

Solution Brief

Molecular Modeling
Intel® MPI Library
Intel® oneAPI Math Kernel Library
Intel® Compiler



Advancing Medical Macromolecular Models and Pharmaceutical Research with ByteDance

The ByteDance Volcano Engine, enhanced by the powerful computing capabilities of the Intel® oneAPI HPC Toolkit, accelerates VASP pharmaceutical macromolecular modeling to help facilitate innovative medical research and drug discovery.



About ByteDance

ByteDance is a leading Internet technology company renowned for driving the success of acclaimed products like TikTok, Douyin, and Toutiao in the global and Chinese markets. Building upon this legacy, ByteDance has unveiled its latest subsidiary brand, Volcano Engine. Volcano Engine offers the comprehensive technology suite responsible for the success of ByteDance's various consumer products, available as a commercialized service designed to help enterprise clients digitize their operations. This solution brief examines how Volcano Engine's AI algorithms, data insight services, and cloud computing solutions contribute to the success of medical macromolecular modeling for pharmaceutical research and drug discovery.

The Challenge of Macromolecular Modeling: Simulating Complex Biological Systems

Medical macromolecular modeling is a highly advanced computational strategy in the medical research and drug discovery field. Specifically, it enables the study and simulation of large biological molecules, or macromolecules, that facilitate essential biological processes in living organisms. By leveraging this computer modeling to understand the structure, function, and interaction of macromolecules, innovative companies can study and predict complex biological systems for biopharmaceutical research, drug discovery, pathology, and many other areas of research. This form of modeling is crucial to gaining medical insights that would be challenging, dangerous, or unethical to obtain through experimental techniques alone, and helps researchers design experiments.

While medical macromolecular modeling is a powerful tool, it also introduces a unique set of challenges regarding computational performance due to the complexity of the biological systems the model is simulating. Consequently, the high performance demands of macromolecular modeling impose significant costs, potentially deterring companies without substantial financial resources from leveraging this transformative technology.

These computational challenges include:



Computational Strain:

At present, simulating entire cellular systems at atomic detail is impossible with existing computational resources due to the sheer volume of data. To overcome this limitation, researchers narrow their simulations to specific areas of interest.

Still, performing complex computations across large macromolecules composed of thousands of atoms, across time scales spanning microseconds or more, requires significant computational resources. These models must be sophisticated enough to explain how molecular-level simulations drive cellular processes and bridge together different simulation scales to provide valuable insights.



Macromolecule Variations:


Macromolecules fluctuate and adopt a multitude of conformations, which impacts the role of their function. Simulating the full range of their possible conformations to predict different outcomes is difficult to capture and further adds to the computational load. The computing power must be high-performance to conduct simulations devoid of oversimplification, as this can compromise the accuracy of the simulation results.



Long Data Timelines:

Many biological processes occur gradually over long periods of time that are challenging to capture through standard data modeling techniques. Modelers may use enhanced techniques that accelerate molecular processes to overcome these limitations but require advanced computational models and power to facilitate these methods.

As evidenced by this range of challenges, the nature of medical macromolecular modeling poses unique computational challenges and high expenses. Researchers in the pharmaceutical field need a more cost-effective solution that offers high compute power and highly sophisticated AI models to continue advancing simulation results that reflect the most accurate biological realities.



The Need for Enhanced Macromolecular Modeling

ByteDance is a multinational Internet technology company who decided to meet this demand. They established Volcano Engine, a subsidiary, to provide powerful computing capabilities for macromolecular modeling. Volcano Engine is a leading cloud technology service provider that delivers high-performance, cost-effective computing and model training capabilities to support the innovation of new drug research and development within pharmaceutical firms. Complementing these offerings, Volcano Engine also provides flexible cloud storage options to strike an optimal balance of cost and performance for medical macromolecular models.

Many of ByteDance's clients rely on the Vienna Ab initio Simulation Package (VASP) software to perform macromolecular modeling on the Volcano Engine server. For biomedical technology clients looking to scale and enhance the performance of VASP macromolecular models for more intricate calculations, ByteDance offers the significant advantage of being able to facilitate this performance optimization thanks to its partnership with Intel. The following analysis explores how ByteDance leverages Intel® technologies to maximize available computational resources to deliver scalable VASP macromolecular modeling.

Empowering VASP Medical Macromolecular Models: Intel and Volcano Engine Join Forces to Deliver High-Performance, Cost-Effective Computing

To run VASP simulations of larger or more intricate macromolecular systems, ByteDance needed to further enhance its computing platform, Volcano Engine, to reduce the latency of its services accessed in the cloud. Based on the characteristics of the VASP macromolecular model, ByteDance's engineers determined that Intel® technologies, specifically, the Intel® oneAPI HPC Toolkit, offered the optimal combination of performance and price to accelerate the VASP macromolecular model without incurring excessive costs on customers.



Volcano Engine provides customers with secure, reliable, high-performance, and easy-to-maintain enterprise-level cloud services. Intel® oneAPI Toolkits brings ByteDance another avenue to achieve better performance for HPC applications.”

— ByteDance Solution Architect in Beijing

Intel® Xeon® Scalable Processors

ByteDance leveraged Intel® Xeon® Scalable Processors to optimize their platform, using a compact, high-performance CPU with high memory bandwidth to support their most demanding workloads in the cloud. Additionally, Intel® Xeon® Scalable Processors provide ByteDance with a valuable resource: the integrated Intel® Accelerator Engines, specifically the Intel® HPC Engine. This accelerator supports fast-growing workloads to accelerate Volcano Engine's performance, improve power efficiency when conducting complex calculations, and reduce the costs of macromolecular modeling to ByteDance's customers. Volcano Engine also adapts VASP parameters based on Intel's CPU core number, further optimizing its performance and resource utilization.



