

MatCloud: Accelerating Matlab Services via a GPU Cloud

Yongpeng Zhang, Xing Wu, Frank Mueller
North Carolina State University

Originally evolving from grid and utility computing, cloud computing has quickly become a ubiquitous phenomenon that provides novel methods of interaction with end-users. Existing commercial cloud platforms such as Amazon's Elastic Cloud (EC2), Google App Engine and Microsoft's Windows Azure (known as Platform as a Service (PaaS)) can greatly reduce the development cycle to build business applications. But their generalities come with the sacrifice of reduced flexibility. As a result, existing scientific cloud services built on those platforms are confined to embarrassingly parallel algorithms or are constrained to paradigms such as Map-Reduce.

The applicability of GPU, on the other hand, has diversified from pure graphic domain to general-purpose and high-performance computing in recent years. Its wide adoption is attributed to a significant performance margin over general-purpose microprocessor and a concise parallel programming abstraction. Currently, general-purpose microprocessors are dominant in supporting aforementioned commercial cloud platforms with no prior work on exposing accelerators such as GPUs in the near future.

MatCloud is a cloud infrastructure and service built from scratch aiming to benefit from the latest GPU technology (Fig. 1). It seeks to not only demonstrate the applicability of GPUs on cloud services, but also to boost the adoption of cloud computing by offering significant performance gains, as shown in our preliminary benchmark results in Fig. 2. The figure compares the execution time of several fundamental matrix operations running on MatCloud with a Matlab script on a contemporary desktop. The matrix size is 2000x2000 in all benchmarks.

We will present the infrastructure of MatCloud and describe our experience designing and implementing MatCloud. MatCloud is under active development and new research results will complement over our early experience reported here.

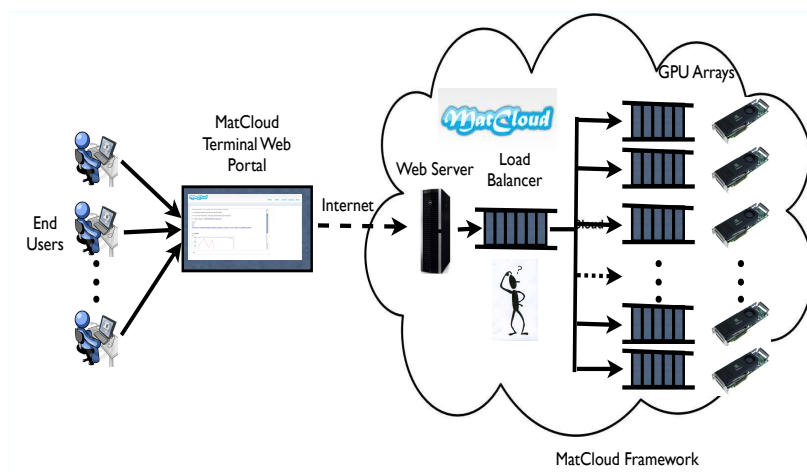


Fig. 1. MatCloud Overview

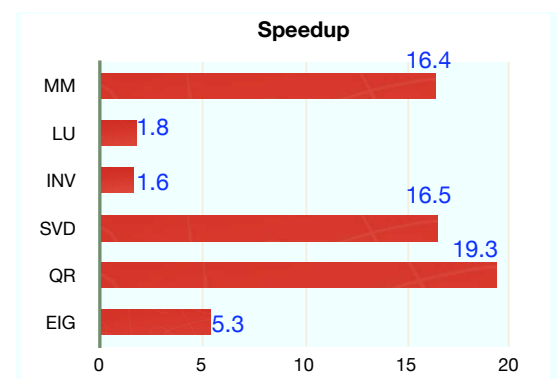


Fig. 2. MatCloud Vs. Matlab