Genomics England grid, Helix cluster	up to 60 nodes (2,160 cores), 2x V100 GPUs	GWAS	Not yet specified	Recruitment of 15,000 participants is ongoing.
The Good Hope Net, 2020	RSC MVS-10P TORNADO supercomputer	up to 28,704 cores, 16,640 GB RAM	<i>in-silico</i> ensemble docking & MD simulations for aptamer development	Unspecified	A group of 256 putative oligonucleotide aptamer leads were generated, resulting in the best one selected for further refinement to target the Coronavirus S-protein RBD.

(Continued)

Table 3. Comparison of COVID-19 research exploiting large-scale distributed computing (Continued)

Ref. / Name	Platform	Scale	Research task	Tools	Outcome
Huang et al., 2020b	On-premises Hadoop cluster	13 Hadoop nodes <sup>1</sup>	Twitter analytics	Hive, Impala	Analysis of over 580 million global geo-tagged tweets demonstrated that twitter data is amenable to quantitatively assess user mobility for epidemiological study, particularly in response to periods of the pandemic and government announcements on mitigating measures. Metric proposed: MRI (Mobility-based Response Index) to act as proxy for human movement.

<sup>1</sup>Although the size of the infrastructure in this project is small, the dataset represents a large-scale study.

Table 4. Comparison of distributed computing architectures (mateescu2011hybrid)

Feature	HPC	Grid	Cloud
Capacity	fixed	average to high; growth by aggregating independently managed resources	high, growth by elasticity of commonly managed resources
Capability	very high	average to high	low to average
VM support	rarely	sometimes	always
Resource sharing	limited	high	limited
Resource heterogeneity	low	average to high	low to average
Workload management	yes	yes	no
Interoperability	n/a	average	low
Security	high	average	low to average

Grids and clouds provide services such as Globus provide a variety of services, for example, reliable file transfer, workflow automation, identity management, publication repositories, and dataset discovery, thereby allowing researchers to focus on research rather than on time-consuming data-management tasks. Furthermore, cloud providers such as AWS, Google, Microsoft and Rescale are offering free credits for COVID-19 researchers.

In the near future, we will be able to assess the ways in which distributed computing technologies have been deployed to solve important problems during the COVID-19 pandemic and we will no doubt learn important lessons that are applicable to a variety of scenarios.

## 7. Acknowledgements

The author wishes to thank Eszter Ábrahám for proofreading the manuscript.

## 8. Conflicts of Interest Statement

The author declares no conflicts of interest.

## References

- Abraham, M. J., Murtola, T., Schulz, R., Páll, S., Smith, J. C., Hess, B., and Lindahl, E., 2015. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX*, 1:19–25.
- Agbehadji, I. E., Awuzie, B. O., Ngowi, A. B., and Millham, R. C., 2020. Review of Big Data Analytics, Artificial Intelligence and Nature-Inspired Computing Models towards Accurate Detection of COVID-19 Pandemic Cases and Contact Tracing. *International journal of environmental research and public health*, 17(15):5330.
- Alnasir, J. J., 2021. Fifteen quick tips for success with HPC, ie, responsibly BASHing that Linux cluster. *PLOS Computational Biology*, 17(8):e1009207.
- Alnasir, J. J. and Shanahan, H. P., 2020. The application of hadoop in structural bioinformatics. *Briefings in bioinformatics*, 21(1):96–105.
- Amaro, R. E., Baudry, J., Chodera, J., Demir, Ö., McCammon, J. A., Miao, Y., and Smith, J. C., 2018. Ensemble docking in drug discovery. *Biophysical journal*, 114(10):2271–2278.
- Amazon, 2016. Amazon EMR (Elastic MapReduce). <https://aws.amazon.com/emr/>. [Online; accessed 14-April-2019].
- Amazon AWS, 2020a. COVID researchers can apply for free cloud services. <https://aws.amazon.com/government-education/nonprofits/disaster-response/diagnostic-dev-initiative/>. [Online; accessed 01-April-2020].
- Amazon AWS, 2020b. Tech Against COVID: Rescale and Microsoft Azure donate supercomputing resources to help researchers combat global pandemic. <https://partner.microsoft.com/ru-ru/case-studies/rescale>. [Online; accessed 01-April-2020].
- Anderson, D. P., 2004. Boinc: A system for public-resource computing and storage. In *Fifth IEEE/ACM international workshop on grid computing*, pages 4–10. IEEE.
- Au, A., Curcin, V., Ghanem, M., Giannadakis, N., Guo, Y., Jafri, M., Osmond, M., Oleynikov, A., Rowe, A., Syed, J. et al., 2004. Why grid-based data mining matters? fighting natural disasters on the grid: from SARS to land slides. In *UK e-science all-hands meeting (AHM 2004), Nottingham, UK*, pages 121–126.
- Ayris, P., Berthou, J.-Y., Bruce, R., Lindstaedt, S., Monreale, A., Mons, B., Murayama, Y., Södergård, C., Tochtermann, K., and Wilkinson, R., 2016. Realising the european open science cloud.

- Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., and Warfield, A., 2003. Xen and the art of virtualization. In *ACM SIGOPS operating systems review*, volume 37, pages 164–177. ACM.
- Beberg, A. L., Ensign, D. L., Jayachandran, G., Khaliq, S., and Pande, V. S., 2009. Folding@ home: Lessons from eight years of volunteer distributed computing. In *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–8. IEEE.
- Bhavsar, M. D. and Pradhan, S. N., 2009. Scavenging idle CPU cycles for creation of inexpensive supercomputing power. *International Journal of Computer Theory and Engineering*, 1(5):602.
- Borgman, C. L., 2015. *Big Data, little data, no data: Scholarship in the networked world*. Mit Press.
- Bragazzi, N. L., Dai, H., Damiani, G., Behzadifar, M., Martini, M., and Wu, J., 2020. How Big Data and Artificial Intelligence Can Help Better Manage the COVID-19 Pandemic. *International Journal of Environmental Research and Public Health*, 17(9):3176.
- Byler, K., Landman, J., and Baudry, J., 2020. High Performance Computing Prediction of Potential Natural Product Inhibitors of SARS-CoV-2 Key Targets.
- Coulouris, G. F., Dollimore, J., and Kindberg, T., 2005. *Distributed systems: concepts and design*. pearson education.
- COVID-19 HPC Consortium, 2020. COVID researchers can apply for free cloud services. <https://www.xsede.org/covid19-hpc-consortium>. [Online; accessed 01-April-2020].
- Dai, L., Gao, X., Guo, Y., Xiao, J., and Zhang, Z., 2012. Bioinformatics clouds for Big Data manipulation. *Biology direct*, 7(1):43.
- Dong, E., Du, H., and Gardner, L., 2020. An interactive web-based dashboard to track COVID-19 in real time. *The Lancet infectious diseases*, 20(5):533–534.
- Dudley, J. T., Pouliot, Y., Chen, R., Morgan, A. A., and Butte, A. J., 2010. Translational bioinformatics in the cloud: an affordable alternative. *Genome medicine*, 2(8):51.
- EPSRC, 2016. An analysis of the impacts and outputs of investment in national HPC. <https://epsrc.ukri.org/newsevents/pubs/impactofnationalhpc/>. [Online; accessed 07-September-2020].
- Ferreira, L., Berstis, V., Armstrong, J., Kendzierski, M., Neukoetter, A., Takagi, M., Bing-Wo, R., Amir, A., Murakawa, R., Hernandez, O. et al., 2003. Introduction to grid computing with globus. *IBM redbooks*, 9.
- Ferretti, L., Wymant, C., Kendall, M., Zhao, L., Nurtay, A., Abeler-Dörner, L., Parker, M., Bonsall, D., and Fraser, C., 2020. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science*, 368(6491).
- Fish, B., Kun, J., Lelkes, A. D., Reyzin, L., and Turán, G., 2015. On the computational complexity of mapreduce. In *International Symposium on Distributed Computing*, pages 1–15. Springer.
- Foster, I., Zhao, Y., Raicu, I., and Lu, S., 2008. Cloud computing and grid computing 360-degree compared. In *Grid Computing Environments Workshop, 2008. GCE'08*, pages 1–10. Ieee.
- Gagliardi, F., 2004. The EGEE European grid infrastructure project. In *International Conference on High Performance Computing for Computational Science*, pages 194–203. Springer.
- Garvin, M. R., Alvarez, C., Miller, J. I., Prates, E. T., Walker, A. M., Amos, B. K., Mast, A. E., Justice, A., Aronow, B., and Jacobson, D., 2020. A mechanistic model and therapeutic interventions for COVID-19 involving a RAS-mediated bradykinin storm. *Elife*, 9:e59177.

