Pre-Exascale HPC- approaches for Molecular Dynamics simulations. Covid-19 research: a use case

Milosz Wieczor[1&], Vito Gena[1&], Juan Aranda[1&], Rosa M. Badia[2], Josep Lluis Gelpi[2,3], Vytautas Gapsys[4], Bert L. de Groot[4], Erik Lindahl[5,6], Adam Hospital[1*] and Modesto Orozco[1,3*]

Exascale computing has been a dream for ages and is close to become a reality that will impact the way in which molecular simulations are being performed and the quantity and quality of the information derived for them. We review how the biomolecular simulations field is anticipating these new architectures, making emphasis in recent work from groups in the BioExcel Center of Excellence for High Performance Computing. We exemplified the power of these simulation strategies with the work done by the HPC simulation community to fight Covid-19 pandemics.

[1] Institute for Research in Biomedicine (IRB Barcelona). The Barcelona Institute of Science and Technology.
[2] Barcelona Supercomputing Center
[3] Department of Biochemistry and Biomedicine. University of Barcelona. Barcelona, Spain
[4] Max Planck Institute for Multidisciplinary Sciences. Computational biomolecular dynamics group. Goettingen, Germany
[5] Swedish e-Science Research Center. Dept. of Applied Physics. KTH Royal Institute of Technology. Stockholm. Sweden.
[6] Science for Life Laboratory, Dept. of Biochemistry and Biophysics, Stockholm University. Stockholm. Sweden.

[&] Equally contributing authors
* Correspondence to Adam Hospital: adam.hospital@irbbarcelona.org or Modesto Orozco: modesto.orozco@irbbarcelona.org

NEW COMPUTER ARCHITECTURES

Biomolecular simulations have linked to computer power since their origin in the sixties. Thus, access to Golem computer explained (in part) the leadership of Weizmann's institute in the development of the first force-field based methods oriented towards the study of biomacromolecules [1], [2], and a large computer placed at CECAM headquarters allowed the first molecular dynamics (MD) simulation of a protein by McCammon, Gellin and Karplus [3]. Since the seventies, continuous improvements in computers have allowed molecular simulations, and in particular MD to be established in the structural biology and biophysics worlds.

The increase in computer power seemed unstoppable for decades but, unfortunately, current chip technology has reached a plateau and we cannot expect much more powerful processors emerging from current computer technology. Thus, while waiting for quantum computers, hardware developers have tried to maintain the increase in processing power by three different strategies: i) to build specific purpose processors, ii) to take advantages of secondary processors specialized in basic numerical operations (like the graphical processing units; GPUs) and iii) to increase the number of cores in the computers.

MD specific computers such as Anton [4], [5] were a major promise for the MD field when first appeared in 2008, but thirteen years later, their global impact has been small. On the other hand, GPUs have revolutionized the field allowing small groups access to significant compute power. Yet, no single GPU can perform massive MD simulations, thus requiring the combination of a large number of CPUs (and eventually GPUs) in the simulation. Thus, most of massive MD simulations are performed in High-Performance Computing (HPC) environments, especially in gigantic supercomputers created by the aggregation of a massive number of CPUs or CPU/GPUs. Two excellent examples of Tier-0 supercomputers are Japanese Fugaku and American Summit. Rikken supercomputer Fugaku has 7.6 million general-purpose cores A64FX 48C and provides 0.5 ExaFlop of peak power, while American Summit has a hybrid CPU/GPU architecture based on Power9 and NVIDIA Tesla V100 cards, with 2.4 million cores and 0.2 ExaFlop of peak power. Both are amazing pieces of engineering and are approaching to the mythical ExaFlop barrier, but we cannot ignore that Fugaku is only 4 times faster than the top-of-the line computer five years ago (Chinese Sunway) and has an electrical consumption (close to 30 MW excluding cooling) that raises doubts on the sustainability of a much bigger computer [6].

In summary, the biosimulation community has largely benefitted from decades of continuous increase in computer power, but we cannot expect much faster processors in the next years. Most likely, increases in aggregated computer power would appear linked to massively parallel architectures. Our community, and in particular the molecular dynamics (MD) field should redefine its objectives and tools having in mind realistic expectations how new high-performance computers will be.

TAKING ADVANTAGE OF HIGH-PERFORMANCE COMPUTING

The quality of an MD simulation depends mainly on three factors: i) the accuracy of the Hamiltonian representing molecular interactions, ii) the similarity between the simulated system and the physical reality, and iii) the exhaustiveness of the sampling of the conformational space. Most MD simulation engines use the same type of classical Hamiltonian defined in the sixties by Lifson and coworkers

[7], which have been largely refined, reaching good agreement with experimental measurements when sufficient convergence is achieved for some macromolecular systems [8]–[11]. It means that in practice, the quality of MD calculations mainly depends on the ability of the simulated model to represent the reality, and on the completeness of the sampling of the relevant configurations. As some of the biologically relevant systems may contain millions of atoms, a tendency in the MD field in the last decade has been to simulate very big systems [12], [13], which are expected to be more realistic and better fitting with supercomputer architectures. Unfortunately, the larger the system, the longer the simulation generally should be, and the more complex the force evaluation is. In brief, the reliable treatment of very large systems is limited to a few computational groups with access to massive computer resources, while in everyday practice the simulation system sizes range from 10^4 to 10^5 atoms. The challenge is how to efficiently use supercomputers on biological problems involving such "small" systems.

Strategies for parallelization of a single calculation

Molecular dynamics simulation approaches are based on the idea of sampling phase space, and as such the efficiency of the method ultimately comes down to the performance of individual simulations. The total throughput of methods relying e.g. on ensembles is multiplied by the speedup due to parallelization of individual simulations, which in turn is multiplied by speedups from acceleration and algorithmic advances. Several modern MD implementations such as Amber [14], NAMD [15], Desmond [16] and GROMACS [17] have been accelerated either for GPU hardware – using CUDA, OpenCL or SYCL APIs – or single-instruction multiple-data instruction units available on modern CPUs [18]. However, both hardware and software landscapes have become extremely diverse, and while almost all codes can be compiled on arbitrary systems, it cannot be taken for granted that a code is able to make good use of the hardware available on a particular system. For large projects, it has become an important decision to select a suitable combination of HPC hardware along with optimally accelerated simulation software already at the planning stage. GPU accelerators have been a revolution for MD simulations, just as for many other scientific fields, which is effectively due to new algorithms enabling parallelization of the force evaluation in each timestep over ~10,000 functional units in each GPU [19]–[21], although this is opaque to the user.

On the other end of the spectrum, algorithms to explicitly parallelize simulations by communicating between multiple different nodes have evolved to use advanced load balancing and advanced domain decomposition techniques that limit communication to the neighboring nodes [22] and minimize communication e.g. by allowing an interaction between atoms present on two specific nodes to be evaluated on a third ("neutral-territory" methods) [23], [24]. While benchmarks occasionally use gigantic systems to showcase codes, biomolecular applications like Covid-19 are typically focused on achieving sampling e.g. for a protein by first reducing the system size as much as possible and the parallelizing the simulation of this given system over as many nodes as possible – so-called "strong scaling", which is considerably more difficult. The algorithms above have enabled truly impressive strong scaling where simulations have been able to make use of 10,000 cores or more [15], [25], resulting in long time scale simulations of individual SARS-CoV-2 proteins [26], sampling of their binding site dynamics [27], conformational changes [28], modeling of SARS-CoV-2 proteins whose structure is not yet known [29], and multiscale simulations e.g. of the SARS-CoV-2 spike protein [30].

One of the currently most critical bottlenecks is that the particle-mesh Ewald summation [31] algorithms used for long-range electrostatics involves global 3D fast fourier transforms, where the number of communication messages scale as the square of the number of nodes. Fortunately, there

are several promising new algorithms in development based e.g. on multipole expansion methods [32], [33] or multi-level summation [34]. However, since molecular dynamics is inherently an iterative stepwise method, high-end parallelism is subject to hard challenges of communication bandwidth and latency – it will not help to add more CPU resources to a simulation waiting for communication. These problems are even more severe when compounded with accelerators. First, some GPU codes do not even support multi-node parallelism since the random-access nature of indexing atoms to be communicated is not a great fit for accelerator hardware. Second, it introduces another layer of communication complexity and latency when data must first be sent between GPUs in each node and between CPU and GPU before it can be communicated to other nodes. Third, since the force calculation can be an order of magnitude faster on GPUs, the relative impact of network bandwidth and latency will also become tenfold higher, and thus dominate for much smaller node counts compared to CPU-only simulations. As an example, for GROMACS the current effective limit to strong scaling is somewhere around 40 atoms per CPU core, while it is typically not meaningful to go to less than 3000 atoms per GPU [19], and other major codes have qualitatively similar behavior. Notably, this means the relative strong scaling is typically considerably worse for GPU-accelerated simulations, and somewhat paradoxically they might not improve the absolute performance of the simulations in the limit of high-end parallelism, but they make it possible to reach the same performance with order-of-magnitude smaller hardware resources. These freed up resources can then instead be utilized much more efficiently to parallelize the solution to the scientific problem through combination with ensembles of many accelerated and parallel simulations.

Ensemble simulations

In effect, a direct alternative to escape the technical shortcomings of direct parallelism is based on the use of ensembles obtained from independent or loosely coupled simulations. The use of replicas is now well established and considered part of the "good practices" in the field [35], [36], especially when one is concerned with simulations of rare events [37]. As an example, in our 2010 study of spontaneous DNA hairpin folding [38], we were lucky to have found a fast "downhill" folding pathway in our first trajectory (1 µs long). However, to obtain good statistics, we had to collect more than 20 such trajectories, and yet four of them never reached the folded state (after 4 microseconds), highlighting the variability in the outcomes of individual simulations. In fact, thanks to the chaotic nature of MD simulations, even initially correlated replicas rapidly lose memory of their common initial state and can thus be analyzed independently to obtain statistical bounds on a given MD-derived quantity. Even more interestingly, they can be also combined to create meta-trajectories which provide ensembles much richer than expected from the length of individual replicas. In this spirit, it was shown that the multiple replica approach can be used not only to describe equilibrium properties, but also non-equilibrium transitions, even when their characteristic timescales are much larger than those sampled by the individual replicas [39]–[41]. In this spirit, Pande, Noe and others [42], [43] have demonstrated that short replicas can be combined into a Markov State Model (MSM) to trace complex kinetic pathways from the transition probabilities between a discrete number of states in the system, even when no single trajectory samples transitions between all states. In recent years, many generalizations of this scheme have been developed, including history-augmented MSMs that go beyond the simple theory of memoryless Markov processes [44], or the transition-based reweighting analysis method (TRAM) that combines data from multiple ensembles, including explicitly biased ones [45].

Indeed, although an unbiased multiple-replica approach is often the most easily interpretable, the use of biasing schemes usually offers better sampling of the configurational space. This is especially important for the exploration of slow conformational transitions that are unlikely to happen within

currently accessible simulation times. An example of this approach is <u>Hamitonian or Temperature replica exchange</u> [46]–[49] where either the Hamiltonian or the kinetic energy is perturbed across replicas to improve sampling efficiency by pushing the systems out of metastable states with long residence times. An equally popular intermediate approach, called Solute Tempering (REST), combines these two ideas to perturb the effective temperature of a subsystem of interest [50], [51]. In <u>Free Energy Perturbation (FEP) or Thermodynamic Integration (TI)</u>, the Hamiltonian is perturbed in discrete steps (FEP, discrete TI) or smoothly (slow growth TI, non-equilibrium TI) from one state to the other, allowing for the calculation of the associated work – which, depending on the Hamiltonian at hand, can be coupled to a conformational transition or an alchemical change [52]–[55]. In <u>umbrella sampling</u> or <u>metadynamics</u> [56]–[60], additional terms are added to the energy function to favor (in a guided or unguided manner) the exploration of a selected collective variable, often in conjunction with the replica exchange scheme [47], [61]. In the <u>string method</u>, trajectories are pinned to points forming a path in a low-dimensional subspace, and the displacements tangent and perpendicular to the path are used to iteratively refine this path, until a lowest-energy pathway is found between two initially defined states [62]–[64]. Subtler biasing techniques imply the use of information-based biasing as in <u>Maxwell-Demon Dynamics</u>, where after a certain time only trajectories approaching to a given target are extended or cloned [65], [66], or adaptive approaches (very popular in metadynamics or MSMs), which launch new trajectories from regions of the configurational space that have been undersampled in the original ensemble. As a notable example, the Weighted Ensemble (WE) method provides a rigorous scheme for evolving large batches of short trajectories to obtain kinetic and thermodynamic data without the application of external forces [67], [68]. Recently, machine learning approaches hold promise to enhance sampling efficiency by learning a refined biasing scheme from the collected trajectories [69], [70], by re-optimizing the reaction coordinate that is being explored on-the-fly [71], [72], by learning an approximate structural ensemble from which new uncorrelated seeds can be drawn [73], or by simplifying the structural description of the system without compromising on its physical accuracy [74]. Note that while unbiased replica approaches rely on completely independent trajectories, biased multiple replica approaches usually imply a loose coupling between trajectories, as decisions on a given trajectory should be made periodically based on the results obtained in the others.

The multiple replica approach fits perfectly into massively parallel computers, as each of the simulations is independent (or nearly independent) from the others and as the CPU time required to compute the forces is quite independent on the specific coordinates, a good processor load balance is expected. In fact, unbiased multiple replica approaches can be run in distributed platforms, where volunteers provide access to their personal computers to run independent simulations, which are then integrated to obtain the representation of the system. As the first such distributed setup to achieve Exaflop performance in 2020 [75], the <u>Folding@home</u> project [76] is a paradigm for this type of initiatives that has advanced the study of many complex biophysical processes [77]–[81] and provided excellent results in SARS-Cov2 research (see below).

There are several desirable properties that should characterize a well-scaling ensemble approach fitted to the Exascale era:

- Direct parallelism should ideally be replaced with "trivial" parallelism, meaning that ideally ensembles should be generated by a very large number of short, independent trajectory segments.
- The ideal method should be able to keep the number of parallel runs constant to make resource allocation efficient and predictable.

