

# Accelerating Scientific Workflow Applications with GPUs

Dustin Shahidehpour, Scott J. Krieder, Benjamin Grimmer, Jeffrey Johnson  
Justin M. Wozniak\*\*, Michael Wilde\*\*, Ioan Raicu\*

Department of Computer Science - Illinois Institute of Technology, Chicago, IL, USA

\*Mathematics and Computer Science Division, Argonne National Laboratory

\*\*Computation Institute – University of Chicago, Chicago, IL, USA

## Abstract

This work analyzes the performance increases gained from enabling Swift applications to utilize the GPU through the GeMTC Framework. By identifying computationally intensive portions of Swift applications, we can easily turn these code blocks into GeMTC microkernels. Users can then call these microkernels throughout the lifetime of their Swift application. The GeMTC API handles task overlap and data movement, providing transparent GPU acceleration for the user. This work highlights preliminary performance results from the scientific application MDPProxy. This application determines the energy of particles in a modeled universe as they move around in space.

## What is GeMTC?

GeMTC (GPU enabled Many-Task Computing), is a CUDA-based framework which provides efficient support for Many-Task Computing workloads on accelerators. The GeMTC framework has been integrated into Swift/T, a parallel programming framework from Argonne National Laboratory and the University of Chicago, providing GPU functionality for the Swift language.

