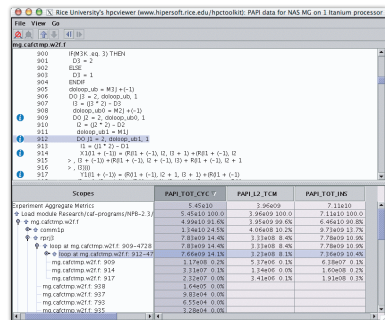


PAPI

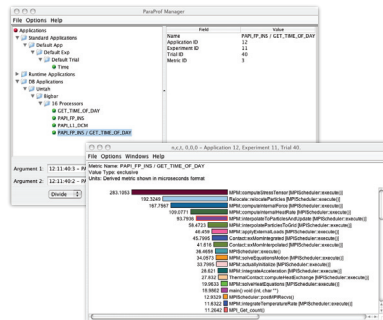
R&D100
2001 WINNER

PERFORMANCE APPLICATION PROGRAMMING INTERFACE

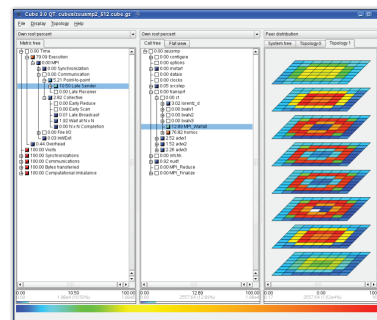
PERFORMANCE ANALYSIS TOOLS



HPCView



TAU



Scalasca

MORE TOOLS THAT USE PAPI

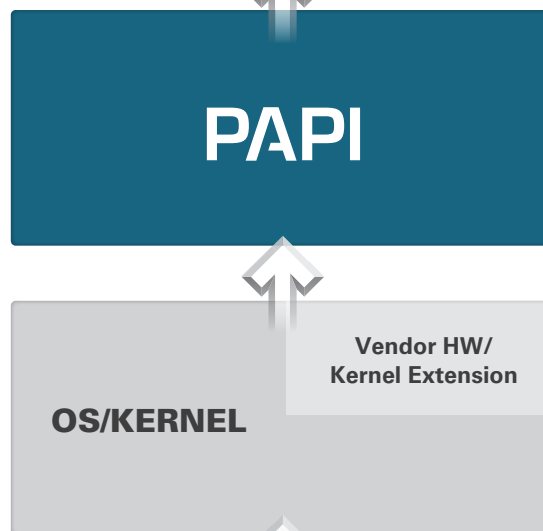
- HPCToolkit
- IPM
- Open|SpeedShop
- PerfSuite
- SCALEA
- Titanium
- Vampir

STANDARD FEATURES


- Standardized Performance Metrics
- Easy Access to Platform-Specific Metrics
- Multiplexed Event Measurement
- Dispatch on Overflow
- Overflow & Profiling on Multiple Simultaneous Events
- Bindings for C, Fortran, Matlab and Java
- User Definable Metrics after Platform-Specific Metrics

SUPPORTED ARCHITECTURES

- AMD Families 0fh, 10h, 14h and 15h
- CRAY XT and XE Series
- IBM POWER AIX, Linux
- IBM BG/L, BG/P
- Intel Nehalem, Westmere, Sandy Bridge
- SGI Altix
- ARM Cortex A8, A9




NETWORK COMMUNICATIONS



CURRENT CPU



TEMPERATURE/SYSTEMS DIAGNOSTICS



GPU

INNOVATIVE
COMPUTING LABORATORY
THE UNIVERSITY OF TENNESSEE

FIND OUT MORE AT <http://icl.eecs.utk.edu/papi/>

WITH SUPPORT FROM



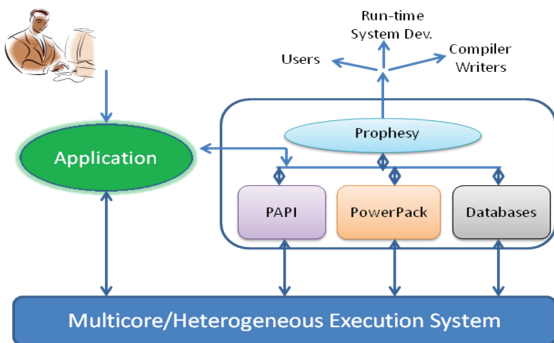
SPONSORED BY



MUMMI

MULTIPLE METRICS MODELING INFRASTRUCTURE

MuMMI, or Multiple Metrics 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 and heterogeneous systems. MuMMI offers 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.



MuMMI 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.

USER-DEFINED EVENTS/PERFORMANCE MODELING

Facilities are being developed within the PAPI project to enable the user to define modeling metrics at a high level and have these metrics mapped to underlying hardware events and characteristics. The idea of user-defined events is to allow users to define their own metrics and to have these metrics mapped to events on a platform without the need to re-install PAPI. This capability allows the user to define a metric by means of an arithmetic expression that can include measured hardware counter values as well as system constants. An example of such an expression is the following which represents the cycles per instruction due to data memory accesses, assuming two levels of cache:

$$(PAPI_L1_DCA * Data_L1_lat + PAPI_L2_DCA * L2_lat + PAPI_L2_DCM * Mem_lat) / PAPI_TOT_INS$$

Using PAPI user-defined events, users can work with any performance model that is based on hardware counters using data collected by any tool that supports PAPI.

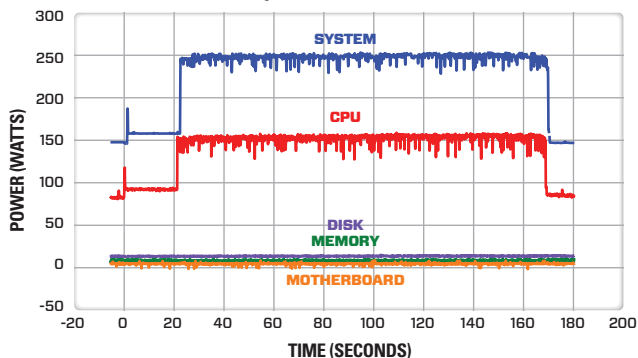
POWER MEASUREMENT

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.

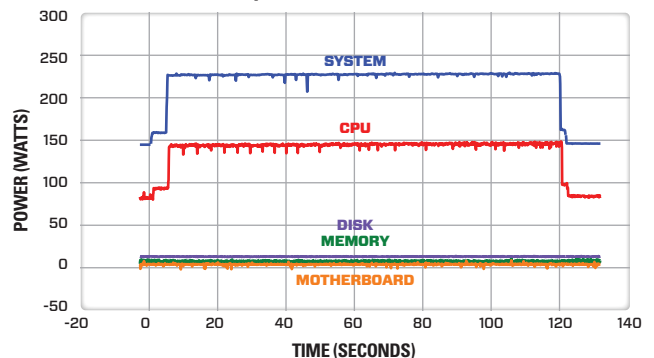
In the graphs below, the LAPACK program consumes system power of about 250 watts and CPU power of 152 watts. The TBLAS program consumes system power of 228 watts and CPU power of 145 watts. The experimental result

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



INNOVATIVE
COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE

FIND OUT MORE AT <http://www.mummi.org/>

A COLLABORATION OF



VirginiaTech

WITH SUPPORT FROM



SPONSORED BY