- Explicit communication between runs should be avoided or minimized, and ideally performed asynchronously as not all computing units perform equally even within a single machine.
- Disk space requirements should be minimized, since 1000s of individual runs can easily overwhelm even purpose-built filesystems.
- Efficient GPU implementations should be considered alongside more trivially parallelized CPU ones to ensure compatibility across all types of HPC resources.
- Finally, popular, out-of-the-box parallelization schemes should be used to ensure interoperability across the increasingly diverse computational cluster architectures.

When choosing a method, one should always consider the availability of resources, the existence of "natural" reaction coordinates, the expected complexity of the free energy landscape, the natural timescale of the process in question, and the desired levels of detail and accuracy in the final model.

Non-equilibrium simulations

Biased replica approaches are typically implemented in the context of equilibrium simulations. For example, in metadynamics, umbrella sampling, or FEP/TI, we expect that the trajectory collected at a given point of the transition (geometrical or alchemical) samples well the neighboring points along the reaction coordinate, and the reversible work associated with this microscopic change will not change by extending the trajectory, which often requires long simulations. Similarly, producing a well-converged MSM requires that the transition probabilities between any two states reach local equilibrium, i.e. the trajectory samples an equilibrium situation, which requires long aggregated trajectory times. The need to run a limited number of resource-intensive replicas is undesirable for HPC, as it limits the number of processors that can be allocated to a specific scientific problem. A potential solution arises from the use of non-equilibrium methods, which allow us to obtain transition free energies by combining a very large number of "irreversible works" derived from short non-equilibrium simulations starting from equilibrium start and end states. Non-equilibrium methods have a solid foundation in Jarzynski equation [82] and its later extension, the Crooks fluctuation theorem [83]. They provide access to equilibrium properties, e.g. free energy differences [84], and are efficient in an HPC environment, as they allow for a distributed use of processors working independently. Within the BioExcel consortium [85], De Groot and coworkers have pushed the limits of these non-equilibrium techniques in the context of alchemical perturbation, developing a full methodology that goes from defining an optimum pathway for the transition to collecting converged free energy results for the alchemical transition [86], [87]. This methodology is becoming very successful in exploring relative binding free energies of related drugs or the impact of mutation in protein interactions [84], [88] and has been extensively used in the context of Covid19 research (see below). Note that the technique can also be adapted to biasing schemes related to geometrical rather than Hamiltonian changes, even though initial attempts exposed a number of specific challenges that need to be addressed if a general workflow is to be established [89], [90].

Multi-simulation projects and workflows

As already discussed, a typical research projects cannot be solved by running a single simulation, but require combining many of them. For example, to determine the structural role of sequence variants on a short loop of a protein, at least 20n (with n equal to the length of the loop) independent simulations need to be done. The Folding@home strategy is well suited for this type of problems as we are referring to completely independent and asynchronous simulations. There are, however,

other cases that require combining different types of simulations that must be arranged using flexible workflows. For example, before running a free energy calculation on the impact of sequence variants on the binding of a large series of drugs to a given target, many calculations need to be done: 3D models and topologies of the drugs need to be generated, force-field parameters have to be determined, and equilibrated structures need to be obtained. In parallel, 3D models of the target protein and its mutants have to be made, optimized, thermalized and equilibrated through MD simulations of both apo and holo forms of the protein. Finally, all these processes must be repeated for hundreds of drugs and sequence variants to reach biologically relevant conclusions, and after some analytics on all the results can be needed. In this context, the cost of running ultra-efficient free energy calculations might have no impact in the global project, as the limiting step can be the manual setup of chains of calculations. Thus, in order to be efficient, manual intervention need to be reduced to a minimum by using suitable applications [91]–[93] automating all the preparation steps. Furthermore, the execution of the different programs has to be coordinated by a workflow manager to guarantee no lag time between calculations.

Workflows have been very popular in bioinformatics [94]–[98], but their adoption in the biomolecular simulation fields has been slower [99], [100], mainly due to the powerful computational resources needed and the large amounts of data generated. The recent advances in the informatics technology, including faster network bandwidths, optimized and distributed file systems and pre-exascale HPC supercomputers opened the door to biomolecular simulation execution pipelines. The combination of automated workflows with new optimized codes and HPC supercomputers [17], [101] allows an efficient usage of the available resources and the implementation of large studies through tasks parallelization. We anticipate that workflows will be in the core of biosimulations projects in a near future.

Execution of workflows is typically controlled by Workflow Management Systems (WFMS): the software that regulates the execution of a defined sequence of tasks, arranges them in the most efficient way, manages intermediate data, and monitors the whole execution. The list of available WFMS is huge, and is still increasing every year (see the list of workflow managers compiled and regularly updated by the Common Workflow Language (CWL) team [102] at https://github.com/common-workflow-language/common-workflow-language/wiki/Existing-Workflow-systems). Some of these tools also allow for workflow setup by means of Graphical User Interfaces (GUIs) [103]–[109]. Very popular in the bioinformatics field, these drag-and-drop interfaces (KNIME [103], Pipeline-Pilot [104], Taverna [105], Galaxy [106], Unipro UGENE [107], Kepler [108]) allow for very easy workflow development based on interconnected nodes (building blocks). The chemoinformatics and MD fields have just recently developed GUI-based WFMS, with some available implementations in Galaxy [109], KNIME [110]–[112], Kepler [113] and Pipeline-Pilot [104], the most popular being integrated into commercial software. Although most of these WFMS can be used to launch jobs in HPC clusters through integrated remote execution functionalities, they were not designed to run massive HPC jobs.

For this reason, the biomolecular simulation field should rely on a different type of WFMS adapted to work with HPC infrastructure, eventually reaching the Exascale regime. BioExcel CoE joined together some of the most important HPC WFMS around the world in a dedicated workshop (2018), with the goal of identifying specific software gaps and opportunities for improved workflow practices [114]. A representation of the most popular WFMS in the field were present in the workshop: Parsl [115], AdaptiveMD [116], PyCOMPSs [117], BioSimSpace [118], LongBow [119], FireWorks [120], Swift [121], ExTASY [122], [123], and RADICAL [124], [125]. Although sharing many of the most important features (automation, data control flow, task management, dependency graph), each WFMS has its own design goal, and adoption of one versus another implies a thorough

analysis on the particular research project needs and the WFMS features. Some of the differentiating elements that were identified are the possibility to run adaptive algorithms (ability to change their behavior in runtime), fault tolerance (ability to distinguish between critical and non-critical task failures), flexibility (ability to extend/reduce the number of resources used at runtime), and provenance (ability to store input metadata for reproducibility purposes). Some of these tools have been designed to cover very specific applications, whereas others are able to handle broader, more generic workflows. Finally, the installation and configuration overhead is a critical point that should be carefully considered before choosing a WFMS.

HPC-focused WFMS have demonstrated their power in biomolecular research projects, and Molecular simulation pipelines are now regularly being used to gain insights into the binding free energy and the residence time of ligands in drug development studies [126]. They paved the way for the classical protein-folding problem before AlphaFold2 appeared [127], [128] and have been extensively used in recent studies helping the fight against the SARS-CoV-2 virus [30], [129]–[132] (see below). However, despite their obvious impact, HPC workflows are still not widely adopted by the biosimulation community, mainly due to the difficulty of integrating them with the most common simulation tools. One step towards such integration is the BioExcel Building Blocks (BioBB) library: a collection of portable wrappers on top of common biomolecular simulation tools [133]. The library is designed to: i) increase the interoperability between the tools wrapped; ii) facilitate the implementation of biomolecular simulation workflows; and iii) increase the reusability and reproducibility of the generated workflows. The library is being developed following the FAIR principles for research software development best practices. The result is a collection of building block modules, classified according to the functionalities offered (e.g. MD, analysis, chemistry). The BioBB library is compatible with different WFMS, including both GUIs (Galaxy, KNIME) and HPC-focused (PyCOMPSs). Installation of BioBB workflows controlled by PyCOMPSs WFMS can be easily done using BioConda packages [134].

The community has no doubt that biomolecular simulation workflows will have a clear impact on the future grand challenges in science [135], and that they will be essential tools for the efficient exploitation of the future Exascale supercomputer generation. However, important aspects linked to the pipeline execution still need further research and discussion to reach a community agreement. Data provenance, historical record of the data [136], [137]; workflow metadata, important information from the input, output and intermediate data; reproducibility, possibility to generate the same results using a particular pipeline [138], [139]; reusability, possibility to use a particular pipeline in different resources [102]; and FAIR computational workflows [140] are all crucial points whose treatment is still under debate.

THE PROBLEM OF DATA

As computer power increases, the management of data is becoming the Achilles heel of molecular simulation. As described in a recent review [141], one microsecond of a medium-sized protein system (10^5 atoms including solvent) generates around 5 Pb of data, and even if only 1‰ of coordinates are stored, a single trajectory would require 5 Tb of disk space. The common practice in the field has been to store trajectories locally and delete them after some period, which means that very valuable (and expensive to obtain) information is lost. FAIR (Findable, Accessible, Interoperable and Reusable) requirements for scientific data [142] forces the storage of plain trajectories and associated simulation metadata. This generates a formidable problem that requires at least: i) efficient metadata to specify the purpose and conditions of the simulation, ii) efficient ways to store

the raw data, iii) efficient transfer mechanisms and iv) flexible analysis portals including virtual research environments (VRE) for remote programmatic access to the data [143].

The BioExcel CoE and the ELIXIR European Initiative have defined a prototype of EDAM ontologies for metadata of MD simulations [144]. This should be extended and refined to be adopted by already existing [10], [145]–[151] or new simulation databases. Some recent databases such as BigNASim already implement an ontology metadata [150] which helps to classify trajectories, favoring then meta-analysis and the detection of artefactual simulations. Defining what should be stored from a trajectory is a major decision, which is often guided by traditions maintained since times when the throughput of HPC systems was a fraction of the current one. A significant work has been done by communities such as ABC [147], [152], [153] to define the maximum acceptable frequency for storing data and the amount of solvent that needs to be maintained for reanalysis. In parallel, software developers are creating compression approaches and importing methods used to compress other type of massive data [154]. Again, the community should reach a consensus on the type of compression strategy and the amount of data to be stored, otherwise the FAIR requirement will not be fulfilled.

Despite cleaver approaches to improve transfer efficiency (GridFTP, Aspera and others; see [141]), current network limitations are severely hampering transferring raw trajectories. Streaming methods which would allow basic analysis in remote before downloading the trajectory are going to be useful to avoid unnecessary downloads [141]. However, it seems that the most sensible approach is to reduce the movement of MD files, with big data centers placed physically close to the computer power concentrating the data, mimicking current situation in other fields such as bioinformatics or genomics. The first MD databases were available around 10-15 years ago and aimed to cover all representative structures in the protein data bank: Dynameomics, MoDEL, or Dynasome [10], [145]–[147] are example of these families of databases. More recent efforts have been focused on specific families of proteins of special importance, such as GPCRmd that includes trajectories of G-protein coupled receptors, MemProtMD [152], which contains simulations on 3,500 membrane proteins in realistic membrane environments, or TMB-iBIOMES that includes hundreds of nucleosome trajectories [148]. A few databases are focused on small molecules with potential biomedical or biomolecular impact; an example is Cyclo-lib [149] that is focused on cyclodextrins, or BCE [151], which contains classical trajectories and quantum structural data on small drugs. Finally, BigNASim [150] is, to our knowledge, the only database focused on nucleic acids simulations, and contains hundreds of simulations including the latest ones obtained by the ABC consortium [155]. Recent efforts to create repositories and associated databases developed by the BioExcel CoE to help Covid19 research will be commented below.

Most of these databases are coupled to flexible graphical web interfaces allowing access to key analyses of the trajectories. However, to be really useful, these user interface should be flexible and interactive, something made possible with recent database architectures. For example in the BioExcel-CV19 web server, an associated REST API allows the user to extract slices of the trajectory as well as fragments of the molecule, which are then shown by the NGL visualizer. Using these tools, the user can focus their analysis on configurations fulfilling certain criteria, or perform meta-analyses combining parts of different trajectories. The next step towards gaining interactivity for analysis tools coupled to databases is probably the generalization of virtual research environments (VRE), infrastructures that allow the users to deploy and share their software to perform customized analyses on databases. These infrastructures make it possible to perform specific analyses on the computers supporting the databases, eventually only downloading analysis results instead of the original trajectories.

Entering the Exascale era, we must not forget the Exabyte problem that will be created by the future generation of Exaflop computers. Otherwise, massive amounts of resources will be wasted without providing useful information for the community.

COVID-19: A USE CASE FOR HPC-MD SIMULATIONS

SARS-CoV-2, an enveloped and positive-strand RNA coronavirus, was first identified in the city of Wuhan in December 2019 and rapidly spread worldwide due to a high inter-human transmission rate and a relevant percentage of asymptomatic infections [156]–[158]. As of December 2021, almost 300 million cases and more than 5 million deaths have been registered worldwide, with dramatic implications for the global economy and human well-being. Since its emergence in late 2019, both the SARS-CoV-2 virus and the host response – starting at the main receptor, the angiotensin-converting enzyme 2 (ACE2) – have been extensively studied to understand the disease etiopathogenesis, the structure of the functional viral proteins, and the viral-host interaction pattern to guide a rapid development of both direct-acting antiviral and host-directed therapeutic agents. In this context, the theoreticians' community, assisted by the constantly increasing number of X-ray and high-resolution cryogenic electron microscopy (Cryo-EM) structures available, has promptly reacted and a massive worldwide computational endeavour has been made to unravel the molecular details of the virus' functioning at a multiscale level, from force-field based all-atom simulations to hybrid quantum-mechanics/molecular-mechanics approaches. Here, we will summarize selected studies that have provided key information on the functioning of the virus obtainable uniquely via scientific HPC simulations.