## What is a microkernel?

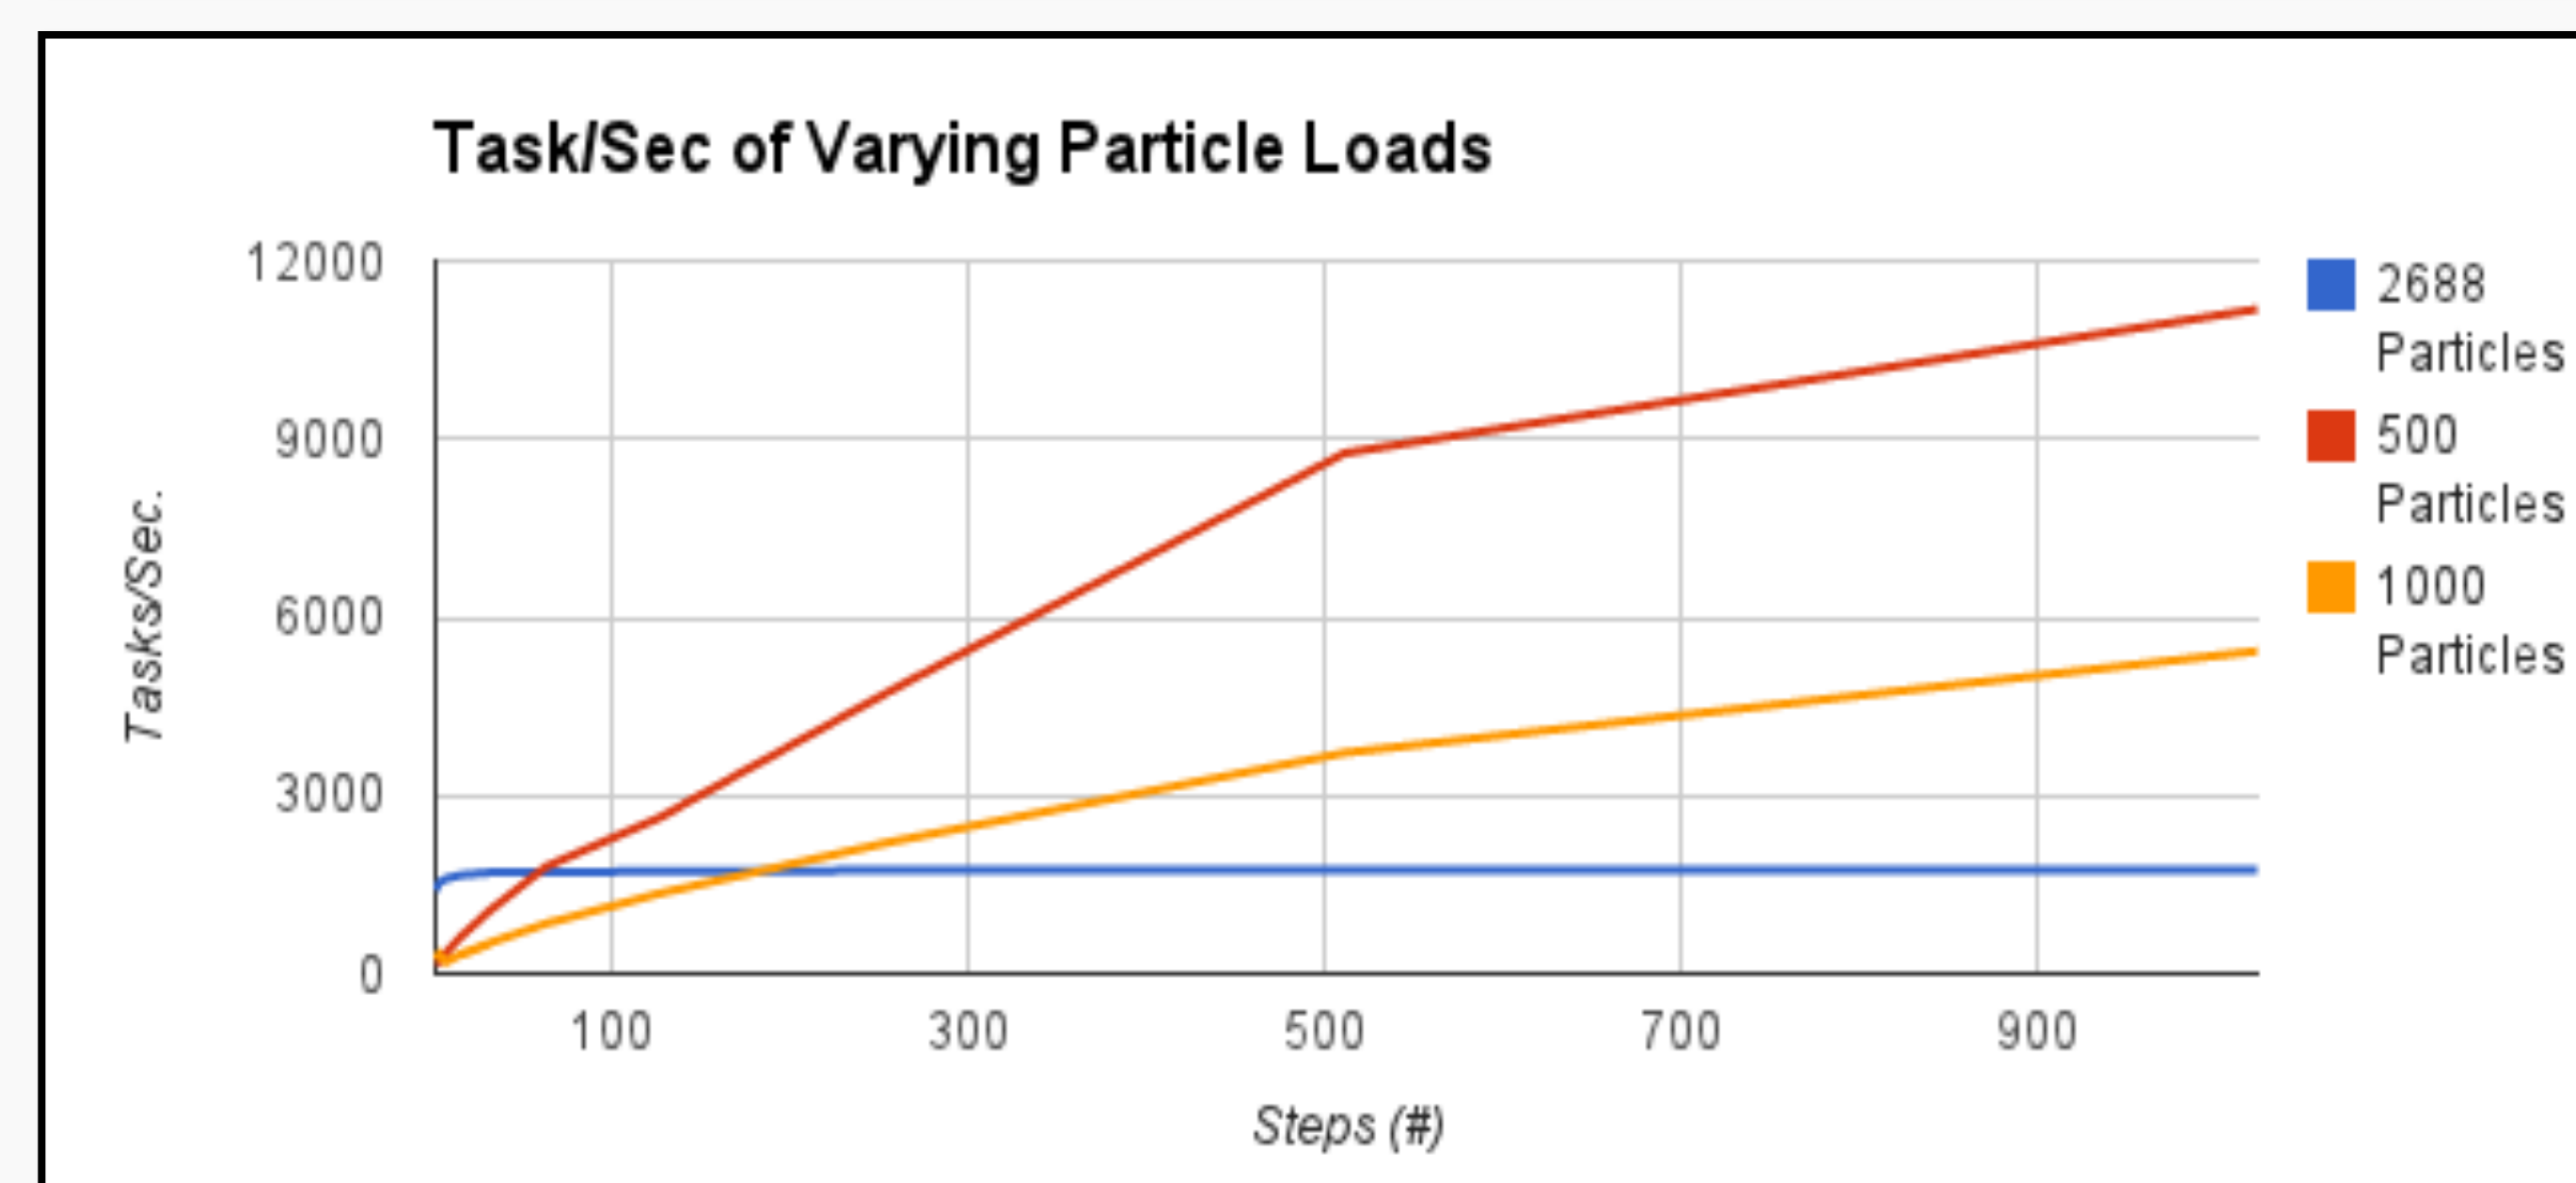
