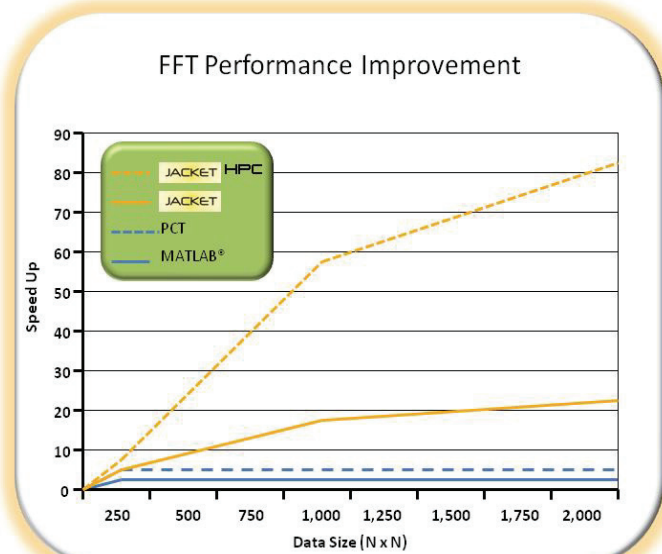
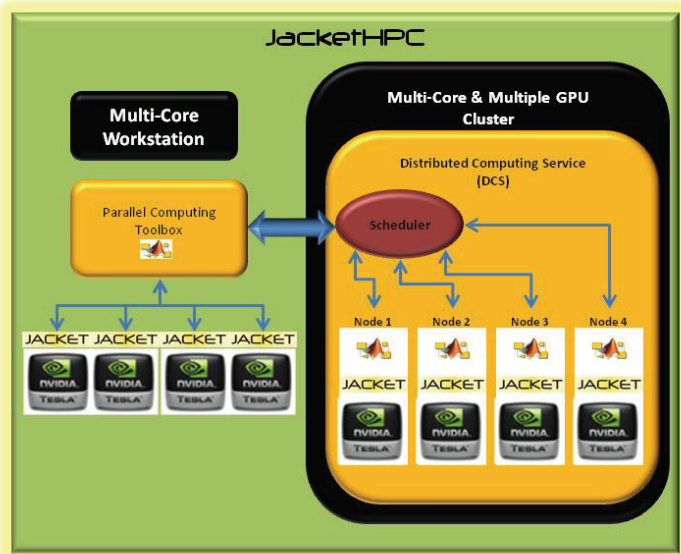


JacketHPC

Jacket is a software platform specifically designed for engineers, scientists, and analysts who need maximum application performance with minimal programming difficulty. *Jacket* is a runtime and language processing system. Through its GPU data types, *Jacket* automatically optimizes application code for GPU acceleration. During execution, the runtime system analyzes the program and balances computation between the CPU and GPU resources in a single computer or across a network. The *Jacket* platform may be used in new algorithms or for acceleration of existing applications. *Jacket* currently supports the MATLAB language as a frontend to the platform. The language processing system maps MATLAB code to GPU hardware. After initial analysis, *Jacket's* runtime system optimizes memory usage, data transfer, and computation. *Jacket* leverages enterprise-wide technical computing resources including notebooks, desktops, servers, clusters, and the Cloud by enhancing and extending the computational cycles accessible to domain experts.

JacketHPC delivers the ability to span computation across multiple GPUs, either on a local machine or over a network. *JacketHPC* allows for an unprecedented ability to transparently scale GPU and CPU computing resources simultaneously. Additional GPUs added to a host may now be instantly utilized without additional code modification. When a host is not capable of driving more GPUs, simply add GPUs to another host on the network and *JacketHPC* will take care of the rest. With the addition of *JacketHPC*, pre-existing CPU clusters may be upgraded through the installation of GPUs, significantly increasing the cluster's computational capability without investing in new development for specialized GPU code.

JacketHPC is built atop Mathworks' Parallel Computing Toolbox (PCT) and Distributing Computing Server (DCS). PCT and DCS product licenses are required for executing *JacketHPC* on network based HPC resources. With the addition of parallel constructs such as *parfor*, *spmd*, or *co-distributed arrays*, pre-existing code may be dispatched across all GPUs and CPUs in a cluster or a Cloud service. In many cases, little to no code revision is required to take advantage of this parallel computing capability.



Optional capabilities are also available with *JacketHPC*, including a Developer SDK and a Compiler that enables the creation of executables for deployment of applications to a large user base. The Developer SDK makes integration of custom CUDA code into *Jacket's* runtime very easy. With a few *Jacket*-specific functions, CUDA code will benefit from *Jacket's* automated optimizations. When *Jacket* applications have completed the development, test, and optimization stages and are ready for deployment, the Compiler option allows users to create standalone, license-free executables for deployment to larger user bases.

AccelerEyes, based in Atlanta, Georgia provides a variety of support and services for its customers. In addition to traditional product documentation, an active User Forum and many applications examples are available on the company website. Optional software maintenance and phone support is also available. Contact AccelerEyes or an authorized partner for additional information on *Jacket*, optional products, or other company support and services.