RNA-dependent RNA polymerase

The RNA-dependent RNA polymerase (RdRp) is the core protein of the replication-transcription complex, which is the machinery in charge of the transcription and replication of SARS CoV-2's viral genome and subgenomic mRNAs [159]. The RdRp holoenzyme is composed of four subunits: one core nsp12 and cofactors, two nsp8s and one nsp7. In addition to RdRp, other important proteins are involved in the replication complex: a helicase (nsp13), an exonuclease that proofreads the newly synthetized RNA (nsp14), mRNA capping enzymes (nsp14, 16), or a specific endoribonuclease (nsp15).

Despite the great effort invested in vaccinating the global population, as of October 2021 still more than 50% of people remain without any dose of the vaccine. Moreover, the rapid spread of variants and its possible resistance to current vaccines still make antiviral drugs important means to alleviate the effects of the pandemic and react swiftly to its future outbursts. Thus, due to its central role in the viral life cycle, the RdRp has been an attractive target for antiviral drugs. Among them, Remdesivir has gained much attention currently being the only small-molecule drug approved for the treatment of Covid-19 by FDA. Remdesivir is an adenine nucleotide analog that acts by inhibiting RdRp through a delayed chain termination mechanism [160]. Other inhibition strategies have also been deployed towards protein partners involved in the replication machinery or by hacking the virus' replication through lethal mutagenesis [161]. Although RdRp has been primarily investigated using experimental techniques, theoretical computations have helped understand different aspects of its inner workings. In order to better understand SARS-CoV-2's RdRp, Shaw and co-workers made use of Cryo-EM microscopy, RNA-protein cross-linking and unbiased molecular dynamics simulations to unveil RdRp's backtracking mechanism [162]. In this study, three structures in which the

backtracked RNA was accommodated by the nucleotide-triphosphate (NTP) entry channel were resolved, a mechanism also found in cellular RNA polymerases. Starting from the cryo-EM structure that featured two nsp8 cofactors, a nsp7 cofactor, two nsp13 helicase cofactors, an RNA template and a product RNA strand with one mismatched nucleotide at its 3' end, three independent 5-μs unbiased molecular dynamics simulations were carried out. Although the backtracking process was not captured in the simulations, these trajectories revealed the spontaneous fraying of the 3' terminal mismatched nucleotide into the NTP entry channel. This positioning would disfavour the RNA translocation and thus the elongation process.

Numerous experimental and computational studies have focused on understanding Remdesivir's inhibitory mechanism inside RdRp. However, Remdesivir's mode of action is still debated. A theoretical study investigated the possible mechanism of action at the molecular level with a total of 34 μs of simulation [163]. Absolute free energy calculations were performed to obtain binding free energies of Remdesivir and two other nucleotide analogues inside RdRp, followed by independent short molecular dynamics simulations (100 ns) for each of the systems where Remdesivir had been incorporated in the nascent RNA strand. Their results would support inhibition coupled by a delayed chain termination, or alternatively by destabilizing the protein complex.

Main Protease

Proteolytic cleavage of the SARS-CoV-2's virally expressed polyproteins pp1a and pp1ab into its functional proteins is carried out by two viral proteases: the main protease Mpro (nsp5) that cuts 11 sites, and the papain-like protease PLpro (nsp3) which cuts 3 sites [159]. This a crucial step that enables the formation of the replication machinery in charge of transcription and replication of the viral genome inside the host cell. In addition, viral proteases can act on cellular proteins to help escape the virus from immune response [164]. Thus, many efforts have been directed towards the inhibition of viral proteases as an effective antiviral strategy. Moreover, due to its dissimilarity to human proteases, Mpro and PLpro are promising drug targets. Mpro and PLpro are cysteine proteases composed of three and four domains respectively. While Mpro is active as a dimer, PLpro acts as a monomer, and as in other cysteine proteases, their active sites are formed by a Cys-His catalytic dyad.

Joint crystallographic and molecular dynamics studies have analyzed the effects and mode of binding/action of substrates or promising inhibitors inside the Mpro/PLpro enzymes [165]. The work by Hummer and co-workers studied SARS-Cov-2 PLpro by means of biochemical, structural and functional analysis. Through multi-μs MD simulations they analysed the interaction between PLpro and two different protein substrates. In addition, 1-μs long MD simulations were obtained for both SARS-CoV and SARS-CoV-2 PLpros in complex with GRL-0617, a non-covalent inhibitor of SARS-CoV-PLpro that was shown to also act against SARS-CoV-2-PLpro. Kovalevsky et al. crystalized Mpro at room temperature and performed 1 μs molecular dynamics simulations to analyse the flexibility of some of the protein's regions [166]. Experimental and theoretical approaches were also combined to study Mpro inhibition by myricetin and its derivatives [167]. Two independent ~1.5 μs long Gaussian accelerated-MD simulations explained myricetin's dynamic interactions inside Mpro's active site. In addition to drug-design efforts, quantum mechanics (QM) calculations were performed on a small active-site QM-cluster model, including implicit solvation. Reaction pathways and transition states were characterized for myricetin and two other derivatives, explaining the formation of a covalent bond. The calculated activation free energies were obtained through thermal correction. Alternative approaches combining MD-metatrajectories, docking experiments and machine learning strategies have been used to explore binding of millions to billion of chemical

compounds as binders of the protease. These massive efforts, which fit into the paradigm of distributed computing, allowed the community to develop chemical moieties that could then undergo lead optimization protocols [28], [168]–[170] .

Additionally, fully theoretical studies have helped understand features not accessible to experiments. Atomistic simulations unveiled the binding preferences of the promising inhibitor ebselen to Mpro through 6 µs of cumulative MD calculations [171]. In another computational study, important insights into structure-function relationships of Mpro were obtained through continuous constant pH MD simulations [172] coupled to replica-exchange enhanced sampling. The study revealed the impact of the protonation of His172, which causes a conformational change of Mpro resulting in its deactivation. Based on their results, the authors provide a set of guidelines to design new inhibitors that could bind to Mpro's residues. Tuñon and co-workers made use of extensive QM/MM-MD simulations to characterize the free energy profiles of the reaction mechanism of Mpro with a peptidic substrate [173]. Exploration of the reaction free energy landscape was performed at the DFT(B3LYP)/MM level making use of the string and umbrella sampling methods, with a total of 2.4 ns of cumulative sampling. They propose a Mpro reaction mechanism consisting of three steps with 4 transition states, with deacylation as the rate-limiting step of the overall reaction. The study unveils the most important residues involved in catalysis which can help guide drug-design.

Spike: the crucial protein for infection

One of the most widely investigated SARS-CoV-2 protein is the viral Spike protein, a transmembrane homotrimeric class-I fusion glycoprotein which exists in a metastable prefusion architecture and comprises two functional subunits for binding to the host cell receptor ACE2 (S1 subunit) and for fusion of the viral and host cell membranes (S2 subunit) [174]. The Spike protein of SARS-CoV is cleaved by a host-cell protease, the transmembrane protease/serine subfamily member 2 (TMPRSS2); an alveolar cell serine protease preferentially expressed on epithelial cells of the respiratory tract [175]–[177]. TMPRSS2-mediated cleavage of SARS-CoV Spike protein is propaedeutic to host ACE2 binding, membrane fusion and cell-entry mechanism; a highly-concerted and regulated step-wise mechanism. Particular attention has been paid to the highly antigenic S1-based receptor-binding domains (RBD) that by undergoing substantial hinge-like conformational rearrangements, transiently expose (Up conformation) or hide (closed conformation) the interface area required for ACE2 receptor binding and subsequent shedding of S1, refolding of S2 for membrane fusion and virus internalization [178]–[180]. Many aspects of Spike have been subjected to massive simulations.

Among them, glycosylation of both viral and host receptor has been quite a debated topic among theoretical chemists and biochemists. Indeed, it is known that extracellular portion of ACE2 contains at least seven N-glycosylated sites (i.e., N53, N90, N103, N322, N432, N546, N690) and several O-glycosylation sites (e.g., T730) [181], [182], with the latter shown to be crucial in the process of protein-protein recognition and binding mode [179], [183], [184]. To shed light on the role of glycans in mediating virus internalization and spread, Hummer and coworkers have recently performed extensive multi-replica approach MD simulations of the fully-glycosylated human ACE2 receptor bound to the RBD of SARS-CoV-2 Spike and its unbound state [185]. Their results have shown that glycosylation at position N90 in human ACE2 structurally hampers virus binding. This explains, at the atomistic level, the experimental data on the augmented susceptibility to SARS-CoV-2 infection when N90 glycosylation is lost [185]. On the other hand, N322 glycan aids infectivity by hiding a key cryptic epitope known to be the target of CR3022 neutralizing antibody [186]. In their work, Hummer and collaborators have used a classical multi-replica approach of the 11 different glycosylated sites

of ACE2. For each of the 11 setups, they carried out three independent MD simulation runs [(1 × 1 μs and 2 × 480 ns) × 11] for a total of ~22 μs trajectory for almost 1 million atoms. The so obtained trajectories were analysed to dissect protein-protein interactions starting from the Cryo-EM native contacts, which were identified accordingly to the Best et al. approach [187].

Glycosylation was also intensively analysed by Amaro and co-workers that have built a full-length model of the glycosylated SARS-CoV-2 S protein, both in the open and closed states and have performed multiple microsecond-long, all-atom molecular dynamics simulations. Their efforts have produces key full-atom structural and dynamical insights on the roles of glycans and their implications on the i) protein geometry, ii) the global plasticity of Spike and iii) the shielding role against antibody epitope recognition [188]. In general, the numerous glycans found on the whole Spike protein play diverse roles and contribute differently to the global and local conformational rearrangements in a time-dependent manner. Indeed, they observed that the two N165 and N234 glycosylation sites play a major role in shifting the equilibrium ensemble from closed to open RBD conformation. The simulated models pinpoint N165 and N234 positions as a major modulator of the RBD conformational plasticity with N234 glycan that promotes RBD exposure and concurrent stabilization. To provide further support for their findings, they performed additional sets of extensive MD simulations on the N165A/N234A mutant of the SARS-CoV-2 Spike protein, which corroborated the importance of these glycosylated positions in the infective dynamics of the protein.This result was also confirmed by bilayer interferometry experiments. From the technical point of view, this work required an enormous computational effort as the full-length structures were embedded into an equilibrated all-atom membrane bilayer with a composition mimicking the one where the virus buds. Subsequently, explicit water molecules and ions were added, resulting in two final systems of ~1.7 million atoms each for which multiple replicas were simulated collecting more than 15 μs of the overall sampling time. Massively paralel computers were used for this study.

Spike conformational transitions and the functional characterization of SARS-CoV-2 rising variants have been an additional question the scientific community has targeted. As anticipated before, the three-protomeric Spike protein can adopt different macrostates prevalently determined by the serial combination of each single RBD arrangement. Since RBD exposure (i.e., closed-to-open) is propaedeutic to ACE2 binding and virus internalization, and considering that the vast majority of evolution-selected mutations fall within the RBD domain (contributing to the creation of higher-infective viral strains), dissecting RBD's opening/closing mechanism and energetics at molecular level is of utmost importance. In this regard, S. Gnanakaran and co-workers [189] have characterized the flexibility and conformational transition of the Spike protein in a work where MD simulations using an HPC resource have provided a unique and detailed perspective that other standard biophysical approaches would have not produced. In particular, they have analyzed the effect of D614G mutation on the Spike protein discovering that the mutation heavily affects the interaction pattern of RBD, especially when this adopts the open-state conformation. Briefly, they found three contributions to structural rearrangements associated with the D614G mutation: (i) variation in the inter-protomer contacts, (ii) alteration in the correlations between the single RBDs, and (iii) variation in the specific inter-protomer hydrogen bond pattern. As a consequence, the infection-capable one-up state conformation of RBD is favored when D614G mutation occurs. Based on these results, the authors suggest that these three effects can lead to an increase in infectivity through an increase in the one-up population of the G-form. In this work, mainly unbiased MD simulations were performed, generating 5 replicas for each of the system analyzed in the study, generating overall ~20 μs of trajectories for ~1 million atoms simulated in explicit solvent. In a similar vein, Ray, Le and Andricioaei studied the long-range allosteric communication within the Spike, using an enforced closed-to-open transition to study dynamical correlations between individual residues [190]. By calculating correlation scores, mutual information-based cross-correlation metrics and connectivity

graph networks from a set of over 40 short trajectories, they traced allosteric pathways governing the global motions of the RBD to identify residues that, upon mutation, would most likely disrupt interdomain communications. Their analysis was able to recover the established effect of the D614G substitution on Spike opening, as well as provide a plausible explanation for the biological effects of other observed mutations such as A570D or P681H. However, the authors note that their model does not yet have predictive power, one that would allow to score new mutations by their potential threat to human health.

Nevertheless, the effect of natural mutations of SARS-CoV-2 on Spike structure, conformation, and antigenicity has been the leading-topic recently due to the rapidly growing number of virus variants with increased transmissibility, higher disease severity, resistance to neutralizing antibodies elicited by current vaccines or from previous infection, reduced efficacy of treatments, or failure of diagnostic methods. For this reason, these variants have been named Variants of Concern (VoCs). Gobeil et al. [191] investigated the Spike VoC involved in transmission between minks and humans, as well as the B1.1.7 (alpha variant), B.1.351 (beta variant), and P1 (gamma variant) Spike variants. All variants showed a remarkably increased ACE2 binding affinity accompanied by a higher propensity of RBDs to stably adopt the functionally-active up states. While adaptation to mink resulted in Spike destabilization, the B.1.1.7 (UK variant) Spike balanced stabilizing and destabilizing mutations. A local destabilizing effect of the RBD E484K mutation was implicated in resistance of the B.1.1.28/P.1 (Brazil Variant) and B.1.351 (South Africa Variant) lineages to neutralizing antibodies. Also, authors have revealed allosteric effects of mutations and mechanistic differences that may govern zoonotic jumps or escaping mechanisms from antibody neutralization. Finally, extensive Adaptive Sampling Molecular Dynamic Simulations (ASMD) were performed by De Fabritis and co-workers [192], [193] to explore more under-sampled regions of the conformational space of Spike, hence overcoming energetic barriers separating free energy wells. In ASMD, simulations are launched in sequential batches called epochs utilizing knowledge of the conformational space obtained from all previous epochs. To generate an ensemble of the RBD tip conformations required to start the adaptive sampling routine, authors have performed 100 distinct 50ns-long simulations in the NVT ensemble with randomized initial velocities for each of the WT and mutant systems. Each iteration consisted of 50 to 100 short independent simulations, conformations from which were selected to further sample the most under-sampled conformational space regions. Time independent component analysis (TICA) [194] was performed to identify slowly changing molecular order parameters and subsequently Markov state models were built. A total of 29 adaptive iterations were performed, yielding total simulation times of 274.8 µs and 256.8 µs for the WT and Mut systems.

