

Case Study: University of Haifa



BACKGROUND

The University of Haifa is focusing on the analysis of the massive amount of digital data in the areas of virology, cancer biology and biofuels.

CHALLENGE

To place an entire graph of genome assembly or processing a huge hash table of regulatory motives into RAM.

SOLUTION

vSMP Foundation software to create a single system with 192 cores and 1.5 TB of memory.

“Thanks to vSMP Foundation memory aggregation technology, we experienced a significant increase in performance for many heavy tasks. We are now able to use a variety of applications on a single system and are able to enhance our research”

Professor Leonid Brodsky, Ph. D., Senior Lecturer, Department of Evolutionary and Environmental Biology and Tauber Bioinformatics Research Center, University of Haifa

BACKGROUND

The Tauber Bioinformatics Research Center at University of Haifa is participating in a number of high profile biological and medical projects (virology, cancer biology, biofuel) and is focusing on the analysis of the massive amount of digital data generated by these projects. As a part of this activity, the Center has developed an alpha version of an expert computational system for the intelligent analysis of the Next Generation sequencing and mass spectrometry (proteomics and metabolomics) data: with this system we intend to service Haifa University life science projects as well as the Center's collaborative projects.

CHALLENGE

The major IT problem we are dealing with is to provide big shared memory for the multiprocessor analysis of high throughput biological data. In particular, a lot of shared memory is required (up to 6 TB) to perform the assembly of large plant genomes and for working with "junk DNA" genome information – regulatory motives distributed over the entire genome. The algorithms require big shared memory on a single machine: for instance, to place an entire graph of genome assembly or processing a huge hash table of regulatory motives into RAM.

ABOUT SCALEMP

ScaleMP is the leader in virtualization for high-end computing, providing maximum performance and lower total cost of ownership. The innovative Versatile SMP™ (vSMP) architecture aggregates multiple independent systems into a single virtual system, delivering an industry-standard, high-end symmetric multiprocessing computer.

SOLUTION

The vSMP Foundation technology has helped solve for us the problems associated with big shared memory. In addition, the aggregation of servers has enlarged the number of CPU cores we can now use on a single machine by our OpenMP-based algorithms.

The vSMP Foundation software currently runs on the system of 3 identical 1U Supermicro servers interconnected via 40 Gbit IB, and aggregates 192 cpu cores and 1.5 TB of RAM.

At the moment our results are not business oriented but are scientifically oriented ones. Together with our collaborators we have published several papers in high profile journals such as Nature and Proceedings of the National Academy of Science.

Several OpenMP-based algorithms coupled with the big memory use were developed and approved on a ScaleMP machine. These algorithms were successfully applied to big genome assembling and whole genome regulation tasks. In addition, a public domain ALLPATHS-LG application (it does not support distributed computing using MPI, instead it uses Shared Memory Parallelization) is being used on our ScaleMP server.