Intel® oneAPI HPC Toolkit

The Intel® oneAPI HPC Toolkit, an add-on to the Intel® oneAPI Base Toolkit, offers a comprehensive range of tools developers can use to build, analyze, optimize, and scale high-performance computing applications, such as the VASP macromolecular modeling. This toolset helped Volcano Engine accelerate VASP macromolecular models on Intel® Xeon® processors with fast, scalable, and reliable parallel code whose workload can be distributed across different nodes to improve development efficiency while reducing R&D costs.

By installing the Intel® oneAPI HPC Toolkit on Volcano Engine's cloud service, ByteDance was able to leverage the Intel® MPI Library, Intel® Compiler, and Intel® oneMKL to take advantage of the latest techniques in vectorization, multithreading, multi-node parallelization, and memory optimization.



Intel® oneAPI Toolkit: Intel® MPI Library

Volcano Engine leveraged the Intel® MPI Library to enable a more collaborative and intelligent drug development program, and optimized the VASP molecular dynamics for NUMA affinity.

Message Passing Interface (MPI) is a set of user-friendly programming interfaces that developers can use to facilitate seamless message exchange between diverse CPUs or server nodes used in calculations in scientific computing programs. Leveraging this set of interfaces, clients' engineers can quickly write parallel computing programs that can be ported across platforms to improve development efficiency. In summary, the integration of the Intel® MPI library with Volcano Engine enables more users to support VASP functions with reports of strong performance. The Intel® MPI Library also undergoes extensive testing and validation to ensure the reliability and stability of MPI communication.


This allows users to run complex applications in large-scale parallel environments without encountering communication errors or deadlocks. Additionally, the Intel® MPI Library integrates many useful features such as process management, data distribution, and communication debugging tools to simplify the development and debugging process of parallel applications.

Intel® oneAPI Toolkit: Intel® oneMKL

The Intel® oneMKL (Math Kernel Library) enhanced the performance of the VASP software by providing optimized mathematical functions. By replacing standard mathematical functions in the VASP medical macromolecular model with their Intel® oneMKL counterparts, ByteDance was able to make Volcano Engine deliver faster and more accurate computations, enhancing the solution's overall performance.

Intel® oneAPI Toolkit: Intel® Compiler

Intel® Compiler, a standards-based compiler designed to function across various architectures, played a pivotal role in aiding Volcano Engine's engineers. They employed it to rebuild the VASP code, evaluate different optimization configurations offered by Intel® Compiler, and successfully compile and rebuild an optimized version of the VASP program. This led to the achievement of faster simulations.



With the integration of the Intel® oneAPI HPC Toolkit and third-generation Intel® Xeon® Scalable processors in Volcano Engine, ByteDance was able to significantly lower the latency of VASP macromolecular modeling conducted on Volcano Engine.¹

Key Benefits

Pharmaceutical research companies can now better leverage advanced medical macromolecular modeling through the collaborative efforts of ByteDance and Intel. This partnership offers a range of benefits designed to reshape research methodologies:

Accelerated Drug Research

Leverage VASP macromolecular modeling to rapid test hypotheses without needing to conduct time- and resource-intensive experiments, resulting in the quicker development of life-changing medications.

Improved R&D Success Rate

Raise the success rate in your team's research and development endeavors to contribute to more impactful outcomes.

Optimized Computing Performance

Leverage cutting-edge technology to improve the accuracy and speed of intricate medical macromolecular simulations to accelerate drug research and development.

Reduced R&D Costs

Optimize resource utilization that translates into cost savings across research and development efforts.

Conclusion

For many biopharmaceutical companies, medical macromolecule modeling is crucial for pioneering medical breakthroughs and transformative drug discovery. With this mission in mind, ByteDance created Volcano Engine —a cloud-native computational platform designed to serve the high-speed and complex computing demands of this field.

Leveraging Intel® technologies, ByteDance further improved the latency and performance of its VASP operation. This collaboration unlocked significant performance gains, paving the way for novel discoveries and breakthroughs. These accomplishments not only set a precedent for future clients but also showcase the transformative impact of cutting-edge technology in scientific research. For organizations looking to optimize their computational capabilities in medical macromolecular modeling and amplify their research outcomes, ByteDance's team invites you to explore the possibilities offered by Volcano Engine.


Learn More

To learn more about the ByteDance Solution visit:

- [Intel Powers Volcano Engine to Accelerate VASP Pharmaceutical Molecular Modeling](#)
- [ByteDance Website](#)

To learn about Intel® technologies visit:

- [Intel® Xeon® Scalable Processors Product Page](#)
- [Intel® oneAPI Toolkit Product Page](#)
- [Intel® MPI Library Product Page](#)
- [Intel® oneAPI Math Kernel Library Product Page](#)



Sources

¹The test date is September 2022. Internal data of ByteDance: Intel joins hands with Volcano Engine to empower VASP medical molecular model computing with high-performance computing power. Intel does not control or audit third-party data. Please review this content, consult other sources, and confirm that the data mentioned is accurate.

Notices & Disclaimers

Intel is committed to respecting human rights and avoiding complicity in human rights abuses.

See Intel's [Global Human Rights Principles](#). Intel® products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

Intel technologies may require enabled hardware, software or service activation. No product or component can be absolutely secure. Your costs and results may vary. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.