An out-scaling Folding@home distributed computing project [76] involved 1 million people to create a distributed exascale computer with the aim of simulating ~0.1 seconds of the entire viral proteome. The brilliant idea of Gregory R. Bowman and collaborators [75] has permitted to observe a dramatic opening of the apo Spike complex, far beyond that seen experimentally, explaining and predicting the existence of 'cryptic' epitopes. The project has also highlighted large conformational changes across the proteome, which reveal over 50 'cryptic' pockets that expand targeting options for the design of antivirals, thus paving the way towards dozens of alternative drug discovery projects. Important observations were also made in other aspects of Spike dynamics. The authors, thanks to extensive sampling, successfully captured this rare event for both glycosylated and unglycosylated protein and found that glycosylation only slightly increases the population of the open state, the effect being smaller than that produced by genetic variation in the protein. The authors also found that opening occurs only for a single RBD at a time and that the scale of spike opening is often substantially larger than has been observed in experimental snapshots in the absence of binding partners. An enormous amount of calculating nodes worldwide were created

and authors measured an impressive Folding@home peak performance of 1.01 exaflops. This was achieved at a point when ~280,000 GPUs and 4.8 million CPU cores were performing simulations simultaneously; a performance five-fold greater than the peak performance of the world's fastest traditional supercomputer at that time, Summit [6].

All the simulations were performed with the Gromacs package [17] and the immense computational power of peers distributed around the globe has permitted to investigate intensively and with the richest sampling ever obtained 36 distinct proteins associated to SARS-CoV-2 virus. The endeavors have produced (so far) a total outstanding simulated time-scale of ~115 ms (> 0,1 s).

The combination of BioExcel Building Blocks and the PyCOMPSs workflow manager illustrates how workflows can help reaching the High Throughput (HT) regime in HPC infrastructures, running hundreds of calculations using hundreds of cores, thus exploiting thousands of cores in one single job. The possibility of running these huge executions will help the efficient usage of large supercomputers such as the ones that are to become available in the coming years. Workflows built using these tools have been used in some of the COVID-19 related research presented in this review. One example is a pipeline for the calculation of binding free energy differences upon protein residue mutations (Fig. 1). The relative changes in binding free energies ($\Delta\Delta G$s) are computed using a so-called alchemical fast-growth thermodynamic integration (TI) method that does not require a quasi-equilibrium simulation during the alchemical transition [195]. The workflow starts from two independent equilibrium simulations: WT and mutant. These simulations need to sufficiently sample the end state ensembles, as the free energy accuracy will depend on the sampling convergence. From the generated trajectories snapshots are selected to start fast (picoseconds-long) transitions driving the system in the forward (WT to mutant) and reverse (mutant to WT) directions. The work values required to perform these transitions are collected and the Crooks Fluctuation Theorem [83] is used to calculate the free energy difference between the two states.


Figure 1: Left: Schematic representation of the non-equilibrium alchemical free energy calculations (taken from Aldeghi, 2019). Right: Alchemical Free Energy calculations workflow implemented with BioExcel Building Blocks.

The whole pipeline was implemented using the BioExcel Building Blocks library, wrapping GROMACS and pmx biomolecular tools [86], [87]. pmx is used to generate hybrid structures and topologies for the mutated residues and GROMACS to perform molecular dynamics simulations. The PyCOMPSs workflow manager is used to automatically distribute and parallelize the high number of transitions between the wild type and mutant states of the protein. Many of these integration steps are independent calculations, and in principple can be run at the same time The flexibility added by PyCOMPSs allows a variable number of HPC cores depending on their availability (see use example in Fig. 2). Here, a particular alchemical free energy calculation was performed on a 218,179-atom system, running a total number of short 1,000 TI runs using 32 nodes (1,536 cores) of the MareNostrum supercomputer in a single job. This particular execution took four hours, using 100% of the CPUs available during most of this time. The number of nodes/cores to be used in the pipeline is completely flexible as PyCOMPSs is taking care of distributing the independent simulations across the number of available cores (see examples of pre-exascale workflows built within the BioExcel CoE including the one described in this section in the BioExcel GitHub repository: https://github.com/bioexcel/biobb_hpc_workflows).
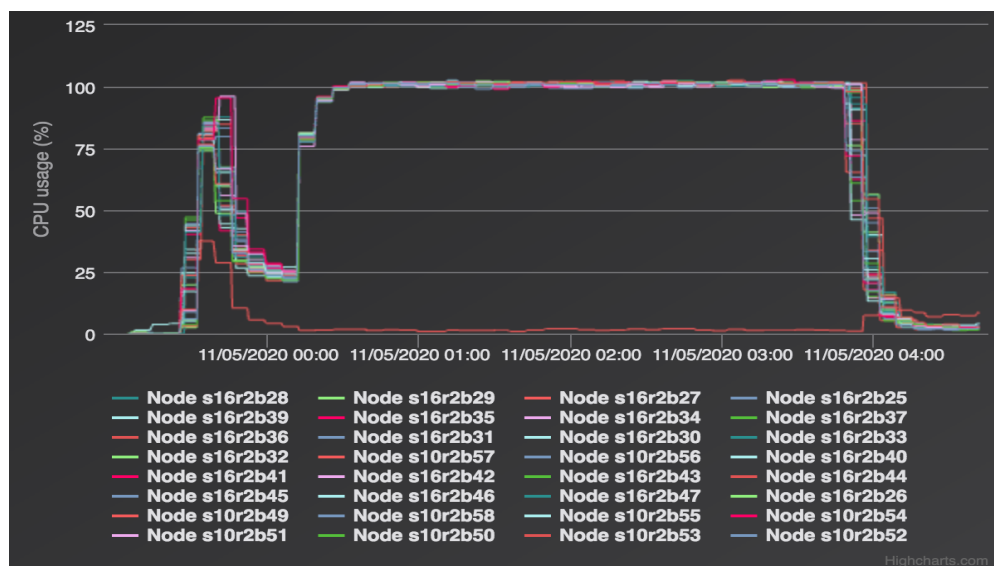

Figure 2: High parallelization reached thanks to the PyCOMPSs workflow manager. Example of a single job using 32 nodes (1,536 cores) of the Marenostrum supercomputer at the BSC. Each line shows the CPU usage in % of a single MareNostrum node.

These workflows were heavily used to understand the effect of mutations in both virus and host on spike recognition by ACE2, finding, for example, that most human ACE2 sequence variants have a negligible effect on the binding affinity of the RBD, but also that the presence of some rare sequence variants would protect a small fraction of the population from virus infection (see Figure 3).

Mutation	Frequency	SARS-CoV-2 RBD hACE2 ΔΔG
Ser19Pro	0.000313	0.86
Ile21Val	0.000011	-0.65
Lys26Arg	0.003883	-1.12
Thr27Ala	0.000011	0.1
Glu35Lys	0.000016	4.25
Phe40Leu	0.000061	1.02
Met82Ile	0.000024	1.44
Glu329Gly	0.000034	0.36
Asp355Asn	0.000012	3.42
Pro389His	0.000038	1.36

Figure 3. Impact of human sequence variants in RBD-hACE2 binding free energy, computed with BioExcel workflows including GROMACS and pmx. Human ACE2 protein mutations with higher frequency in the population are shown. Positive ΔΔG values indicate that a mutation reduces binding affinity. The ΔΔG values are in kcal/mol.

Similarly, Gapsys & de Groot have explored an effect of ACE2 mutations on the ACE2-RBD binding creating a two-level screening strategy (Figure 4). Firstly, the residues at the ACE2-RBD interface were exhaustively scanned with Rosetta Flex ddG [Barlow2018] protocol. This approach provides a good compromise for accuracy and computational efficiency [Aldeghi2018] allowing to predict changes in binding affinity for a large number of mutations. In the second step, a selection of mutations based on the predictions from the previous step was subjected to the pmx/Gromacs based alchemical free energy calculation protocol. While these calculations are computationally demanding, they also offer high prediction accuracy [Gapsys2016]. In addition to estimating changes in the ACE2-RBD binding affinity, relative changes in ACE2 stability upon amino acid mutation were evaluated as well. All in all, this multi-step strategy allows identifying the residue positions in ACE2 and their mutations which would have the largest impact on the binding with the viral RBD. This strategy can be easily adjusted to probe for RBD mutations and their effect on the binding affinity, thus identifying viral variants of higher potency. Furthermore, by replacing ACE2, as the RBD binding partner, to an antibody or peptide, such mutation scan can be used to design novel therapeutics by enhancing their binding affinity to the viral RBD.

Figure 4. A multi-step strategy to predict the effect of ACE2 mutations on the ACE-RBD binding affinity. In the first step, the ACE2 residues at the protein binding interface are scanned by means of the computationally efficient Rosetta Flex ddG protocol. In the second step, a selection of the mutations is probed by the computationally more demanding MD based free energy calculations using pmx and Gromacs software packages. The calculations allow evaluating the effect of mutations on the protein binding affinity, as well as on the stability of the ACE2 protein. This strategy can be further adapted to predict the effects of virus mutations, design high affinity antibodies or peptide therapeutics.

The data

The entire scientific community reacted to the COVID-19 pandemic by redirecting efforts to study the SARS-CoV-2 infection and its mechanisms of action. The biomolecular simulation field contributed since day one with an incredible number of MD simulations, which in turn produced an enormous amount of data distributed around the different groups working on them. BioExcel CoE, in collaboration with the Molecular Sciences Software Institute (MolSSI), developed the COVID-19 Molecular Structure and Therapeutics Hub (https://covid.bioexcel.eu/), a website presented as a community-driven data repository and curation service for molecular structures, models, therapeutics, and simulations related to COVID-19 computational research. The repository was

18

designed from scratch to share data (including MD data) from the scientific community, making them completely open to better tackle the COVID-19 global emergency. The Hub has become a reference repository with huge amounts of useful information gathered in one single portal. Renowned groups in the field (e.g. D.E. Shaw, Riken, Folding@home) have contributed to the repository, summing up to milliseconds of trajectory data. Trajectories stored include different structures involved in the process of virus infection and virus life cycle, such as the ones previously presented in this review (Spike, Protease, RNA Polymerase).

Today the Hub is an essential repository for the field, a central point where to find and download useful data for research. The BioExcel-CV19 database and associated web server (Fig. 5) expand the power of the Hub, including interactive graphical representations of the trajectories and analyses performed on them. The main objective of the BioExcel-CV19 project is to generate a tool for scientists interested in the COVID-19 research to interactively and graphically check key structural and dynamic features stemming from MDs. As these features vary depending on the structure analyzed, specific analyses were performed, uploaded to the database, and represented in the web portal. These analyses and key features were collected by direct interaction with the authors of the simulations. As an example, trajectories corresponding to the viral RBD-hACE2 complex include interface observables (e.g. residue distances, hydrogen bonds), allowing an easy analysis of their behaviour along the simulation (see Figure 5)


Figure 5. Screenshots from the BioExcel-CV19 web server. a) Trajectory representation (*berzosertib* FDA approved drug molecule binding to the ectodomain of human ACE2); b) Analysis of the ligand pocket flexibility; c) Analysis of electrostatic and Van der Waals interactions between the ligand and the ACE2 receptor; d) Analysis of Hydrogen Bonding between the ligand and the ACE2 receptor; e) Electrostatic potential surface in the ligand-binding pocket; f) Insight on the main interaction

between the ligand and the ACE2 receptor involving an Aspartic Acid. URL: https://bioexcel-cv19.bsc.es/#/browse/MCV1900072

All the analyses integrated in the web portal are completely interactive. Whenever possible, a direct link from the analysis to the 3D representation is offered, using NGL viewer tool. The fast extraction of a particular snapshot from the whole trajectory is possible thanks to the NoSQL Mongo database backend powering the server. The whole set of trajectories (atomistic 3D coordinates for every atom and every frame of the simulation) and analyses are stored in this distributed database and efficiently retrieved on the fly from any web portal request. The entire pipeline is automated and new COVID-19 related trajectories are currently being processed. The set of analyses will be continuously extended, according to the suggestions of the authors of MD simulations.

THE ROAD AHEAD

MD and, in general, biomolecular simulation tools are no longer marginal techniques used by a small set of groups to confirm already known facts. Rather, the methods are used by a huge community, helping to understand the mechanisms of life, making predictions and guiding experiments. The Covid-19 pandemic has highlighted the incredible power of MD simulations, either alone or combined with experimental techniques, to reveal atomistic details of biologically relevant mechanisms. These research projects have also shown how large the computational requirements of state-of-the-art MD simulations are. It is not only a question of access to massive computer platforms, but also of strategies for how to use them in an efficient manner. The bio-simulation community will not have the power to tailor new-generation computers, nor to define the nature of the biological problem. However, we have to be the interface between computer science and biology. Exascale supercomputing will be soon a reality, while Exascale distributed platforms are already here. If we learn from the past, computer scientists will soon dream of the YottaFlop computers [196]. The computational science community should escape the sterile discussion between capacity and capability of computers and our community should strive to take the best from whatever new computer technology emerges, and to fully exploit the data derived from simulations.

REFERENCES

[1]   S. Lifson and A. Warshel, "Consistent force field for calculations of conformations, vibrational spectra, and enthalpies of cycloalkane and n-alkane molecules," *J. Chem. Phys.*, vol. 49, no. 11, pp. 5104–5107, Sep. 1968, doi: 10.1063/1.1670007.

