

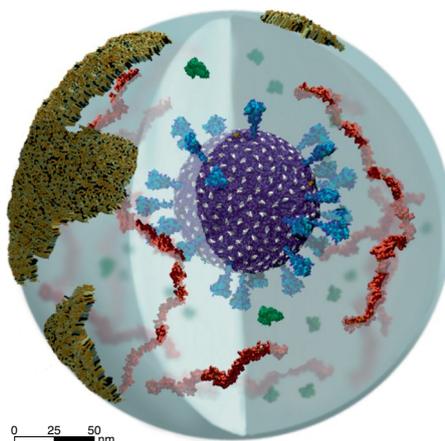
Fighting COVID-19 with HPC

A special prize for outstanding research in high-performance computing (HPC) applied to COVID-19-related challenges sheds light on the computational science community's efforts to fight the ongoing global crisis.

Established in 1987, the ACM Gordon Bell Prize is a prestigious award given each year to recognize outstanding achievement in high-performance computing, rewarding the most innovative research that has applied HPC technology to various applications in science, engineering, and large-scale data analytics. Looking back on the history of the prize gives insights into how scientific computing capabilities have changed and substantially improved over time, as well as into the broad range of applications that can successfully make use of these powerful capabilities, including fluid dynamics simulations, molecular dynamics (MD) simulations, climate analytics, and classical simulations of quantum circuits, to name a few examples. It is certainly remarkable to see how the community comes up with creative solutions that allow for the more effective use of powerful computing facilities to run calculations that would otherwise be hard or even impossible to perform using the more modest resources of a typical desktop computer or workstation.

In response to the ongoing pandemic, a [separate prize](#) was awarded for the first time last year, and awarded again this year, for outstanding research that uses HPC in innovative ways to deepen our understanding of the nature, spread and/or treatment of COVID-19. This prize highlights the efforts of the computational science community, with the help of HPC, in addressing one of the most difficult crises of our recent years. This year's finalists for the special prize, presented at the [SC21 conference](#), showcased the breadth of COVID-19-related challenges being targeted by the community, as well as the range of computational solutions being explored by researchers, which we would like to highlight here.

Some of the finalists focused on SARS-CoV-2 antiviral drug design. Determining potential drug candidates can be very costly: given the vast size of the chemical space, searching for inhibitors that more strongly bind to the target (for instance, SARS-CoV-2 M^{pro} and PL^{pro} proteases) entails running an exhaustive and expensive search. Jens Glaser and colleagues from Oak Ridge National Laboratory (ORNL) used a natural language processing approach in order to accelerate the screening process of potential



One-billion-atom multiscale simulation of the aerosolized version of the SARS-CoV-2 Delta virion. Figure used with permission from Abigail Dommer, Amaro Lab, UCSD.

drug candidates¹. The team generated an unprecedented dataset of approximately 9.6 billion molecules using the SMILES (Simplified Molecular Input Line Entry System) text representation, and pre-trained a large deep learning language model on this input dataset: the model learned a representation for chemical structure in a completely unsupervised manner. This pre-training stage is computationally expensive, and the researchers used the ORNL Summit supercomputer — which is currently the fastest supercomputer in the United States, and the second fastest in the world — to accomplish this task. Then, a smaller dataset of known binding affinities between molecules (potential inhibitors) and targets was used to fine-tune the model for binding affinity prediction: the pre-trained model can be used for candidate generation, and the fine-tuned one can be used for choosing the candidates with greater binding affinity. Both models can run on computers with modest resources, thus making the drug screening stage more broadly accessible to the research community.

The ORNL team was not the only one that focused on drug design: Hai-Bin Luo and colleagues also targeted the large-scale screening process, but by using statistical mechanics-based methods instead of language models². More specifically,

this team focused on the free energy perturbation-absolute binding free energy prediction (FEP-ABFE) method, which samples microscopic states using MD or Monte Carlo simulations to predict macroscopic properties (such as properties related to binding affinity) of the target system. While FEP-ABFE can achieve good accuracy, it has an extremely high demand for computational resources, which hampers its use for large-scale drug screening. To address this issue, the researchers developed, among other techniques, a customized job management system to run this method in a scalable manner on the new generation of the Tianhe system, currently the seventh fastest supercomputer in the world. They virtually screened more than 3.6 million compounds from commercially available databases using docking methods, followed by the FEP-ABFE calculations for about 12,000 compounds to achieve FEP-based binding free energy results.

