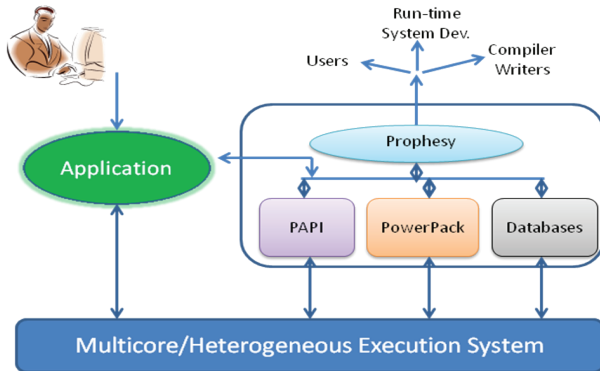


MuMI

MuMI, or Multicore Application Modeling Infrastructure, is a project that is developing a framework to facilitate systematic measurement, modeling, and prediction of performance, power consumption, and performance-power tradeoffs for applications running on multicore systems. MuMI will offer an integrated and extensible performance modeling and prediction framework for use by application developers, system designers, and scientific application users, helping them to make key choices for configuring and selecting appropriate multicore and hybrid systems and to evaluate and optimize power/performance on these systems.




MuMI combines UTK's PAPI hardware performance monitoring capabilities with Texas A&M's Prophesy performance modeling interface and Virginia Tech's Power-Pack power-performance measurement and analysis system.

POWER MEASUREMENT


PAPI has been integrated with Power-Pack and Prophesy, and performance and power consumption data have been collected for a range of benchmarks and applications running on multicore systems.

In the graphs below, the LAPACK program consumes system power of about 250 watts and CPU power of 152 watts. The TBLAS program consumes system power of 228 watts and CPU power of 145 watts. The experimental result demonstrates that the dynamic scheduling program takes less execution time than LAPACK by 21.8%, and uses less power. This is because the dynamic scheduling method has fewer synchronization points and better data locality due to the tile data layout.

LAPACK Cholesky Factorization



TBLAS Cholesky Factorization



WITH SUPPORT FROM



SPONSORED BY



FIND OUT MORE AT <http://www.mumi-tool.org/>