[2]   M. Levitt and S. Lifson, "Refinement of protein conformations using a macromolecular energy minimization procedure," *J. Mol. Biol.*, vol. 46, no. 2, pp. 269–279, Dec. 1969, doi: 10.1016/0022-2836(69)90421-5.

[3]   J. A. McCammon, B. R. Gelin, and M. Karplus, "Dynamics of folded proteins," *Nature*, vol. 267, no. 5612, pp. 585–590, 1977, doi: 10.1038/267585a0.

[4]   D. E. Shaw *et al.*, "Anton, a special-purpose machine for molecular dynamics simulation," *Commun. ACM*, vol. 51, no. 7, pp. 91–97, Jul. 2008, doi: 10.1145/1364782.1364802.

[5]   D. E. Shaw *et al.*, "Anton 2: Raising the Bar for Performance and Programmability in a Special-Purpose Molecular Dynamics Supercomputer," in *International Conference for High Performance Computing, Networking, Storage and Analysis, SC*, Jan. 2014, vol. 2015-Janua, no. January, pp. 41–53, doi: 10.1109/SC.2014.9.

[6]   "TOP500: The List." https://www.top500.org/ (accessed Oct. 06, 2021).

[7]   M. Bixon and S. Lifson, "Potential functions and conformations in cycloalkanes," *Tetrahedron*, vol. 23, no. 2, pp. 769–784, Jan. 1967, doi: 10.1016/0040-4020(67)85023-3.

[8]   F. Palazzesi, M. K. Prakash, M. Bonomi, and A. Barducci, "Accuracy of current all-atom force-fields in modeling protein disordered states," *J. Chem. Theory Comput.*, vol. 11, no. 1, pp. 2–7, Jan. 2015, doi: 10.1021/ct500718s.

[9]   P. D. Dans, I. Ivani, A. Hospital, G. Portella, C. González, and M. Orozco, "How accurate are accurate force-fields for B-DNA?," *Nucleic Acids Res.*, vol. 45, no. 7, pp. 4217–4230, Apr. 2017, doi: 10.1093/nar/gkw1355.

[10]   M. Rueda *et al.*, "A consensus view of protein dynamics," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 104, no. 3, pp. 796–801, Jan. 2007, doi: 10.1073/pnas.0605534104.

[11]   A. Pérez, F. Lankas, F. J. Luque, and M. Orozco, "Towards a molecular dynamics consensus view of B-DNA flexibility," *Nucleic Acids Res.*, vol. 36, no. 7, pp. 2379–2394, Apr. 2008, doi: 10.1093/nar/gkn082.

[12]   J. Jung *et al.*, "New parallel computing algorithm of molecular dynamics for extremely huge scale biological systems," *J. Comput. Chem.*, vol. 42, no. 4, pp. 231–241, Feb. 2021, doi: 10.1002/jcc.26450.

[13]   G. Palermo, A. M. J. J. Bonvin, M. Dal Peraro, R. E. Amaro, and V. Tozzini, "Editorial: Multiscale Modeling From Macromolecules to Cell: Opportunities and Challenges of

Biomolecular Simulations," *Frontiers in Molecular Biosciences*, vol. 7. Frontiers, p. 194, Aug. 28, 2020, doi: 10.3389/fmolb.2020.00194.

[14]   R. Salomon-Ferrer, D. A. Case, and R. C. Walker, "An overview of the Amber biomolecular simulation package," *Wiley Interdiscip. Rev. Comput. Mol. Sci.*, vol. 3, no. 2, pp. 198–210, Mar. 2013, doi: 10.1002/wcms.1121.

[15]   J. C. Phillips *et al.*, "Scalable molecular dynamics on CPU and GPU architectures with NAMD," *J. Chem. Phys.*, vol. 153, no. 4, p. 044130, Jul. 2020, doi: 10.1063/5.0014475.

[16]   K. J. Bowers *et al.*, "Scalable Algorithms for Molecular Dynamics Simulations on Commodity Clusters," Feb. 2007, pp. 43–43, doi: 10.1109/sc.2006.54.

[17]   M. J. Abraham *et al.*, "Gromacs: High performance molecular simulations through multi-level parallelism from laptops to supercomputers," *SoftwareX*, vol. 1–2, pp. 19–25, Sep. 2015, doi: 10.1016/j.softx.2015.06.001.

[18]   S. Páll and B. Hess, "A flexible algorithm for calculating pair interactions on SIMD architectures," *Comput. Phys. Commun.*, vol. 184, no. 12, pp. 2641–2650, Dec. 2013, doi: 10.1016/j.cpc.2013.06.003.

[19]   S. Páll *et al.*, "Heterogeneous parallelization and acceleration of molecular dynamics simulations in GROMACS," *J. Chem. Phys.*, vol. 153, no. 13, p. 134110, Oct. 2020, doi: 10.1063/5.0018516.

[20]   A. W. Götz, M. J. Williamson, D. Xu, D. Poole, S. Le Grand, and R. C. Walker, "Routine microsecond molecular dynamics simulations with AMBER on GPUs. 1. generalized born," *J. Chem. Theory Comput.*, vol. 8, no. 5, pp. 1542–1555, May 2012, doi: 10.1021/ct200909j.

[21]   J. E. Stone, D. J. Hardy, I. S. Ufimtsev, and K. Schulten, "GPU-accelerated molecular modeling coming of age," *J. Mol. Graph. Model.*, vol. 29, no. 2, pp. 116–125, Sep. 2010, doi: 10.1016/j.jmgm.2010.06.010.

[22]   S. Plimpton, "Fast parallel algorithms for short-range molecular dynamics," *J. Comput. Phys.*, vol. 117, no. 1, pp. 1–19, Mar. 1995, doi: 10.1006/jcph.1995.1039.

[23]   K. J. Bowers, R. O. Dror, and D. E. Shaw, "Overview of neutral territory methods for the parallel evaluation of pairwise particle interactions," in *Journal of Physics: Conference Series*, Jan. 2005, vol. 16, no. 1, pp. 300–304, doi: 10.1088/1742-6596/16/1/041.

[24]   K. J. Bowers, R. O. Dror, and D. E. Shaw, "Zonal methods for the parallel execution of range-limited N-body simulations," *J. Comput. Phys.*, vol. 221, no. 1, pp. 303–329, Jan. 2007, doi: 10.1016/j.jcp.2006.06.014.

[25]   C. Kutzner, R. Apostolov, B. Hess, and H. Grubmüller, "Scaling of the GROMACS 4.6 molecular dynamics code on SuperMUC," in *Advances in Parallel Computing*, 2014, vol. 25, pp. 722–727, doi: 10.3233/978-1-61499-381-0-722.

[26]   T. Jaffrelot Inizan *et al.*, "High-resolution mining of the SARS-CoV-2 main protease conformational space: supercomputer-driven unsupervised adaptive sampling," *Chem. Sci.*, vol. 12, no. 13, pp. 4889–4907, Jan. 2021, doi: 10.1039/d1sc00145k.

[27]   J. Gossen *et al.*, "A Blueprint for High Affinity SARS-CoV-2 Mpro Inhibitors from Activity-Based Compound Library Screening Guided by Analysis of Protein Dynamics," *ACS Pharmacol. Transl. Sci.*, vol. 4, no. 3, pp. 1079–1095, Jun. 2021, doi: 10.1021/acsptsci.0c00215.

[28]   A. Grottesi *et al.*, "Computational Studies of SARS-CoV-2 3CLpro: Insights from MD Simulations," *Int. J. Mol. Sci. 2020, Vol. 21, Page 5346*, vol. 21, no. 15, p. 5346, Jul. 2020, doi: 10.3390/IJMS21155346.

[29]   A. Mehregan *et al.*, "Probing effects of the SARS-CoV-2 E protein on membrane curvature and intracellular calcium," *bioRxiv*, p. 2021.05.28.446179, Oct. 2021, doi: 10.1101/2021.05.28.446179.

[30]   L. Casalino *et al.*, "AI-driven multiscale simulations illuminate mechanisms of SARS-CoV-2 spike dynamics," *Int. J. High Perform. Comput. Appl.*, vol. 35, no. 5, pp. 432–451, Apr. 2021, doi: 10.1177/10943420211006452.

[31]   U. Essmann, L. Perera, M. L. Berkowitz, T. Darden, H. Lee, and L. G. Pedersen, "A smooth particle mesh Ewald method," *J. Chem. Phys.*, vol. 103, no. 19, pp. 8577–8593, Aug. 1995, doi: 10.1063/1.470117.

[32]   D. S. Shamshirgar, R. Yokota, A. K. Tornberg, and B. Hess, "Regularizing the fast multipole method for use in molecular simulation," *J. Chem. Phys.*, vol. 151, no. 23, p. 234113, Dec. 2019, doi: 10.1063/1.5122859.

[33]   B. Kohnke *et al.*, "Gromex: A scalable and versatile fast multipole method for biomolecular simulation," in *Lecture Notes in Computational Science and Engineering*, vol. 136, Springer, Cham, 2020, pp. 517–543.

[34]   D. J. Hardy, Z. Wu, J. C. Phillips, J. E. Stone, R. D. Skeel, and K. Schulten, "Multilevel summation method for electrostatic force evaluation," *J. Chem. Theory Comput.*, vol. 11, no. 2, pp. 766–779, Feb. 2015, doi: 10.1021/ct5009075.

[35]   B. Knapp, L. Ospina, and C. M. Deane, "Avoiding False Positive Conclusions in Molecular Simulation: The Importance of Replicas," *J. Chem. Theory Comput.*, vol. 14, no. 12, pp. 6127–6138, Dec. 2018, doi: 10.1021/acs.jctc.8b00391.

[36]   V. Gapsys and B. L. de Groot, "On the importance of statistics in molecular simulations for thermodynamics, kinetics and simulation box size," *Elife*, vol. 9, pp. 1–21, Aug. 2020, doi: 10.7554/ELIFE.57589.

[37]   W. Jiang *et al.*, "Generalized scalable multiple copy algorithms for molecular dynamics simulations in NAMD," *Comput. Phys. Commun.*, vol. 185, no. 3, pp. 908–916, Mar. 2014, doi: 10.1016/j.cpc.2013.12.014.

[38]   G. Portella and M. Orozco, "Multiple routes to characterize the folding of a small DNA hairpin," *Angew. Chemie - Int. Ed.*, vol. 49, no. 42, pp. 7673–7676, Oct. 2010, doi: 10.1002/anie.201003816.

[39]   J. D. Chodera, W. C. Swope, J. W. Pitera, and K. A. Dill, "Long-Time Protein Folding Dynamics from Short-Time Molecular Dynamics Simulations," *http://dx.doi.org/10.1137/06065146X*, vol. 5, no. 4, pp. 1214–1226, Dec. 2006, doi: 10.1137/06065146X.

[40]   V. S. Pande, K. Beauchamp, and G. R. Bowman, "Everything you wanted to know about Markov State Models but were afraid to ask," *Methods*, vol. 52, no. 1. pp. 99–105, 2010, doi: 10.1016/j.ymeth.2010.06.002.

[41]   A. Sirur, D. De Sancho, and R. B. Best, "Markov state models of protein misfolding," *J. Chem. Phys.*, vol. 144, no. 7, p. 75101, Feb. 2016, doi: 10.1063/1.4941579.

[42]   G. R. Bowman, V. S. Pande, and F. Noé, "An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation," *Springer*, vol. 797, p. 148, 2014, Accessed: Aug. 29, 2021. [Online]. Available: https://books.google.com/books/about/An_Introduction_to_Markov_State_Models_a.html ?id=ggnGBAAAQBAJ.

[43]   F. Nüske, H. Wu, J. H. Prinz, C. Wehmeyer, C. Clementi, and F. Noé, "Markov state models from short non-equilibrium simulations - Analysis and correction of estimation bias," *J. Chem. Phys.*, vol. 146, no. 9, p. 094104, Mar. 2017, doi: 10.1063/1.4976518.

[44]   E. Suárez, J. L. Adelman, and D. M. Zuckerman, "Accurate Estimation of Protein Folding and Unfolding Times: Beyond Markov State Models," *J. Chem. Theory Comput.*, vol. 12, no. 8, pp. 3473–3481, Aug. 2016, doi: 10.1021/acs.jctc.6b00339.

[45]   H. Wua, F. Paul, C. Wehmeyer, and F. Noé, "Multiensemble Markov models of molecular thermodynamics and kinetics," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 113, no. 23, pp. E3221–E3230, Jun. 2016, doi: 10.1073/pnas.1525092113.

[46]   G. Bussi, "Hamiltonian replica exchange in GROMACS: A flexible implementation," in *Molecular Physics*, Feb. 2014, vol. 112, no. 3–4, pp. 379–384, doi: 10.1080/00268976.2013.824126.

[47]   Y. Sugita, A. Kitao, and Y. Okamoto, "Multidimensional replica-exchange method for free-energy calculations," *J. Chem. Phys.*, vol. 113, no. 15, pp. 6042–6051, Oct. 2000, doi: 10.1063/1.1308516.

[48]   K. Hukushima and K. Nemoto, "Exchange Monte Carlo Method and Application to Spin Glass Simulations," *J. Phys. Soc. Japan*, vol. 65, no. 6, pp. 1604–1608, Nov. 1996, doi: 10.1143/JPSJ.65.1604.

[49]   Y. Sugita and Y. Okamoto, "Replica-exchange molecular dynamics method for protein folding," *Chem. Phys. Lett.*, vol. 314, no. 1–2, pp. 141–151, Nov. 1999, doi: 10.1016/S0009-2614(99)01123-9.

[50] P. Liu, B. Kim, R. A. Friesner, and B. J. Berne, "Replica exchange with solute tempering: A method for sampling biological systems in explicit water," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 102, no. 39, pp. 13749–13754, Sep. 2005, doi: 10.1073/pnas.0506346102.

[51] L. Wang, R. A. Friesner, and B. J. Berne, "Replica exchange with solute scaling: A more efficient version of replica exchange with solute tempering (REST2)," *J. Phys. Chem. B*, vol. 115, no. 30, pp. 9431–9438, Aug. 2011, doi: 10.1021/jp204407d.