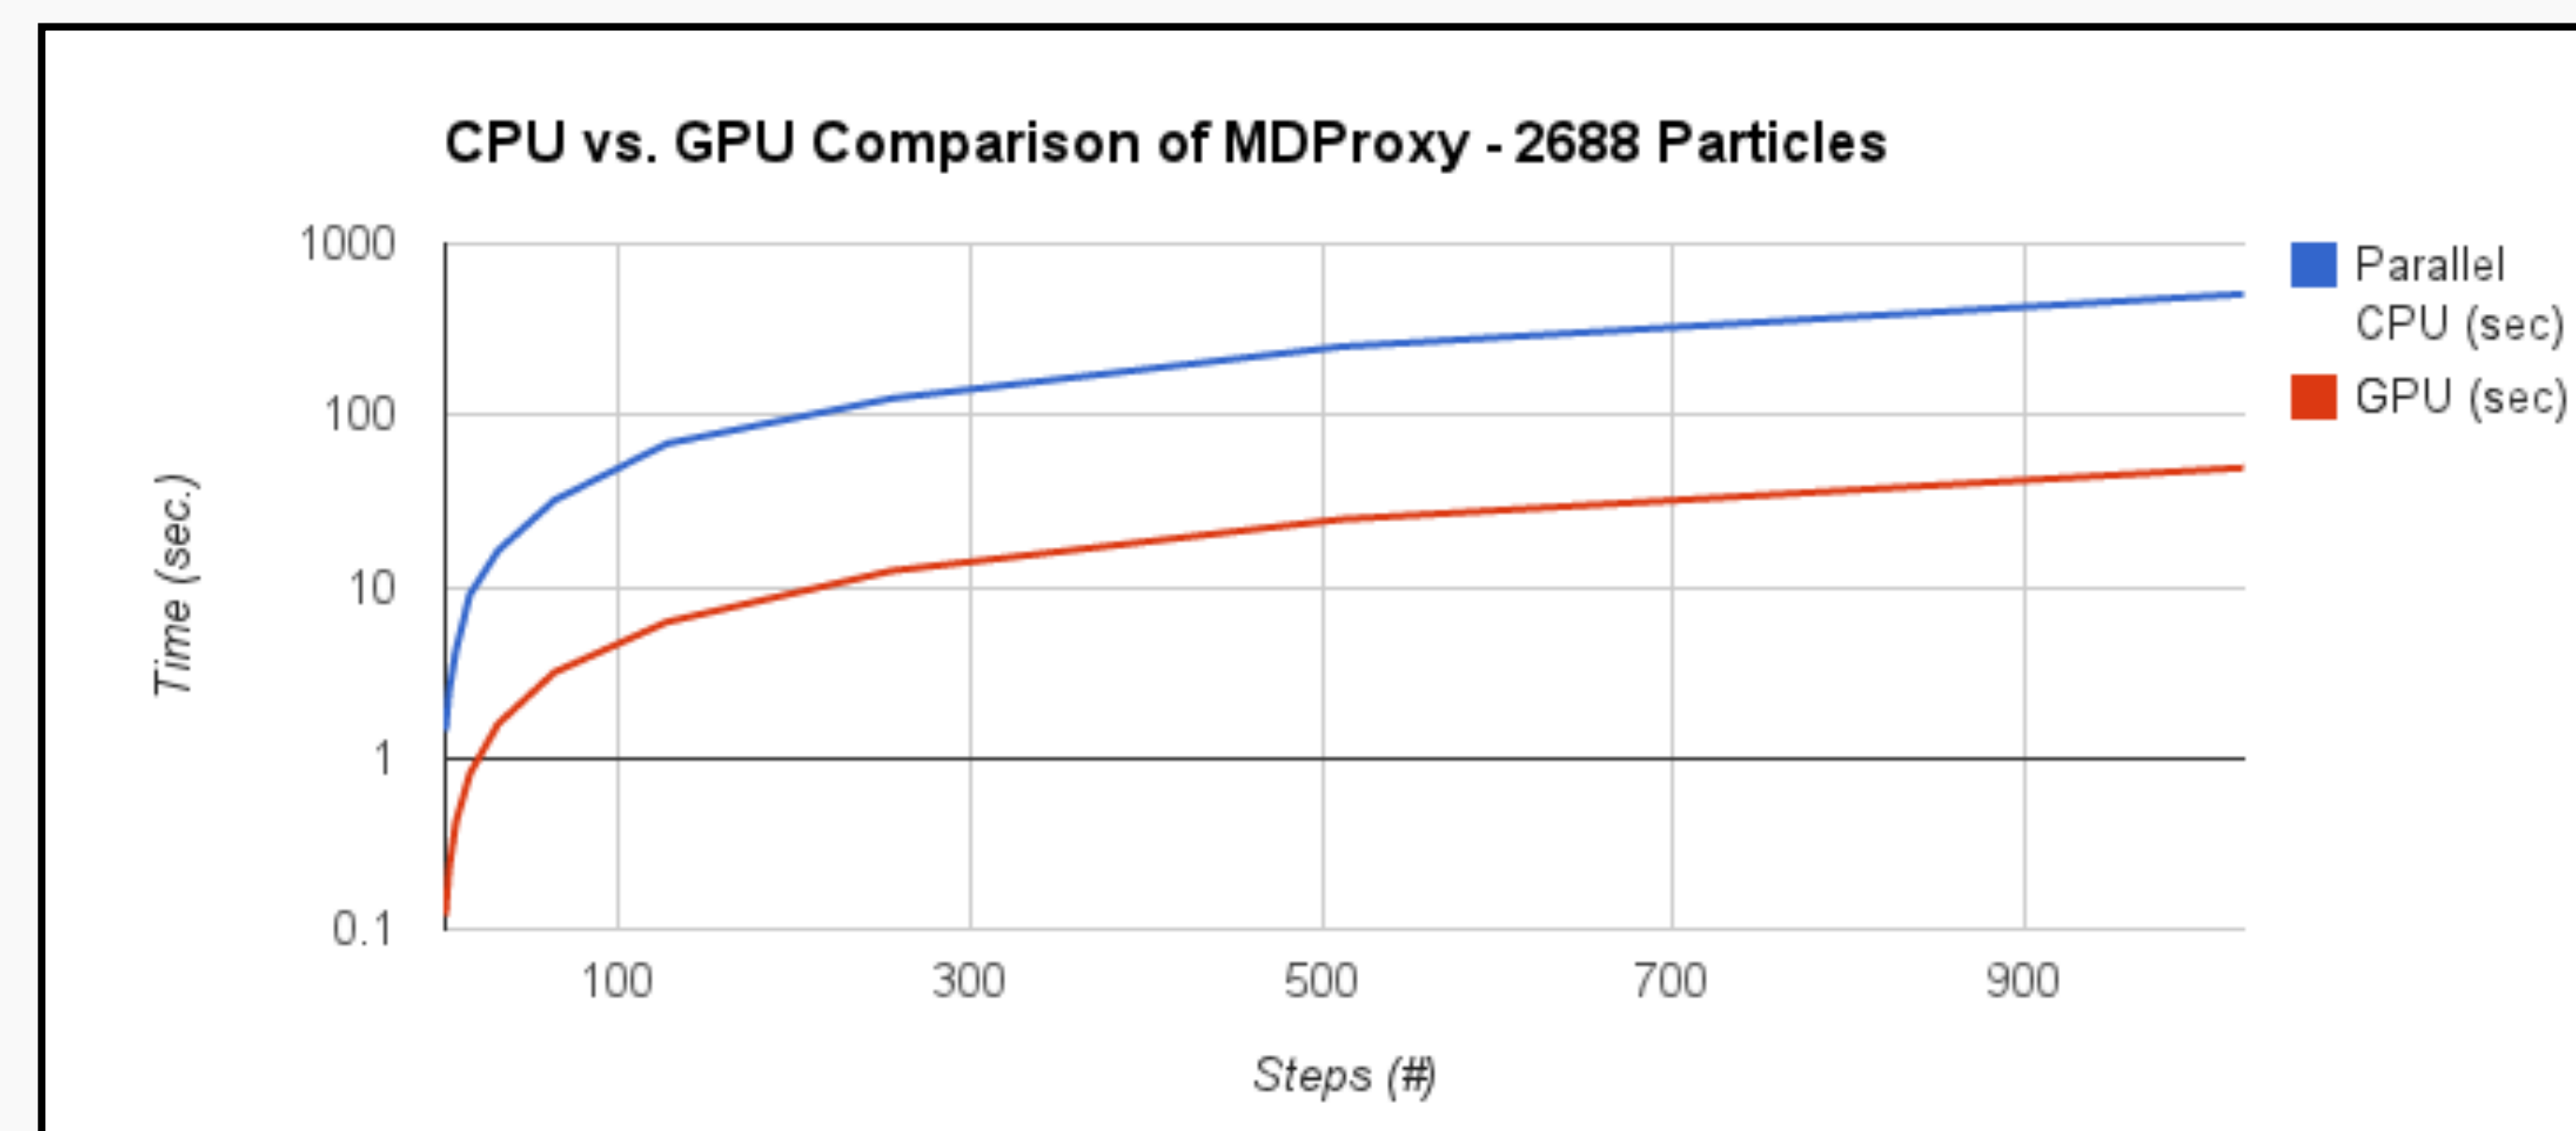
A microkernel is a traditional CUDA kernel that is modified to run in the GeMTC framework. A CUDA kernel is a user-defined function that runs on a NVIDIA GPU.

