



# Cores

**1** Simple Distributed Computing

“Leverage existing investments through increased resource efficiency”

#### **Cores – Simple Distributed Computing**

Cores is a new platform for distributed computing specifically created for working with complex multi-level algorithms. Uniquely, it lets the user focus on obtaining end results, while distributing calculations in the most efficient way to any computational resource you present to it.

#### **It's about results**

Instead of specifying in detail what to run where, and how, Cores allows you to formulate a complex multi-part computational project – such as the calculation of a drug binding affinity, or the rendering of a complex scene – as a single task, and compose them to suit your needs. Cores then takes care of decomposing the task into its parts, and run them. This way, you can focus on the issues behind the question you're trying to solve, not the details of which algorithm is optimal in which situation.

#### **Maximize your infrastructure utilization**

Cores allows you to maximize your hardware investments by consolidating all your computational resources and making them available. It also makes it easy to scale out to external resources, such as public clouds, when more computational power is required.


#### **Inspect, compose and combine calculations**

In computing, it can be important to know what's going on under the hood: what's the underlying cause of an unexpected result?

How does a specific input value affect the final output? With Cores, any input or output of any part of a project can be monitored at any point during a calculation. Also, any output can be used as an input for a new calculation – even while the original calculation is still running. This way, results can be easily, and meaningfully combined to form the basis for new calculations if needed.

#### **Adapt any calculation and run any application**

Inevitably, you will need to run custom-made software that is specific to your needs. Cores has extensive plug-in facilities that not only allow you to save your composed projects, but also allow you to incorporate any program or script into the Cores platform. Any type of calculation can be performed this way, with real-time intermediate and post-processing.



# 3

Cores



# 4

## Cores

### Compute anywhere

Once presented with a project, Cores will divide it into individual calculations or simulations that can be run in parallel. This can be done on one or more local and remote compute clusters, but also on collections of workstations or cloud computing instances. The networked nature of the Cores platform allows it flexibility over where a part of a calculation can run, so it can schedule and match resources so that no cycles are wasted or used inefficiently.

This means that if a user needs more processing power for a calculation while it's running, Cores will redirect and prioritize the data to the most efficient resource or burst it to the cloud, without having to restart any process.

### Distributed Storage and data processing

A Cores network keeps track of where data is stored, and can schedule jobs close to where the data is, making optimal use of high-bandwidth links if they are available. As a high-performance distributed storage solution, it presents all of that data not as a large collection of files, but in the way that they relate to the computations that are being performed. Output data can be accessed as soon as it is generated, even if other parts of a computation are still running.

### Control your data

With Cores, you decide where you run, and where you keep your data. All communication in Cores happens through SSL (TLS) with client and server-side authentication. Because the decision of where to run can be made on a per-project basis, the decision to burst to the Cloud can be made on an as-needed basis with security in mind.

### The team behind Cores

The researchers behind Cores have many years of experience with simulations. The team consists of key developers of Folding@home, the world's most powerful network for distributed computing, and Gromacs, the fastest and one of the most widely used software packages for molecular dynamics.

The daily work of the researchers consists of speeding up processes. They are experts in parallel computing and make them suitable to every kind of resource. Extensive experience has gathered from simulations being done on some the world's largest and powerful supercomputers.



“Boost performance and shorten time to solution”



# 6

## Molecular Dynamics Example: Gromacs on Cores

### A new paradigm in molecular simulations

The idea for Cores began when it was realized that computer resources are more and more parallel, and sampling algorithms (such as BAR free energy calculations, Markov state modeling and transition path sampling) are more and more sophisticated. Combining running, post-processing and partitioning works especially well for molecular simulations.

### Gromacs

Cores is a product developed by several key developers of the Gromacs molecular simulation package, who have years of experience with simulations. Given its background, it should be no surprise that Cores' built-in support for Gromacs is extensive. Gromacs is a high-performance simulation package with the following advantages:

- large academic and non-academic user base
- high performance through hand-tuned assembly kernels and state-of-the-art parallelization
- extensive documentation and large collection of third-party post-processing tools


### Optimized performance

Running a high-performance molecular simulations package like Gromacs in parallel usually requires careful tuning to get the best performance. Because Cores is aware of all connected computational resources, and of the larger context in which individual simulations are

run, Cores can automatically optimize individual Gromacs simulations and the entire calculation to match the available resources. Because the higher-level sampling happens in parallel, Cores achieves total effective scaling that is much better than single MD runs can hope to achieve.

### Folding a protein in 30 hours

One of the algorithms Cores makes available out-of-the box is Markov state modeling: a sampling technique that identifies typical states of a molecule and calculates the rate at which these states change. In an example application, we have folded a protein with full atomic detail (9864 atoms) in 30 hours on 5376 cores, while starting from fully extended conformations. This presents a 24-fold speedup over previous attempts, and effective scaling down to 2 particles per core.



Calculated (blue) and experimental (white) structure of a protein after 30 hours of Markov state modeling in Cores.

Address	BinaryBio AB Stureplan 15 111 45 Stockholm	Telephone Web Email	+46 (0) 707 98 36 39 <a href="http://www.binarybio.com">www.binarybio.com</a> <a href="mailto:info@binarybio.com">info@binarybio.com</a>
---------	--	---------------------------	--

BinaryBio AB is a Swedish company based on scientifically leading research in the field of bioinformatics, molecular dynamics, and high performance computing.

Working in close collaboration with the Center for Biomembrane Research at Stockholm University and the Royal Institute of Technology in Stockholm, BinaryBio offers its customers state of the art technology for high performance parallel and distributed computing with focus on biomolecular simulations and bioinformatics.