[52] R. W. Zwanzig, "High-temperature equation of state by a perturbation method. I. Nonpolar gases," *J. Chem. Phys.*, vol. 22, no. 8, pp. 1420–1426, 1954, doi: 10.1063/1.1740409.

[53] W. L. Jorgensen and L. L. Thomas, "Perspective on free-energy perturbation calculations for chemical equilibria," *J. Chem. Theory Comput.*, vol. 4, no. 6, pp. 869–876, 2008, doi: 10.1021/ct800011m.

[54] P. A. Bash, U. C. Singh, R. Langridge, and P. A. Kollman, "Free energy calculations by computer simulation," *Science (80-.).*, vol. 236, no. 4801, pp. 564–568, May 1987, doi: 10.1126/science.3576184.

[55] T. P. Straatsma and J. A. McCammon, "Multiconfiguration thermodynamic integration," *J. Chem. Phys.*, vol. 95, no. 2, pp. 1175–1188, Aug. 1991, doi: 10.1063/1.461148.

[56] G. M. Torrie and J. P. Valleau, "Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling," *J. Comput. Phys.*, vol. 23, no. 2, pp. 187–199, 1977, doi: 10.1016/0021-9991(77)90121-8.

[57] A. Laio and M. Parrinello, "Escaping free-energy minima," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 99, no. 20, pp. 12562–12566, Oct. 2002, doi: 10.1073/pnas.202427399.

[58] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, "THE weighted histogram analysis method for free-energy calculations on biomolecules. I. The method," *J. Comput. Chem.*, vol. 13, no. 8, pp. 1011–1021, 1992, doi: 10.1002/jcc.540130812.

[59] H. Grubmüller, "Predicting slow structural transitions in macromolecular systems: Conformational flooding," *Phys. Rev. E*, vol. 52, no. 3, pp. 2893–2906, Sep. 1995, doi: 10.1103/PhysRevE.52.2893.

[60] A. Barducci, G. Bussi, and M. Parrinello, "Well-tempered metadynamics: A smoothly converging and tunable free-energy method," *Phys. Rev. Lett.*, vol. 100, no. 2, p. 020603, Jan. 2008, doi: 10.1103/PhysRevLett.100.020603.

[61] S. Piana and A. Laio, "A bias-exchange approach to protein folding," *J. Phys. Chem. B*, vol. 111, no. 17, pp. 4553–4559, May 2007, doi: 10.1021/jp0678731.

[62] L. Maragliano, A. Fischer, E. Vanden-Eijnden, and G. Ciccotti, "String method in collective variables: Minimum free energy paths and isocommittor surfaces," *J. Chem. Phys.*, vol. 125, no. 2, p. 024106, Jul. 2006, doi: 10.1063/1.2212942.

[63] A. C. Pan, D. Sezer, and B. Roux, "Finding transition pathways using the string method with swarms of trajectories," *J. Phys. Chem. B*, vol. 112, no. 11, pp. 3432–3440, Mar. 2008, doi: 10.1021/jp0777059.

[64] L. Maragliano, B. Roux, and E. Vanden-Eijnden, "Comparison between mean forces and swarms-of-trajectories string methods," *J. Chem. Theory Comput.*, vol. 10, no. 2, pp. 524–533, Feb. 2014, doi: 10.1021/ct400606c.

[65] M. Rueda, E. Cubero, C. A. Laughton, and M. Orozco, "Exploring the counterion atmosphere around DNA: What can be learned from molecular dynamics simulations?," *Biophys. J.*, vol. 87, no. 2, pp. 800–811, Aug. 2004, doi: 10.1529/biophysj.104.040451.

[66] P. Sfriso, A. Emperador, L. Orellana, A. Hospital, J. L. Gelpí, and M. Orozco, "Finding conformational transition pathways from discrete molecular dynamics simulations," *J. Chem. Theory Comput.*, vol. 8, no. 11, pp. 4707–4718, Nov. 2012, doi: 10.1021/ct300494q.

[67] D. M. Zuckerman and L. T. Chong, "Weighted Ensemble Simulation: Review of Methodology, Applications, and Software," *Annual Review of Biophysics*, vol. 46, no. 1. Annual Reviews Inc., pp. 43–57, May 22, 2017, doi: 10.1146/annurev-biophys-070816-033834.

[68] M. C. Zwier *et al.*, "WESTPA: An interoperable, highly scalable software package for weighted ensemble simulation and analysis," *J. Chem. Theory Comput.*, vol. 11, no. 2, pp. 800–809, Feb. 2015, doi: 10.1021/ct5010615.

[69] M. I. Zimmerman and G. R. Bowman, "FAST Conformational Searches by Balancing Exploration/Exploitation Trade-Offs," *J. Chem. Theory Comput.*, vol. 11, no. 12, pp. 5747–5757, Nov. 2015, doi: 10.1021/acs.jctc.5b00737.

[70] M. I. Zimmerman, J. R. Porter, X. Sun, R. R. Silva, and G. R. Bowman, "Choice of Adaptive Sampling Strategy Impacts State Discovery, Transition Probabilities, and the Apparent Mechanism of Conformational Changes," *J. Chem. Theory Comput.*, vol. 14, no. 11, pp. 5459–5475, Nov. 2018, doi: 10.1021/acs.jctc.8b00500.

[71] S. Brandt, F. Sittel, M. Ernst, and G. Stock, "Machine Learning of Biomolecular Reaction Coordinates," *J. Phys. Chem. Lett.*, vol. 9, no. 9, pp. 2144–2150, May 2018, doi: 10.1021/acs.jpclett.8b00759.

[72] H. Jung, R. Covino, and G. Hummer, "Artificial Intelligence Assists Discovery of Reaction Coordinates and Mechanisms from Molecular Dynamics Simulations," Jan. 2019, Accessed: Aug. 30, 2021. [Online]. Available: https://arxiv.org/abs/1901.04595v1.

[73] F. Noé, S. Olsson, J. Köhler, and H. Wu, "Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning," *Science (80-.)*, vol. 365, no. 6457, 2019, doi: 10.1126/science.aaw1147.

[74] J. Wang *et al.*, "Machine Learning of Coarse-Grained Molecular Dynamics Force Fields," *ACS Cent. Sci.*, vol. 5, no. 5, pp. 755–767, May 2019, doi: 10.1021/acscentsci.8b00913.

[75] M. I. Zimmerman *et al.*, "SARS-CoV-2 simulations go exascale to predict dramatic spike opening and cryptic pockets across the proteome," *Nat. Chem. 2021 137*, vol. 13, no. 7, pp. 651–659, May 2021, doi: 10.1038/s41557-021-00707-0.

[76] A. L. Beberg, D. L. Ensign, G. Jayachandran, S. Khaliq, and V. S. Pande, "Folding@home: Lessons from eight years of volunteer distributed computing," 2009, doi: 10.1109/IPDPS.2009.5160922.

[77] M. M. Sultan, R. A. Denny, R. Unwalla, F. Lovering, and V. S. Pande, "Millisecond dynamics of BTK reveal kinome-wide conformational plasticity within the apo kinase domain," *Sci. Rep.*, vol. 7, no. 1, Dec. 2017, doi: 10.1038/s41598-017-10697-0.

[78] X. Sun, S. Singh, K. J. Blumer, and G. R. Bowman, "Simulation of spontaneous G protein activation reveals a new intermediate driving GDP unbinding," *Elife*, vol. 7, Oct. 2018, doi: 10.7554/eLife.38465.

[79] Y. Meng, D. Shukla, V. S. Pande, and B. Roux, "Transition path theory analysis of c-Src kinase activation," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 113, no. 33, pp. 9193–9198, Aug. 2016, doi: 10.1073/pnas.1602790113.

[80] D. Shukla, A. Peck, and V. S. Pande, "Conformational heterogeneity of the calmodulin binding interface," *Nat. Commun.*, vol. 7, Apr. 2016, doi: 10.1038/ncomms10910.

[81] S. M. Ryckbosch, P. A. Wender, and V. S. Pande, "Molecular dynamics simulations reveal ligand-controlled positioning of a peripheral protein complex in membranes," *Nat. Commun.*, vol. 8, no. 1, Dec. 2017, doi: 10.1038/s41467-016-0015-8.

[82] C. Jarzynski, "Nonequilibrium equality for free energy differences," *Phys. Rev. Lett.*, vol. 78, no. 14, pp. 2690–2693, Apr. 1997, doi: 10.1103/PhysRevLett.78.2690.

[83] G. E. Crooks, "Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences," *Phys. Rev. E - Stat. Physics, Plasmas, Fluids, Relat. Interdiscip. Top.*, vol. 60, no. 3, pp. 2721–2726, Sep. 1999, doi: 10.1103/PhysRevE.60.2721.

[84] H. M. Baumann, V. Gapsys, B. L. De Groot, and D. L. Mobley, "Challenges Encountered Applying Equilibrium and Nonequilibrium Binding Free Energy Calculations," *J. Phys. Chem. B*, vol. 125, no. 17, pp. 4241–4261, May 2021, doi: 10.1021/acs.jpcb.0c10263.

[85] "BioExcel – Centre of Excellence for Computation Biomolecular Research." https://bioexcel.eu/ (accessed Aug. 30, 2021).

[86] V. Gapsys, S. Michielssens, D. Seeliger, and B. L. De Groot, "pmx: Automated protein structure and topology generation for alchemical perturbations," *J. Comput. Chem.*, vol. 36, no. 5, pp. 348–354, Feb. 2015, doi: 10.1002/jcc.23804.

[87] V. Gapsys and B. L. De Groot, "pmx Webserver: A User Friendly Interface for Alchemistry," *J. Chem. Inf. Model.*, vol. 57, no. 2, pp. 109–114, Feb. 2017, doi: 10.1021/acs.jcim.6b00498.

[88]   V. Gapsys *et al.*, "Large scale relative protein ligand binding affinities using non-equilibrium alchemy," *Chem. Sci.*, vol. 11, no. 4, pp. 1140–1152, Jan. 2020, doi: 10.1039/c9sc03754c.

[89]   S. Wolf and G. Stock, "Targeted Molecular Dynamics Calculations of Free Energy Profiles Using a Nonequilibrium Friction Correction," *J. Chem. Theory Comput.*, vol. 14, no. 12, pp. 6175–6182, Dec. 2018, doi: 10.1021/acs.jctc.8b00835.

[90]   E. Pohjolainen, S. Malola, G. Groenhof, and H. Häkkinen, "Exploring Strategies for Labeling Viruses with Gold Nanoclusters through Non-equilibrium Molecular Dynamics Simulations," *Bioconjug. Chem.*, vol. 28, no. 9, pp. 2327–2339, Sep. 2017, doi: 10.1021/acs.bioconjchem.7b00367.

[91]   S. Jo, T. Kim, V. G. Iyer, and W. Im, "CHARMM-GUI: A web-based graphical user interface for CHARMM," *J. Comput. Chem.*, vol. 29, no. 11, pp. 1859–1865, Aug. 2008, doi: 10.1002/JCC.20945.

[92]   A. Hospital, P. Andrio, C. Fenollosa, D. Cicin-Sain, M. Orozco, and J. L. Gelpí, "MDWeb and MDMoby: an integrated web-based platform for molecular dynamics simulations," *Bioinformatics*, vol. 28, no. 9, pp. 1278–1279, May 2012, doi: 10.1093/BIOINFORMATICS/BTS139.

[93]   A. Hospital, I. Faustino, R. Collepardo-Guevara, C. González, J. L. Gelpí, and M. Orozco, "NAFlex: a web server for the study of nucleic acid flexibility," *Nucleic Acids Res.*, vol. 41, no. W1, pp. W47–W55, Jul. 2013, doi: 10.1093/NAR/GKT378.

[94]   O. Spjuth *et al.*, "Experiences with workflows for automating data-intensive bioinformatics," *Biology Direct*, vol. 10, no. 1. BioMed Central, pp. 1–12, Aug. 19, 2015, doi: 10.1186/s13062-015-0071-8.

[95]   B. Fjukstad and L. A. Bongo, "A Review of Scalable Bioinformatics Pipelines," *Data Sci. Eng.*, vol. 2, no. 3, pp. 245–251, Oct. 2017, doi: 10.1007/s41019-017-0047-z.

[96]   J. Leipzig, "A review of bioinformatic pipeline frameworks," *Brief. Bioinform.*, vol. 18, no. 3, pp. 530–536, May 2017, doi: 10.1093/bib/bbw020.

[97]   M. Atkinson, S. Gesing, J. Montagnat, and I. Taylor, "Scientific workflows: Past, present and future," *Future Generation Computer Systems*, vol. 75. North-Holland, pp. 216–227, Oct. 01, 2017, doi: 10.1016/j.future.2017.05.041.

[98]   T. Reiter *et al.*, "Streamlining data-intensive biology with workflow systems," *GigaScience*, vol. 10, no. 1. Oxford Academic, pp. 1–19, Jan. 29, 2021, doi: 10.1093/gigascience/giaa140.

[99]   T. E. Cheatham and D. R. Roe, "The impact of heterogeneous computing on workflows for biomolecular simulation and analysis," *Computing in Science and Engineering*, vol. 17, no. 2. IEEE Computer Society, pp. 30–39, Mar. 01, 2015, doi: 10.1109/MCSE.2015.7.

[100]   A. Shade and T. K. Teal, "Computing Workflows for Biologists: A Roadmap," *PLoS Biol.*, vol. 13, no. 11, p. e1002303, Nov. 2015, doi: 10.1371/journal.pbio.1002303.

[101] C. Kutzner, S. Páll, M. Fechner, A. Esztermann, B. L. de Groot, and H. Grubmüller, "More bang for your buck: Improved use of GPU nodes for GROMACS 2018," *J. Comput. Chem.*, vol. 40, no. 27, pp. 2418–2431, Oct. 2019, doi: 10.1002/jcc.26011.

[102] M. R. Crusoe *et al.*, "Methods Included: Standardizing Computational Reuse and Portability with the Common Workflow Language," May 2021, doi: 10.1145/3486897.