- Genomics England, 2014. 100,000 Genomes project by numbers. <https://www.genomicsengland.co.uk/the-100000-genomes-project-by-numbers/>. [Online; accessed 24-November-2019].
- Genomics England, 2018. Secretary of State for Health and Social Care announces ambition to sequence 5 million genomes within five years. <https://www.genomicsengland.co.uk/matt-hancock-announces-5-million-genomes-within-five-years/>. [Online; accessed 30-October-2020].
- Genomics England, 2020. Genomics England Research Environment - HPC (Helix) Migration 2020. [https://cnfl.extge.co.uk/display/GERE/HPC+%28Helix%29+Migration+2020#HPC\(Helix\)Migration2020-ChangeToThePhysicalComputeNodes](https://cnfl.extge.co.uk/display/GERE/HPC+%28Helix%29+Migration+2020#HPC(Helix)Migration2020-ChangeToThePhysicalComputeNodes). [Online; accessed 22-September-2020].
- George, L., 2011. HBase: The Definitive Guide: Random Access to Your Planet-Size Data. «O'Reilly Media, Inc.».
- Gunarathne, T., Wu, T.-L., Qiu, J., and Fox, G., 2010. MapReduce in the Clouds for Science. In *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*, pages 565–572. IEEE.
- Hassan, Y., Ogg, S., and Ge, H., 2018. Expression of novel fusion antiviral proteins ricin A chain-pokeweed antiviral proteins (RTA-PAPs) in *Escherichia coli* and their inhibition of protein synthesis and of hepatitis B virus in vitro. *BMC biotechnology*, 18(1):47.
- Hassan, Y., Ogg, S., and Ge, H., 2020. Novel anti-SARS-CoV-2 mechanisms of fusion broad range anti-infective protein ricin A chain mutant-pokeweed antiviral protein 1 (RTAM-PAP1) in silico.
- Hilgenfeld, R., 2014. From SARS to MERS: crystallographic studies on coronaviral proteases enable antiviral drug design. *The FEBS journal*, 281(18):4085–4096.
- Hill, M. D., Jouppe, N. P., and Sohi, G., 2000. *Readings in computer architecture*. Gulf Professional Publishing.
- HPC wire, 2020. Genomics England Scales Up Genomic Sequencing with Quantum ActiveScale Object Storage. <https://www.hpcwire.com/off-the-wire/genomics-england-scales-up-genomic-sequencing-with-quantum-activescale-object-storage/>. [Online; accessed 22-September-2020].
- HPC Wire (2020). The Good Hope Net Project and Russian Supercomputer Achieve New Milestone in COVID-19 Fight. <https://www.hpcwire.com/off-the-wire/the-good-hope-net-project-and-russian-supercomputer-achieve-new-milestone-in-covid-19-fight/>. [Online; accessed 08-August-2021].
- Huang, C., Wang, Y., Li, X., Ren, L., Zhao, J., Hu, Y., Zhang, L., Fan, G., Xu, J., Gu, X. et al., 2020a. Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. *The lancet*, 395(10223):497–506.
- Huang, X., Li, Z., Jiang, Y., Li, X., and Porter, D., 2020b. Twitter reveals human mobility dynamics during the COVID-19 pandemic. *PLoS one*, 15(11):e0241957.
- Hussain, H., Malik, S. U. R., Hameed, A., Khan, S. U., Bickler, G., Min-Allah, N., Qureshi, M. B., Zhang, L., Yongji, W., Ghani, N. et al., 2013. A survey on resource allocation in high performance distributed computing systems. *Parallel Computing*, 39(11):709–736.
- IBM (2020). IBM World Community Grid - about page. [https://www.worldcommunitygrid.org/about\\_us/viewAboutUs.do](https://www.worldcommunitygrid.org/about_us/viewAboutUs.do). [Online; accessed 16-September-2020].
- James Gallagher, BBC, 2014. DNA project 'to make UK world genetic research leader'. <http://www.bbc.co.uk/news/health-28488313>. [Online; accessed 21-January-2019].
- Joshua, J., Alao, D., Okolie, S., and Awodele, O., 2013. Software Ecosystem: Features, Benefits and Challenges.