Other teams focused on better understanding different stages of the life cycle of SARS-CoV-2 with modeling and simulations. For instance, Arvind Ramanathan and colleagues explored the replication mechanism of SARS-CoV-2 in the host cell, which can provide insights into drug design³. Cryo-EM techniques have helped to elucidate the structural organization of the viral-RNA replication mechanism, but the overall resolution of the data is often poor, hindering the complete understanding of this mechanism. This team of researchers developed an iterative approach to improve the resolution within cryo-EM datasets by using MD simulations and finite element analysis. One of the challenges of the approach was the coupling of different resolutions, which was intelligently done by leveraging machine learning algorithms. In order to help balance the workload, the researchers used a single coordinated workflow across multiple geographically dispersed supercomputing facilities: Perlmutter, which is currently the fifth fastest supercomputer in the world and located in Berkeley, California, and ThetaGPU, which is an extension of Theta, currently the seventieth fastest supercomputer in the world and located at the Argonne National Laboratory in Illinois.

Makoto Tsubokura and colleagues, the winners of this year's special prize,

turned their attention to how COVID-19 is transmitted via droplets and aerosol particles⁴. To better understand and evaluate the risk of infection caused by droplets and aerosols, this team of researchers focused on simulating how droplets reach other individuals after being emitted from an infected person and transported through the air. These end-to-end simulations must take into account complex phenomena and geometries, including the surrounding environment, the physics of droplets and aerosols, any air flow activity induced by other elements (such as air-conditioning systems), the amount of people around, and so forth. The researchers implemented different computational fluid dynamics techniques to be able to scale the simulations to the Fugaku system, which is currently the fastest supercomputer in the world. Such simulations generated digital twins representing different transmission scenarios, and the results were widely communicated by the media and used to implement public policies in Japan.

Rommie Amaro and colleagues also focused their work on the airborne transmission of SARS-CoV-2: they developed a multiscale framework to study aerosolized viruses⁵. Studying these complex systems requires taking into account different resolutions (from nanometers to approximately one micron in size) and longer timescales (spanning microseconds to seconds): the multi-resolution requirement makes all-atom MD simulations very challenging and computationally expensive. Among the many technical contributions, the researchers ran and scaled the MD simulations on the Summit supercomputer,

allowing them to develop an impressive one-billion-atom simulation of the aerosolized version of the SARS-CoV-2 Delta virion — the first ever simulation of a respiratory aerosol (see image). Such a simulation allows one to explore the composition, structure, and dynamics of respiratory aerosols, and can serve as a basis to develop new therapeutic solutions for COVID-19 (for instance, to find potential binding sites).

Last but not least, another finalist team focused on a different challenge: performing epidemic simulations. Madhav Marathe and colleagues developed a framework to generate real-time scenario projections, assessing the likelihood of epidemiological outcomes for possible future scenarios⁶. This can be used, for instance, to better allocate vaccine supplies, to evaluate the role of vaccine hesitancy, and to understand the impact of waning immunity, among other important studies for public health use. As part of their framework, the researchers built a digital twin of a time-varying social contact network of the United States using various national-scale datasets. This digital twin can then be brought to life by contextualizing it with current real-world conditions, using, again, different datasets of varying scales. After initializing the digital twin, a parallel agent-based socio-epidemic simulator, also developed by the researchers, can then be used to generate and analyze different scenarios. Because the simulations are very computationally intensive, a meta-scheduler for HPC clusters was explored in order to allow for the use of multiple clusters. In their analyses, the team used two supercomputers: Bridges-2, located at the Pittsburgh Supercomputing Center,

and Rivanna, located at the University of Virginia. It is worth noting that the team has been doing scenario projections from the start of the pandemic for various state and federal agencies in the United States.

Overall, all of these research works presented not only technical contributions, which can certainly be used in other HPC applications, but also important frameworks and studies that can be used to improve our understanding of the ongoing pandemic and to better implement policies to decrease the spread of the virus. As new HPC technologies and new computing architectures are developed, such as supercomputers with exascale capabilities, more remarkable advances are expected from the computational science community. □

Published online: 13 December 2021
<https://doi.org/10.1038/s43588-021-00180-2>

References

1. Blanchard, A. E. et al. Language models for the prediction of SARS-CoV-2 inhibitors. In *Int. Conf. High Performance Computing, Networking, Storage, and Analysis* (ACM, 2021); https://sc21.supercomputing.org/proceedings/tech_paper/tech_paper_pages/gbv102.html
2. Li, Z. et al. FEP-based large-scale virtual screening for effective drug discovery against COVID-19. In *Int. Conf. High Performance Computing, Networking, Storage, and Analysis* (ACM, 2021); https://sc21.supercomputing.org/proceedings/tech_paper/tech_paper_pages/gbv105.html
3. Trifan, A. et al. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.10.09.463779> (2021).
4. Ando, K. et al. Preprint at <https://arxiv.org/abs/2110.09769v3> (2021).
5. Dommer, A. et al. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.11.12.468428> (2021).
6. Bhattacharya, P. et al. Data-driven scalable pipeline using national agent-based models for real-time pandemic response and decision support. In *Int. Conf. High Performance Computing, Networking, Storage, and Analysis* (ACM, 2021); https://sc21.supercomputing.org/proceedings/tech_paper/tech_paper_pages/gbv103.html