[103] M. R. Berthold *et al.*, "KNIME: The Konstanz information miner," in *Studies in Classification, Data Analysis, and Knowledge Organization*, 2008, pp. 319–326, doi: 10.1007/978-3-540-78246-9_38.

[104] W. A. Warr, "Scientific workflow systems: Pipeline Pilot and KNIME," *J. Comput. Aided. Mol. Des.*, vol. 26, no. 7, pp. 801–804, May 2012, doi: 10.1007/s10822-012-9577-7.

[105] T. Oinn *et al.*, "Taverna: A tool for the composition and enactment of bioinformatics workflows," *Bioinformatics*, vol. 20, no. 17, pp. 3045–3054, Nov. 2004, doi: 10.1093/bioinformatics/bth361.

[106] E. Afgan *et al.*, "The Galaxy platform for accessible, reproducible and collaborative biomedical analyses: 2016 update," *Nucleic Acids Res.*, vol. 44, no. W1, pp. W3–W10, Jul. 2016, doi: 10.1093/nar/gkw343.

[107] K. Okonechnikov *et al.*, "Unipro UGENE: A unified bioinformatics toolkit," *Bioinformatics*, vol. 28, no. 8. Oxford Academic, pp. 1166–1167, Apr. 15, 2012, doi: 10.1093/bioinformatics/bts091.

[108] B. Ludäscher *et al.*, "Scientific workflow management and the Kepler system: Research Articles," *Concurr. Comput. Pract. Exp.*, vol. 18, no. 10, pp. 1039–1065, 2006, doi: 10.1002/cpe.v18:10.

[109] S. A. Bray, X. Lucas, A. Kumar, and B. A. Grüning, "The ChemicalToolbox: Reproducible, user-friendly cheminformatics analysis on the Galaxy platform," *J. Cheminform.*, vol. 12, no. 1, pp. 1–7, Jun. 2020, doi: 10.1186/s13321-020-00442-7.

[110] S. Beisken, T. Meinl, B. Wiswedel, L. F. de Figueiredo, M. Berthold, and C. Steinbeck, "KNIME-CDK: Workflow-driven cheminformatics," *BMC Bioinformatics*, vol. 14, no. 1, pp. 1–4, Aug. 2013, doi: 10.1186/1471-2105-14-257.

[111] V. A. Eyrich, "The Schrödinger KNIME Extensions," Accessed: Sep. 29, 2021. [Online]. Available: https://www.schrodinger.com/products/knime-extensions.

[112] J. M. Gally, S. Bourg, Q. T. Do, S. Aci-Sèche, and P. Bonnet, "VSPrep: A General KNIME Workflow for the Preparation of Molecules for Virtual Screening," *Mol. Inform.*, vol. 36, no. 10, p. 1700023, Oct. 2017, doi: 10.1002/minf.201700023.

[113] S. Purawat *et al.*, "A Kepler Workflow Tool for Reproducible AMBER GPU Molecular Dynamics," *Biophys. J.*, vol. 112, no. 12, pp. 2469–2474, Jun. 2017, doi: 10.1016/j.bpj.2017.04.055.

[114] L. Naden, S. Ellis, and S. Jha, "MolSSI and BioExcel Workflow Workshop 2018 Report," May 2019, Accessed: Sep. 29, 2021. [Online]. Available: https://arxiv.org/abs/1905.11863v1.

[115] Y. Babuji *et al.*, "Parsl: Pervasive parallel programming in Python," in *HPDC 2019-Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing*, 2019, pp. 25–36, doi: 10.1145/3307681.3325400.

[116] J. Ossyra, A. Sedova, A. Tharrington, F. Noé, C. Clementi, and J. C. Smith, "Porting Adaptive Ensemble Molecular Dynamics Workflows to the Summit Supercomputer," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Jun. 2019, vol. 11887 LNCS, pp. 397–417, doi: 10.1007/978-3-030-34356-9_30.

[117] E. Tejedor *et al.*, "PyCOMPSs: Parallel computational workflows in Python," *Int. J. High Perform. Comput. Appl.*, vol. 31, no. 1, pp. 66–82, Aug. 2017, doi: 10.1177/1094342015594678.

[118] L. Hedges *et al.*, "BioSimSpace: An interoperable Python framework for biomolecular simulation," *J. Open Source Softw.*, vol. 4, no. 43, p. 1831, Nov. 2019, doi: 10.21105/joss.01831.

[119] J. Gebbie-Rayet, G. Shannon, H. H. Loeffler, and C. A. Laughton, "Longbow: A Lightweight Remote Job Submission Tool," *J. Open Res. Softw.*, vol. 4, no. 1, p. 1, Jan. 2016, doi: 10.5334/jors.95.

[120] A. Jain *et al.*, "FireWorks: a dynamic workflow system designed for high-throughput applications," *Concurr. Comput. Pract. Exp.*, vol. 27, no. 17, pp. 5037–5059, Dec. 2015, doi: 10.1002/CPE.3505.

[121] Y. Zhao *et al.*, "A Service Framework for Scientific Workflow Management in the Cloud," *IEEE Trans. Serv. Comput.*, vol. 8, no. 6, pp. 930–944, Nov. 2015, doi: 10.1109/TSC.2014.2341235.

[122] V. Balasubramanian *et al.*, "ExTASY: Scalable and flexible coupling of MD simulations and advanced sampling techniques," in *Proceedings of the 2016 IEEE 12th International Conference on e-Science, e-Science 2016*, Mar. 2017, pp. 361–370, doi: 10.1109/eScience.2016.7870921.

[123] E. Hruska, V. Balasubramanian, H. Lee, S. Jha, and C. Clementi, "Extensible and Scalable Adaptive Sampling on Supercomputers," *J. Chem. Theory Comput.*, vol. 16, no. 12, pp. 7915–7925, Dec. 2020, doi: 10.1021/acs.jctc.0c00991.

[124] M. Turilli, V. Balasubramanian, A. Merzky, I. Paraskevakos, and S. Jha, "Middleware Building Blocks for Workflow Systems," *Comput. Sci. Eng.*, vol. 21, no. 4, pp. 62–75, Jul. 2019, doi: 10.1109/MCSE.2019.2920048.

[125] A. Merzky, M. Turilli, M. Maldonado, M. Santcroos, and S. Jha, "Using pilot systems to execute many task workloads on supercomputers," in *Lecture Notes in Computer Science*

(including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), May 2019, vol. 11332 LNCS, pp. 61–82, doi: 10.1007/978-3-030-10632-4_4.

[126] J. Dakka *et al.*, "High-throughput binding affinity calculations at extreme scales," *BMC Bioinformatics*, vol. 19, no. 18, pp. 33–45, Dec. 2018, doi: 10.1186/s12859-018-2506-6.

[127] H. Lee, M. Turilli, S. Jha, D. Bhowmik, H. Ma, and A. Ramanathan, "DeepDriveMD: Deep-learning driven adaptive molecular simulations for protein folding," in *Proceedings of DLS 2019: Deep Learning on Supercomputers - Held in conjunction with SC 2019: The International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov. 2019, pp. 12–19, doi: 10.1109/DLS49591.2019.00007.

[128] H. Ma *et al.*, "Deep generative model driven protein folding simulations," in *Advances in Parallel Computing*, 2020, vol. 36, pp. 45–55, doi: 10.3233/APC200023.

[129] J. V. Vermaas *et al.*, "Supercomputing pipelines search for therapeutics against covid-19," *Comput. Sci. Eng.*, vol. 23, no. 1, pp. 7–16, Jan. 2021, doi: 10.1109/MCSE.2020.3036540.

[130] A. P. Bhati *et al.*, "Pandemic Drugs at Pandemic Speed: Accelerating COVID-19 Drug Discovery with Hybrid Machine Learning- and Physics-based Simulations on High Performance Computers," vol. 20, pp. 1–19, Oct. 2021, Accessed: Sep. 30, 2021. [Online]. Available: https://arxiv.org/abs/2010.06574v1.

[131] S. Wan, A. Bhati, A. Wade, D. Alfè, and P. Coveney, "Thermodynamic and structural insights into the repurposing of drugs that bind to SARS-CoV-2 main protease," pp. 1–10, Aug. 2021, doi: 10.33774/CHEMRXIV-2021-03NRL-V2.

[132] A. Al Saadi *et al.*, "IMPECCABLE: Integrated Modeling PipelinE for COVID Cure by Assessing Better LEads," vol. 20, Oct. 2020, Accessed: Sep. 30, 2021. [Online]. Available: https://arxiv.org/abs/2010.06574v1.

[133] P. Andrio *et al.*, "BioExcel Building Blocks, a software library for interoperable biomolecular simulation workflows," *Sci. Data*, vol. 6, no. 1, pp. 1–8, Sep. 2019, doi: 10.1038/s41597-019-0177-4.

[134] R. Dale *et al.*, "Bioconda: Sustainable and comprehensive software distribution for the life sciences," *Nat. Methods*, vol. 15, no. 7, pp. 475–476, Jul. 2018, doi: 10.1038/s41592-018-0046-7.

[135] A. Krylov *et al.*, "Perspective: Computational chemistry software and its advancement as illustrated through three grand challenge cases for molecular science," *J. Chem. Phys.*, vol. 149, no. 18, p. 180901, Nov. 2018, doi: 10.1063/1.5052551.

[136] R. Ferdous, B. Roy, C. K. Roy, and K. A. Schneider, "Workflow Provenance for Big Data: From Modelling to Reporting," Springer, Cham, 2020, pp. 185–200.

[137] F. Z. Khan, S. Soiland-Reyes, R. O. Sinnott, A. Lonie, C. Goble, and M. R. Crusoe, "Sharing interoperable workflow provenance: A review of best practices and their practical

application in CWLProv," *Gigascience*, vol. 8, no. 11, pp. 1–27, Nov. 2019, doi: 10.1093/gigascience/giz095.

[138] H. Meng and D. Thain, "Facilitating the Reproducibility of Scientific Workflows with Execution Environment Specifications," in *Procedia Computer Science*, Jan. 2017, vol. 108, pp. 705–714, doi: 10.1016/j.procs.2017.05.116.

[139] S. Cohen-Boulakia *et al.*, "Scientific workflows for computational reproducibility in the life sciences: Status, challenges and opportunities," *Futur. Gener. Comput. Syst.*, vol. 75, pp. 284–298, Oct. 2017, doi: 10.1016/j.future.2017.01.012.

[140] C. Goble *et al.*, "FAIR Computational Workflows," *Data Intell.*, vol. 2, no. 1–2, pp. 108–121, Jan. 2020, doi: 10.1162/dint_a_00033.

[141] A. Hospital, F. Battistini, R. Soliva, J. L. Gelpí, and M. Orozco, "Surviving the deluge of biosimulation data," *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 10, no. 3. John Wiley & Sons, Ltd, p. e1449, May 01, 2020, doi: 10.1002/wcms.1449.

[142] S. Stall *et al.*, "Make scientific data FAIR," *Nat. 2021 5707759*, vol. 570, no. 7759, pp. 27–29, Jun. 2019, doi: 10.1038/d41586-019-01720-7.

[143] A. Elofsson, B. Hess, E. Lindahl, A. Onufriev, D. van der Spoel, and A. Wallqvist, "Ten simple rules on how to create open access and reproducible molecular simulations of biological systems," *PLoS Computational Biology*, vol. 15, no. 1. Public Library of Science, p. e1006649, Jan. 17, 2019, doi: 10.1371/journal.pcbi.1006649.

[144] J. Ison *et al.*, "EDAM: An ontology of bioinformatics operations, types of data and identifiers, topics and formats," *Bioinformatics*, vol. 29, no. 10, pp. 1325–1332, May 2013, doi: 10.1093/bioinformatics/btt113.

[145] T. Meyer *et al.*, "MoDEL (Molecular Dynamics Extended Library): A Database of Atomistic Molecular Dynamics Trajectories," *Structure*, vol. 18, no. 11, pp. 1399–1409, Nov. 2010, doi: 10.1016/j.str.2010.07.013.

[146] M. W. van der Kamp *et al.*, "Dynameomics: A Comprehensive Database of Protein Dynamics," *Structure*, vol. 18, no. 4, pp. 423–435, Mar. 2010, doi: 10.1016/j.str.2010.01.012.

[147] U. Hensen, T. Meyer, J. Haas, R. Rex, G. Vriend, and H. Grubmüller, "Exploring protein dynamics space: The dynasome as the missing link between protein structure and function," *PLoS One*, vol. 7, no. 5, p. e33931, May 2012, doi: 10.1371/journal.pone.0033931.

[148] R. Sun, Z. Li, and T. C. Bishop, "TMB-iBIOMES: An iBIOMES-Lite database of Nucleosome Trajectories and meta-analysis," *ChemRxiv*, Mar. 2019, doi: 10.26434/chemrxiv.7793939.

[149] E. Mixcoha, R. Rosende, R. Garcia-Fandino, and Á. Piñeiro, "Cyclo-lib: A database of computational molecular dynamics simulations of cyclodextrins," *Bioinformatics*, vol. 32, no. 21, pp. 3371–3373, Nov. 2016, doi: 10.1093/bioinformatics/btw289.

[150] A. Hospital *et al.*, "BIGNASim: A NoSQL database structure and analysis portal for nucleic acids simulation data," *Nucleic Acids Res.*, vol. 44, no. D1, pp. D272–D278, Jan. 2016, doi: 10.1093/nar/gkv1301.

[151] S. Zivanovic *et al.*, "Bioactive Conformational Ensemble Server and Database. A Public Framework to Speed up in Silico Drug Discovery," *J. Chem. Theory Comput.*, vol. 16, no. 10, pp. 6586–6597, Oct. 2020, doi: 10.1021/acs.jctc.0c00305.

[152] T. D. Newport, M. S. P. Sansom, and P. J. Stansfeld, "The MemProtMD database: a resource for membrane-embedded protein structures and their lipid interactions," *Nucleic Acids Res.*, vol. 47, no. D1, pp. D390–D397, Jan. 2019, doi: 10.1093/NAR/GKY1047.