- Kalil, A. C., 2020. Treating COVID-19—off-label drug use, compassionate use, and randomized clinical trials during pandemics. *Jama*, 323(19):1897–1898.
- Kerner, S.M., 2018. IBM Unveils Summit, the World's Fastest Supercomputer (For Now). <https://www.serverwatch.com/server-news/ibm-unveils-summit-the-worlds-faster-supercomputer-for-now.html>. [Online; accessed 07-September-2020].
- Kissler, S. M., Tedijanto, C., Goldstein, E., Grad, Y. H., and Lipsitch, M., 2020. Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. *Science*, 368(6493):860–868.
- Kong, R., Wang, F., Zhang, J., Wang, F., and Chang, S., 2019. CoDockPP: A Multistage Approach for Global and Site-Specific Protein–Protein Docking. *Journal of chemical information and modeling*, 59(8):3556–3564.
- Krieger, E. and Vriend, G., 2002. Models@ Home: distributed computing in bioinformatics using a screensaver based approach. *Bioinformatics*, 18(2):315–318.
- Kuba, K., Imai, Y., Rao, S., Gao, H., Guo, F., Guan, B., Huan, Y., Yang, P., Zhang, Y., Deng, W. et al., 2005. A crucial role of angiotensin converting enzyme 2 (ACE2) in SARS coronavirus–induced lung injury. *Nature medicine*, 11(8):875–879.
- Lake, M. A., 2020. What we know so far: COVID-19 current clinical knowledge and research. *Clinical Medicine*, 20(2):124.
- Laney, D., 2001. 3D Data Management: Controlling Data Volume, Velocity, and Variety. Technical report, META Group.
- Li, J. W.-H. and Vederas, J. C., 2009. Drug discovery and natural products: end of an era or an endless frontier? *Science*, 325(5937):161–165.
- Li, W., Moore, M. J., Vasilieva, N., Sui, J., Wong, S. K., Berne, M. A., Somasundaran, M., Sullivan, J. L., Luzuriaga, K., Greenough, T. C. et al., 2003. Angiotensin-converting enzyme 2 is a functional receptor for the SARS coronavirus. *Nature*, 426(6965):450–454.
- Lifebit, 2020. Lifebit Provides Free Cloud Operating System, Data Hosting & Analysis Tools to COVID-19 Researchers. <https://blog.lifebit.ai/2020/03/30/lifebit-provides-free-cloud-operating-system-data-hosting-analysis-tools-to-covid-19-researchers> [Online; accessed 01-April-2020].
- Lu, H., Stratton, C. W., and Tang, Y.-W., 2020. Outbreak of pneumonia of unknown etiology in Wuhan, China: The mystery and the miracle. *Journal of medical virology*, 92(4):401–402.
- Lyubimov, D. and Palumbo, A., 2016. *Apache Mahout: Beyond MapReduce*. CreateSpace Independent Publishing Platform.
- Marx, V., 2015. The DNA of a nation. *Nature*, 524(7566):503–505.
- Mell, P., Grance, T. et al., 2011. The NIST definition of cloud computing.
- Meng, X.-Y., Zhang, H.-X., Mezei, M., and Cui, M., 2011. Molecular docking: a powerful approach for structure-based drug discovery. *Current computer-aided drug design*, 7(2):146–157.
- Messerschmitt, D. G., Szyperski, C. et al., 2005. Software ecosystem: understanding an indispensable technology and industry. *MIT Press Books*, 1.
- Miller, S., 2020. COVID researchers can apply for free cloud services. <https://gcn.com/articles/2020/03/24/cloud-vendors-covid-research.aspx>. [Online; accessed 09-September-2020].
- Morris, G. M. and Lim-Wilby, M., 2008. Molecular docking. *Molecular modeling of proteins*, pages 365–382.
- Moses, H., Dorsey, E. R., Matheson, D. H., and Thier, S. O., 2005. Financial anatomy of biomedical research. *Jama*, 294(11):1333–1342.

