

Solving large Nonlinear Systems of ODE with Hierarchical Structure Using Multi-GPGPUs and an Adaptive Runge Kutta(ARK)

Ahmad Al-omari, Jonathan Arnold, Thiab Taha, Heinz-Bernd Schuttler
 Institute of Bioinformatics, The University of Georgia, Athens, GA 30602, aomari@uga.edu

Introduction

A graphics processing unit (GPU) offers a solution to a very important and fundamental problem of interest to many researchers, a problem that would be prohibitive to solve without the technology of GPUs. The problem is how the biological clock controls the rhythms of ~2436 genes in the genome (with 11,000 genes) of a model system, the filamentous fungus, *Neurospora crassa* (Dong, et al.,2008). Ultimately, we want to be able to predict and hence understand the dynamics of all of these genes and their products in a genetic network describing how the clock functions. The problem of identifying such large genetic networks is beyond the current capability of the largest and fastest serial computers in existence. In order to overcome this problem, we developed an algorithm that is described below helps us to begin to chip away at the entire circadian network, a network that is found from bacteria to humans(Lakin-Thomas, et al., 2011)

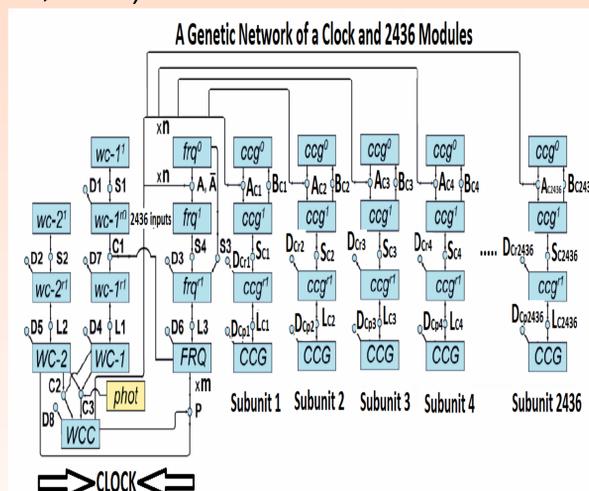


Figure 1. The genetic network to be solved by the GPUs[2436 subunits]. The subunits are independent from each other but on the clock. Solving this huge network is beyond the fastest computer in existence.

Methods

Procedure Flowchart:

All of the slave modules appear in Figure1 has the same mathematical ODE form but different parameters:

$$\begin{aligned} dg_0/dt &= B_c g_1 - A_c g_0 w(t) \\ dg_1/dt &= A_c g_0 w(t) - B_c g_1 \\ dg_r/dt &= S_c g_1 - D_c g_r \\ dg_p/dt &= L_c g_r - D_{cp} g_p \end{aligned}$$

CPU sends the master module solution[w(t)] and the parameter set to each slave module

GPU(s) solve(s) the 2436 ODEs of the slave modules in parallel using the(ARK) method

Each thread block executes a kernel, which contains the ARK ODE solver. All threads in the same block hold the same data and execute the same system of ODE for an assigned slave module.

GPU sends the solution of the 2436 systems of ODE back to the CPU

CPU calculates a cost function (Chi-square using Metropolis procedure) that minimizes the difference between the experimental data and the fitted solution

Results

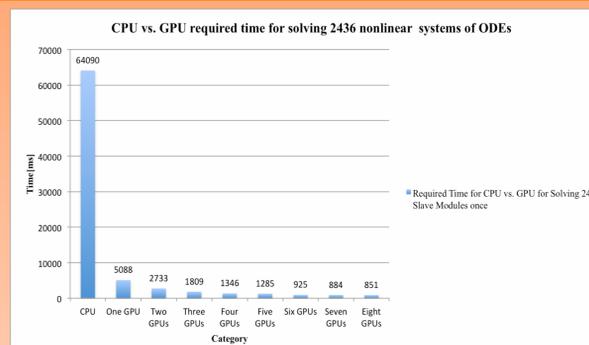


Figure 2. The time required for solving 2436 slave modules just one time using a NVIDIA GPU(s) [Kepler K-20x Tesla] over an extreme edition of optimized CPU [Quad-cores CPU [Intel(R) Core(TM) i5-2400 CPU@ 3.10GHz]. Using four of the GPUs to solve our target genetic network 800,000 times shown in the (Fig. 1) to fit the observed data requires just twelve days while using the CPU requires a year and six months.

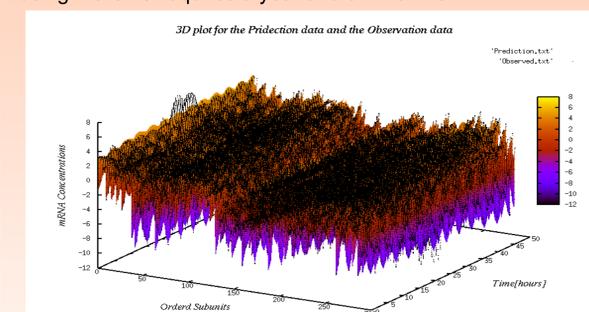


Figure 3. Sample of our output for just 292 subunits, all of them are under the control of the WCC. Black dots shows the observed data that fit the prediction data (in orange and purple) very well.

Model of solution procedure

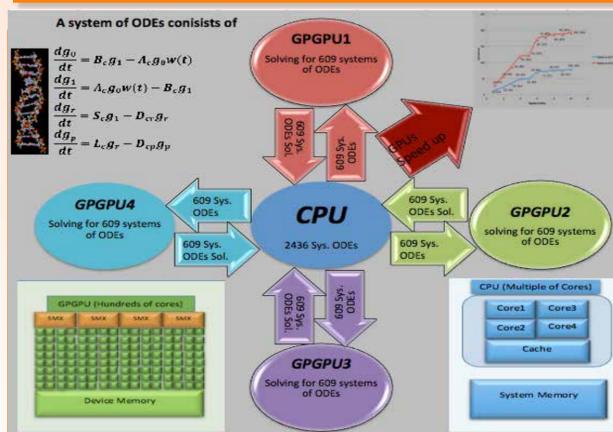


Figure 4. Solution's Model

Conclusion

We have demonstrated that we can implement the ensemble method (Yu et al., 2007). We have also demonstrated that we can solve the ODEs on GPUs for the slave module in **Figure3**. We propose to:

- ◆ Implement an ensemble method on the CPU and GPUs for identifying the network in **Figure1**.

We have developed an algorithm on the GPU able to solve the genetic network of 2436 genes shown in **Figure1** within about 12 days instead of more than 18 months on the CPU.

We have more than 31,668 data points for fitting the proposed large genetic network for the clock (Dong et al., 2008). We propose to:

- ◆ Fit the genetic network in **Figure1** by the best parallelization strategy on the CPU and GPUs to our microarray data. The resulting network will enable us to examine how the clock controls ribosome biogenesis by a detailed molecular mechanism.

References

Dong, W., et al. (2008) Systems biology of the clock in *Neurospora crassa*, *PLoS one*, 3, e3105.
 Huang, Y. (2007) Parameter estimation of chemical reaction networks. A Master thesis, Physics Department, The University of Georgia.
 Lakin-Thomas, P.L., Bell-Pedersen, D. and Brody, S. (2011) The Genetics of Circadian Rhythms in *Neurospora*, *Advances in genetics*, 74, 55.
 Yu, Y., et al. (2007) A genetic network for the clock of *Neurospora crassa*, *Proceedings of the National Academy of Sciences*, 104, 2809-2814.

Acknowledgment

CUDA Teaching Center at The University of Georgia