[153] M. Pasi *et al.*, "µABC: A systematic microsecond molecular dynamics study of tetranucleotide sequence effects in B-DNA," *Nucleic Acids Res.*, vol. 42, no. 19, pp. 12272–12283, Oct. 2014, doi: 10.1093/nar/gku855.

[154] M. Lundborg, R. Apostolov, D. Spångberg, A. Gärdenäs, D. Van Der Spoel, and E. Lindahl, "An efficient and extensible format, library, and API for binary trajectory data from molecular simulations," *J. Comput. Chem.*, vol. 35, no. 3, pp. 260–269, Jan. 2014, doi: 10.1002/JCC.23495.

[155] P. D. Dans *et al.*, "The static and dynamic structural heterogeneities of B-DNA: extending Calladine-Dickerson rules," *Nucleic Acids Res.*, vol. 47, no. 21, pp. 11090–11102, Dec. 2019, doi: 10.1093/nar/gkz905.

[156] S. Tabata *et al.*, "Clinical characteristics of COVID-19 in 104 people with SARS-CoV-2 infection on the Diamond Princess cruise ship: a retrospective analysis," *Lancet Infect. Dis.*, vol. 20, no. 9, pp. 1043–1050, Sep. 2020, doi: 10.1016/S1473-3099(20)30482-5.

[157] P. Zhou *et al.*, "A pneumonia outbreak associated with a new coronavirus of probable bat origin," *Nat. 2020 5797798*, vol. 579, no. 7798, pp. 270–273, Feb. 2020, doi: 10.1038/s41586-020-2012-7.

[158] P. Zhou *et al.*, "Addendum: A pneumonia outbreak associated with a new coronavirus of probable bat origin," *Nat. 2020 5887836*, vol. 588, no. 7836, pp. E6–E6, Nov. 2020, doi: 10.1038/s41586-020-2951-z.

[159] P. V'kovski, A. Kratzel, S. Steiner, H. Stalder, and V. Thiel, "Coronavirus biology and replication: implications for SARS-CoV-2," *Nature Reviews Microbiology*, vol. 19, no. 3. Nature Publishing Group, pp. 155–170, Oct. 28, 2021, doi: 10.1038/s41579-020-00468-6.

[160] C. J. Gordon *et al.*, "Remdesivir is a direct-acting antiviral that inhibits RNA-dependent RNA polymerase from severe acute respiratory syndrome coronavirus 2 with high potency," *Journal of Biological Chemistry*, vol. 295, no. 20. J Biol Chem, pp. 6785–6797, May 15, 2020, doi: 10.1074/jbc.RA120.013679.

[161] C. J. Gordon, E. P. Tchesnokov, R. F. Schinazi, and M. Götte, "Molnupiravir promotes SARS-CoV-2 mutagenesis via the RNA template," *Journal of Biological Chemistry*, vol. 297, no. 1. Elsevier, Jul. 01, 2021, doi: 10.1016/j.jbc.2021.100770.

[162] B. Malone *et al.*, "Structural basis for backtracking by the SARS-CoV-2 replication-transcription complex," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 118, no. 19, May 2021, doi: 10.1073/pnas.2102516118.

[163] F. Byléhn, C. A. Menéndez, G. R. Perez-Lemus, W. Alvarado, and J. J. De Pablo, "Modeling the Binding Mechanism of Remdesivir, Favilavir, and Ribavirin to SARS-CoV-2 RNA-Dependent RNA Polymerase," *ACS Cent. Sci.*, vol. 7, no. 1, pp. 164–174, Jan. 2021, doi: 10.1021/acscentsci.0c01242.

[164] D. Shin *et al.*, "Papain-like protease regulates SARS-CoV-2 viral spread and innate immunity," *Nature*, vol. 587, no. 7835, pp. 657–662, Jul. 2020, doi: 10.1038/s41586-020-2601-5.

[165] M. D. Sacco *et al.*, "Structure and inhibition of the SARS-CoV-2 main protease reveal strategy for developing dual inhibitors against Mpro and cathepsin L," *Sci. Adv.*, vol. 6, no. 50, Dec. 2020, doi: 10.1126/sciadv.abe0751.

[166] D. W. Kneller *et al.*, "Structural plasticity of SARS-CoV-2 3CL Mpro active site cavity revealed by room temperature X-ray crystallography," *Nat. Commun.*, vol. 11, no. 1, pp. 1–6, Jun. 2020, doi: 10.1038/s41467-020-16954-7.

[167] H. Su *et al.*, "Identification of pyrogallol as a warhead in design of covalent inhibitors for the SARS-CoV-2 3CL protease," *Nat. Commun.*, vol. 12, no. 1, pp. 1–12, Jun. 2021, doi: 10.1038/s41467-021-23751-3.

[168] Y. L. Weng, S. R. Naik, N. Dingelstad, M. R. Lugo, S. Kalyaanamoorthy, and A. Ganesan, "Molecular dynamics and in silico mutagenesis on the reversible inhibitor-bound SARS-CoV-2 main protease complexes reveal the role of lateral pocket in enhancing the ligand affinity," *Sci. Reports 2021 111*, vol. 11, no. 1, pp. 1–22, Apr. 2021, doi: 10.1038/s41598-021-86471-0.

[169] A. T. Ton, F. Gentile, M. Hsing, F. Ban, and A. Cherkasov, "Rapid Identification of Potential Inhibitors of SARS-CoV-2 Main Protease by Deep Docking of 1.3 Billion Compounds," *Mol. Inform.*, vol. 39, no. 8, p. 2000028, Aug. 2020, doi: 10.1002/MINF.202000028.

[170] N. Mekni, C. Coronnello, T. Langer, M. De Rosa, and U. Perricone, "Support Vector Machine as a Supervised Learning for the Prioritization of Novel Potential SARS-CoV-2 Main Protease Inhibitors," *Int. J. Mol. Sci. 2021, Vol. 22, Page 7714*, vol. 22, no. 14, p. 7714, Jul. 2021, doi: 10.3390/IJMS22147714.

[171] C. A. Menéndez, F. Byléhn, G. R. Perez-Lemus, W. Alvarado, and J. J. de Pablo, "Molecular characterization of ebselen binding activity to SARS-CoV-2 main protease," *Sci. Adv.*, vol. 6, no. 37, pp. 345–356, Sep. 2020, doi: 10.1126/sciadv.abd0345.

[172] N. Verma, J. A. Henderson, and J. Shen, "Proton-Coupled Conformational Activation of SARS Coronavirus Main Proteases and Opportunity for Designing Small-Molecule Broad-Spectrum

Targeted Covalent Inhibitors," *J. Am. Chem. Soc.*, vol. 142, no. 52, pp. 21883–21890, Dec. 2020, doi: 10.1021/JACS.0C10770.

[173] C. A. Ramos-Guzmán, J. J. Ruiz-Pernía, and I. Tuñón, "Unraveling the SARS-CoV-2 Main Protease Mechanism Using Multiscale Methods," *ACS Catal.*, vol. 10, no. 21, pp. 12544–12554, Nov. 2020, doi: 10.1021/acscatal.0c03420.

[174] S. M. C. Gobeil *et al.*, "Effect of natural mutations of SARS-CoV-2 on spike structure, conformation, and antigenicity," *Science (80-.).*, vol. 373, no. 6555, Aug. 2021, doi: 10.1126/SCIENCE.ABI6226.

[175] S. Matsuyama, N. Nagata, K. Shirato, M. Kawase, M. Takeda, and F. Taguchi, "Efficient Activation of the Severe Acute Respiratory Syndrome Coronavirus Spike Protein by the Transmembrane Protease TMPRSS2," *J. Virol.*, vol. 84, no. 24, pp. 12658–12664, Dec. 2010, doi: 10.1128/jvi.01542-10.

[176] I. Glowacka *et al.*, "Evidence that TMPRSS2 Activates the Severe Acute Respiratory Syndrome Coronavirus Spike Protein for Membrane Fusion and Reduces Viral Control by the Humoral Immune Response," *J. Virol.*, vol. 85, no. 9, pp. 4122–4134, May 2011, doi: 10.1128/jvi.02232-10.

[177] A. Sternberg and C. Naujokat, "Structural features of coronavirus SARS-CoV-2 spike protein: Targets for vaccination," *Life Sciences*, vol. 257. Elsevier, p. 118056, Sep. 15, 2020, doi: 10.1016/j.lfs.2020.118056.

[178] S. S. Hwang *et al.*, "MRNA destabilization by BTG1 and BTG2 maintains T cell quiescence," *Science (80-.).*, vol. 367, no. 6483, pp. 1255–1260, Mar. 2020, doi: 10.1126/science.abb2507.

[179] A. C. Walls, Y. J. Park, M. A. Tortorici, A. Wall, A. T. McGuire, and D. Veesler, "Structure, Function, and Antigenicity of the SARS-CoV-2 Spike Glycoprotein," *Cell*, vol. 181, no. 2, pp. 281-292.e6, Apr. 2020, doi: 10.1016/j.cell.2020.02.058.

[180] L. Duan, Q. Zheng, H. Zhang, Y. Niu, Y. Lou, and H. Wang, "The SARS-CoV-2 Spike Glycoprotein Biosynthesis, Structure, Function, and Antigenicity: Implications for the Design of Spike-Based Vaccine Immunogens," *Frontiers in Immunology*, vol. 11. Front Immunol, Oct. 07, 2020, doi: 10.3389/fimmu.2020.576622.

[181] A. Shajahan, S. Archer-Hartmann, N. T. Supekar, A. S. Gleinich, C. Heiss, and P. Azadi, "Comprehensive characterization of N- and O- glycosylation of SARS-CoV-2 human receptor angiotensin converting enzyme 2," *Glycobiology*, vol. 31, no. 4, pp. 410–424, May 2021, doi: 10.1093/GLYCOB/CWAA101.

[182] R. Yan, Y. Zhang, Y. Li, L. Xia, Y. Guo, and Q. Zhou, "Structural basis for the recognition of SARS-CoV-2 by full-length human ACE2," *Science (80-.).*, vol. 367, no. 6485, pp. 1444–1448, Mar. 2020, doi: 10.1126/science.abb2762.

[183] J. Lan *et al.*, "Structure of the SARS-CoV-2 spike receptor-binding domain bound to the ACE2 receptor," *Nature*, vol. 581, no. 7807, pp. 215–220, Mar. 2020, doi: 10.1038/s41586-020-2180-5.

[184] E. P. Barros *et al.*, "The flexibility of ACE2 in the context of SARS-CoV-2 infection," *Biophys. J.*, vol. 120, no. 6, pp. 1072–1084, Mar. 2021, doi: 10.1016/J.BPJ.2020.10.036.

[185] A. R. Mehdipour and G. Hummer, "Dual nature of human ACE2 glycosylation in binding to SARS-CoV-2 spike," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 118, no. 19, May 2021, doi: 10.1073/pnas.2100425118.

[186] M. Yuan *et al.*, "A highly conserved cryptic epitope in the receptor binding domains of SARS-CoV-2 and SARS-CoV," *Science (80-.).*, vol. 368, no. 6491, pp. 630–633, May 2020, doi: 10.1126/science.abb7269.

[187] R. B. Best, G. Hummer, and W. A. Eaton, "Native contacts determine protein folding mechanisms in atomistic simulations," *Proc. Natl. Acad. Sci.*, vol. 110, no. 44, pp. 17874–17879, Oct. 2013, doi: 10.1073/PNAS.1311599110.

[188] L. Casalino *et al.*, "Beyond Shielding: The Roles of Glycans in the SARS-CoV-2 Spike Protein," *ACS Cent. Sci.*, vol. 6, no. 10, pp. 1722–1734, Oct. 2020, doi: 10.1021/ACSCENTSCI.0C01056.

[189] R. A. Mansbach, S. Chakraborty, K. Nguyen, D. C. Montefiori, B. Korber, and S. Gnanakaran, "The SARS-CoV-2 Spike variant D614G favors an open conformational state," *Sci. Adv.*, vol. 7, no. 16, pp. 3671–3687, Apr. 2021, doi: 10.1126/sciadv.abf3671.

[190] D. Ray, L. Le, and I. Andricioaei, "Distant residues modulate conformational opening in SARS-CoV-2 spike protein," *Proc. Natl. Acad. Sci.*, vol. 118, no. 43, p. e2100943118, Oct. 2021, doi: 10.1073/pnas.2100943118.

[191] S. M.-C. Gobeil *et al.*, "D614G Mutation Alters SARS-CoV-2 Spike Conformation and Enhances Protease Cleavage at the S1/S2 Junction," *Cell Rep.*, vol. 34, no. 2, p. 108630, Dec. 2021, doi: 10.1016/j.celrep.2020.108630.

[192] S. Doerr and G. De Fabritiis, "On-the-fly learning and sampling of ligand binding by high-throughput molecular simulations," *J. Chem. Theory Comput.*, vol. 10, no. 5, pp. 2064–2069, May 2014, doi: 10.1021/ct400919u.

[193] S. Doerr, M. J. Harvey, F. Noé, and G. De Fabritiis, "HTMD: High-Throughput Molecular Dynamics for Molecular Discovery," *J. Chem. Theory Comput.*, vol. 12, no. 4, pp. 1845–1852, Apr. 2016, doi: 10.1021/acs.jctc.6b00049.

[194] G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, "Identification of slow molecular order parameters for Markov model construction," *J. Chem. Phys.*, vol. 139, no. 1, p. 015102, Jul. 2013, doi: 10.1063/1.4811489.

[195] M. Aldeghi, B. L. de Groot, and V. Gapsys, "Accurate Calculation of Free Energy Changes upon Amino Acid Mutation," in *Methods in Molecular Biology*, vol. 1851, Methods Mol Biol, 2019, pp. 19–47.

[196] R. Goñi and M. Orozco, "The Yottaflop Frontier of Atomistic Molecular Dynamics Simulations," in *Theoretical and Quantum Chemistry at the Dawn of the 21st Century*, no. 2018, Apple Academic Press, 2018, pp. 597–616.