- Navale, V. and Bourne, P. E., 2018. Cloud computing applications for biomedical science: A perspective. *PLoS computational biology*, 14(6):e1006144.
- Nicola, M., Alsafi, Z., Sohrabi, C., Kerwan, A., Al-Jabir, A., Iosifidis, C., Agha, M., and Agha, R., 2020. The socio-economic implications of the coronavirus pandemic (COVID-19): A review. *International journal of surgery (London, England)*, 78:185.
- Novick, P. A., Ortiz, O. F., Poelman, J., Abdulhay, A. Y., and Pande, V. S., 2013. SWEETLEAD: an in silico database of approved drugs, regulated chemicals, and herbal isolates for computer-aided drug discovery. *PLoS One*, 8(11):e79568.
- Olston, C., Reed, B., Srivastava, U., Kumar, R., and Tomkins, A., 2008. Pig latin: a not-so-foreign language for data processing. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1099–1110. ACM.
- Organization, W. H. et al., 2020. WHO Director-General's opening remarks at the media briefing on COVID-19- 11 March 2020.
- Ossyra, J., Sedova, A., Tharrington, A., Noé, F., Clementi, C., and Smith, J. C., 2019. Porting adaptive ensemble molecular dynamics workflows to the summit supercomputer. In *International Conference on High Performance Computing*, pages 397–417. Springer.
- Pairo-Castineira, E., Clohisey, S., Klaric, L., Bretherick, A., Rawlik, K., Parkinson, N., Pasko, D., Walker, S., Richmond, A., Fourman, M. H. et al., 2020. Genetic mechanisms of critical illness in Covid-19. *medRxiv*.
- Perez, G. I. P. and Abadi, A. T. B., 2020. Ongoing Challenges Faced in the Global Control of COVID-19 Pandemic. *Archives of Medical Research*.
- Pordes, R., Petravick, D., Kramer, B., Olson, D., Livny, M., Roy, A., Avery, P., Blackburn, K., Wenaus, T., Würthwein, F. et al., 2007. The open science grid. In *Journal of Physics: Conference Series*, volume 78, page 012057. IOP Publishing.
- Rawlins, M. D., 2004. Cutting the cost of drug development? *Nature reviews Drug discovery*, 3(4): 360–364.
- Roche, J. A. and Roche, R., 2020. A hypothesized role for dysregulated bradykinin signaling in COVID-19 respiratory complications. *The FASEB Journal*.
- Rowe, A., Kalaitzopoulos, D., Osmond, M., Ghanem, M., and Guo, Y., 2003. The discovery net system for high throughput bioinformatics. *Bioinformatics*, 19(suppl\_1):i225–i231.
- Savin, G., Shabanov, B., Telegin, P., and Baranov, A., 2019. Joint supercomputer center of the Russian Academy of Sciences: Present and future. *Lobachevskii Journal of Mathematics*, 40(11):1853–1862.
- Schatz, M. C., Langmead, B., and Salzberg, S. L., 2010. Cloud computing and the DNA data race. *Nature biotechnology*, 28(7):691.
- Schwede, T., Kopp, J., Guex, N., and Peitsch, M. C., 2003. SWISS-MODEL: an automated protein homology- modeling server. *Nucleic acids research*, 31(13):3381–3385.
- Sciaba, A., Campana, S., Litmaath, M., Donno, F., Moscicki, J., Magini, N., Renshall, H., and Andreeva, J., 2010. Computing at the Petabyte scale with the WLCG. Technical report.
- Shanahan, H. P., Owen, A. M., and Harrison, A. P., 2014. Bioinformatics on the cloud computing platform Azure. *PLoS one*, 9(7):e102642.