## References

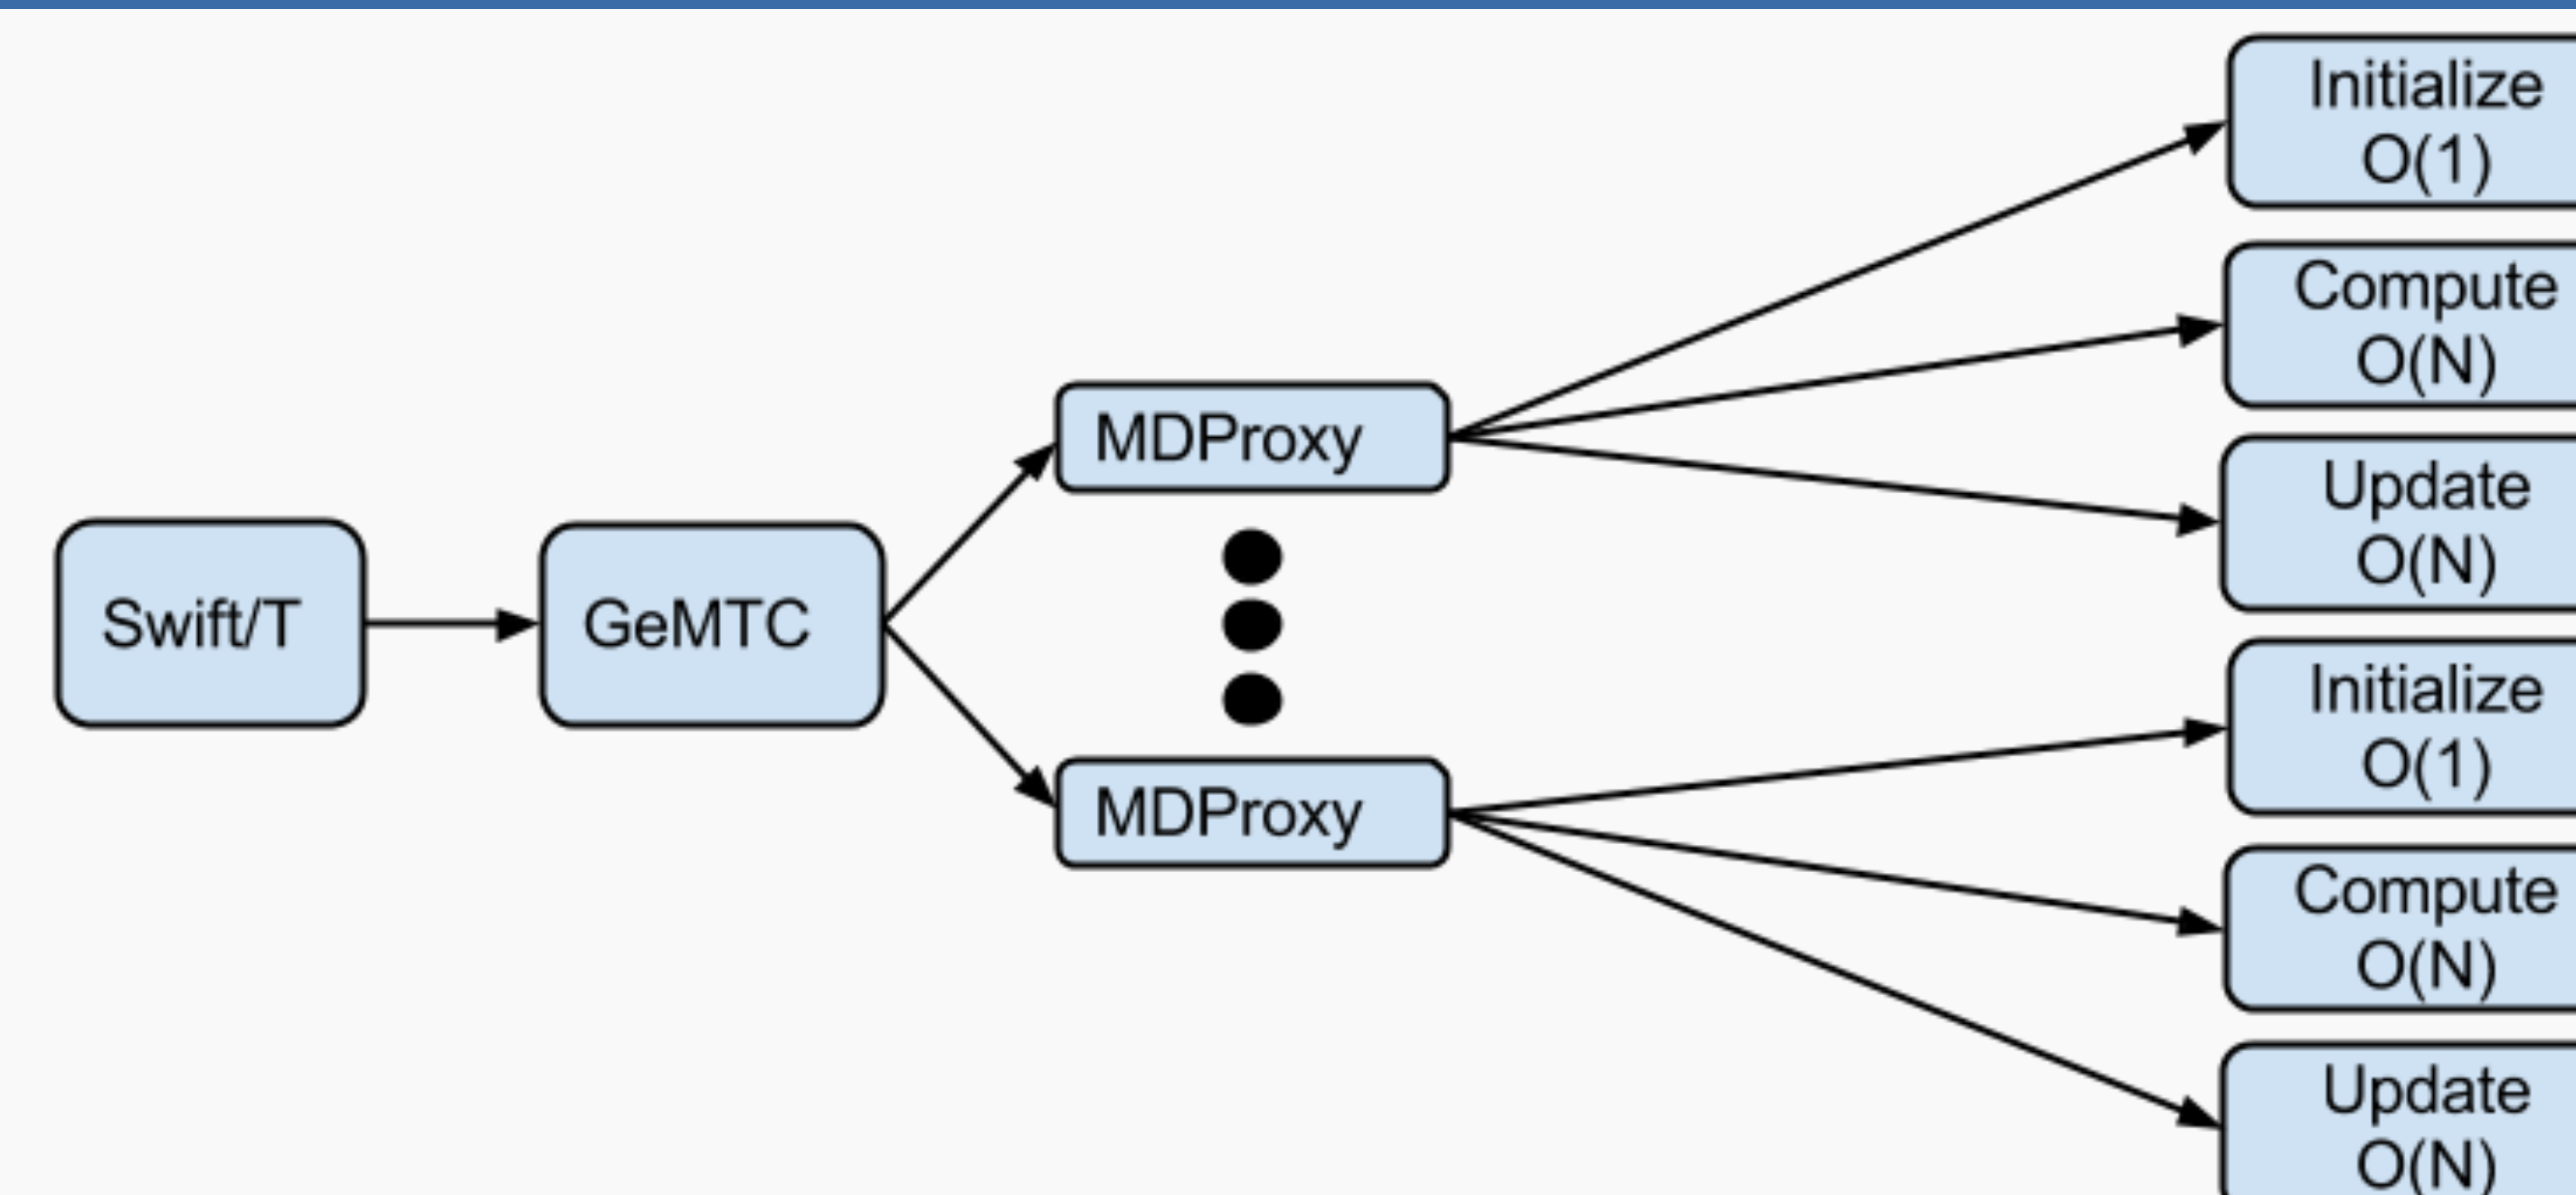
NVIDIA - <http://www.nvidia.com/object/nvidia-kepler.html>

GeMTC– <http://www.datasys.cs.iit.edu/projects/GeMTC>

## MDProxy Benchmarks



## MDProxy Architecture



## GeForce GTX 670



- 7 SMXs
- 84 Warps
- 1344 CUDA Cores
- 2GB DDR5

## Conclusions

- Evaluate Science Application
- MDPProxy highlights GeMTC potential
- GeMTC 10x faster than threaded CPU

## Future Work

- Improve MD algorithm
- Enable Multi-Node Performance
- Investigate optimal MD Task Size
- Compare performance against other GeMTC-enabled accelerators
- Develop high level abstractions for the Swift/T + GeMTC stack
- Expand library of GeMTC microkernels