- Shanahan, J. G. and Dai, L., 2015. Large scale distributed data science using apache spark. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2323–2324. ACM.
- Shipman, G. M., Woodall, T. S., Graham, R. L., Maccabe, A. B., and Bridges, P. G., 2006. Infiniband scalability in Open MPI. In *Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International*, pages 10–pp. IEEE.
- Siva, N., 2015. UK gears up to decode 100 000 genomes from NHS patients. *The Lancet*, 385(9963):103–104.
- Smith, J. E. and Nair, R., 2005. The architecture of virtual machines. *Computer*, 38(5):32–38.
- Smith, M. and Smith, J. C., 2020. Repurposing therapeutics for COVID-19: supercomputer-based docking to the SARS-CoV-2 viral spike protein and viral spike protein-human ACE2 interface.
- Smith, T., 2020. IA Supercomputer Analyzed Covid-19 — and an Interesting New Theory Has Emerged. <https://elemental.medium.com/a-supercomputer-analyzed-covid-19-and-an-interesting-new-theory-has-emerged-31cb8eba9d63>. [Online; accessed 09-September-2020].
- Sorokina, M. and Steinbeck, C., 2020. COlleCtion of Open NatUral producTs. <http://doi.org/10.5281/zenodo.3778405>. [Online; accessed 11-September-2020].
- Sun, K., Chen, J., and Viboud, C., 2020. Early epidemiological analysis of the coronavirus disease 2019 outbreak based on crowdsourced data: a population-level observational study. *The Lancet Digital Health*.
- Sun, M., Liu, S., Wei, X., Wan, S., Huang, M., Song, T., Lu, Y., Weng, X., Lin, Z., Chen, H. et al., 2021. Aptamer Blocking Strategy Inhibits SARS-CoV-2 Virus Infection. *Angewandte Chemie*, 133(18):10354–10360.
- Taylor, R. C., 2010. An overview of the Hadoop/MapReduce/HBase framework and its current applications in bioinformatics. *BMC bioinformatics*, 11(Suppl 12):S1.
- The Good Hope Net (2020). The Good Hope Net project uses Russian supercomputer to develop treatment against coronavirus infection. <https://thegoodhope.net>. [Online; accessed 05-August-2021].
- Thusoo, A., Sarma, J. S., Jain, N., Shao, Z., Chakka, P., Anthony, S., Liu, H., Wyckoff, P., and Murthy, R., 2009. Hive: a warehousing solution over a map-reduce framework. *Proceedings of the VLDB Endowment*, 2(2):1626–1629.
- Vazhkudai, S. S., De Supinski, B. R., Bland, A. S., Geist, A., Sexton, J., Kahle, J., Zimmer, C. J., Atchley, S., Oral, S., Maxwell, D. E. et al., 2018. The design, deployment, and evaluation of the CORAL pre-exascale systems. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 661–672. IEEE.
- Vilar, S., Cozza, G., and Moro, S., 2008. Medicinal chemistry and the molecular operating environment (MOE): application of QSAR and molecular docking to drug discovery. *Current topics in medicinal chemistry*, 8(18):1555–1572.
- Wassenaar, T. A., Van Dijk, M., Loureiro-Ferreira, N., Van Der Schot, G., De Vries, S. J., Schmitz, C., Van Der Zwan, J., Boelens, R., Giachetti, A., Ferella, L. et al., 2012. WeNMR: structural biology on the grid. *Journal of Grid Computing*, 10(4):743–767.
- WHO, 2020. Statement on the second meeting of the International Health Regulations (2005) Emergency Committee regarding the outbreak of novel coronavirus (2019-nCoV).
- Wong, Z. S., Zhou, J., and Zhang, Q., 2019. Artificial intelligence for infectious disease big data analytics. *Infection, disease & health*, 24(1):44–48.



- World Health Organisation (2020). Off-label use of medicines for COVID-19. <https://www.who.int/publications/i/item/off-label-use-of-medicines-for-covid-19-scientific-brief>. [Online; accessed 11-September-2020].
- Younge, A. J., Henschel, R., Brown, J. T., Von Laszewski, G., Qiu, J., and Fox, G. C., 2011. Analysis of virtualization technologies for high performance computing environments. In *Cloud Computing (CLOUD), 2011 IEEE International Conference on*, pages 9–16. IEEE.
- Zaharia, M., Chowdhury, M., Das, T., Dave, A., Ma, J., McCauley, M., Franklin, M. J., Shenker, S., and Stoica, I., 2012. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 2–2. USENIX Association.
- Zhang, D., Hu, M., and Ji, Q., 2020. Financial markets under the global pandemic of COVID-19. *Finance Research Letters*, page 101528.
- Zhang, X., Wong, S. E., and Lightstone, F. C., 2013. Message passing interface and multithreading hybrid for parallel molecular docking of large databases on petascale high performance computing machines. *Journal of computational chemistry*, 34(11):915–927.
- Zhang, Y., 2020. Initial genome release of novel coronavirus.
- Zimmerman, M. I., Porter, J. R., Ward, M. D., Singh, S., Vithani, N., Meller, A., Mallimadugula, U. L., Kuhn, C. E., Borowsky, J. H., Wiewiora, R. P., Hurley, M. F. D., Harbison, A. M., Fogarty, C. A., Coffland, J. E., Fadda, E., Voelz, V. A., Chodera, J. D., & Bowman, G. R. (2020). SARS-CoV-2 Simulations Go Exascale to Capture Spike Opening and Reveal Cryptic Pockets Across the Proteome. *BioRxiv*, 2020.06.27.175430. <https://doi.org/10.1101/2020.06.27.175430>.

## Author's Biography



Dr. **Jamie J. Alnasir** is a post-doctoral research associate at the SCALE Lab, Department of Computing, Imperial College London and gained his Ph.D. from the University of London. His research interests are distributed computing, high-performance computing, DNA-storage, computational biology, next-generation sequencing, scientific workflows and bioinformatics. He is also a Genomics England Clinical Interpretation Partnership (Gecip) member, investigating viral insertions into human genomes. At the ICR (Institute of Cancer Research) London, he worked at the Scientific Computing department helping researchers leverage HPC, training them in the use of workflow languages and consulting in scientific software engineering.