Multiple Metrics Modeling Infrastructure (MuMMI)

Presented by

Shirley Moore
Dan Terpstra

University of Tennessee
Oak Ridge National Laboratory

MuMMI Framework

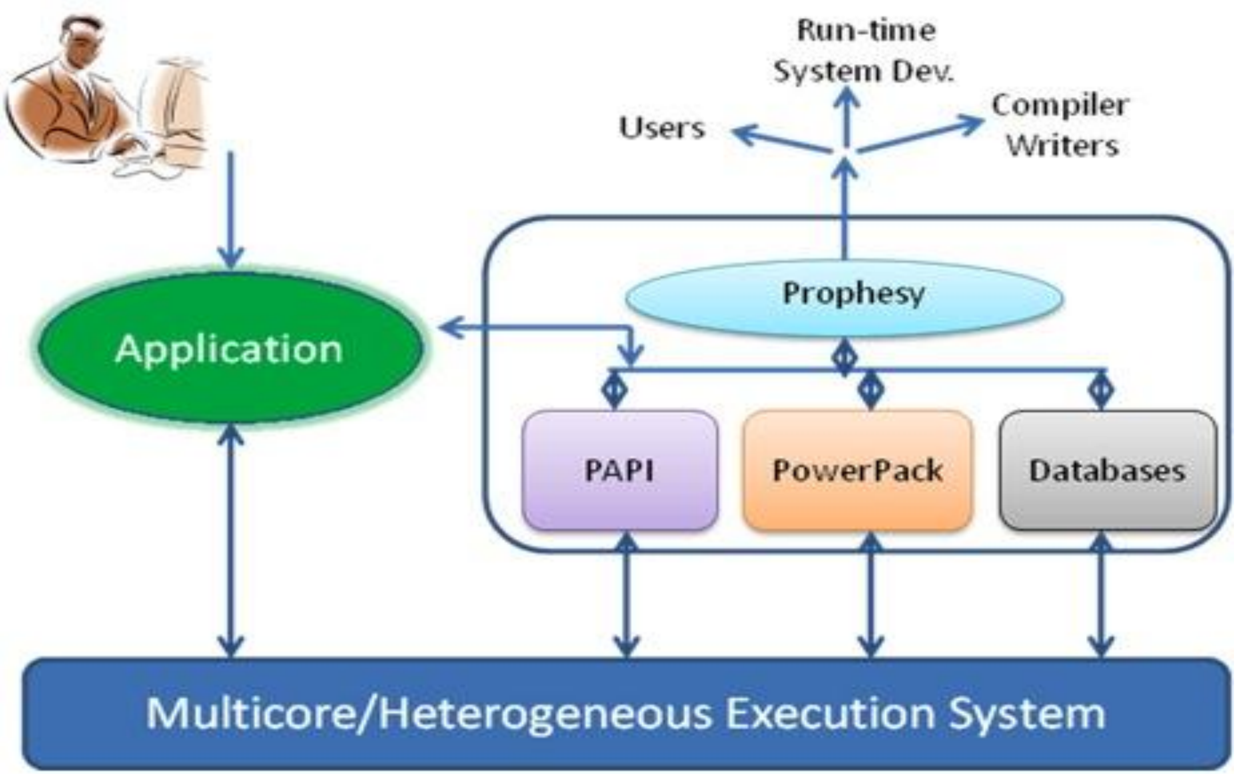

Multiple Metrics Modeling Infrastructure (MuMMI)

SystemG

Configuration of SystemG	
Mac Pro Model Number	MA970LL/A
Total Cores	2,592
Total Nodes	324
Cores/Socket	4
Cores/Node	8
CPU Type	Intel Xeon 2.8Ghz Quad-Core
Memory/Node	8GB
L1 Inst/D-Cache per core	32-kB/32-kB
L2 Cache/Chip	12MB
Interconnect	QDR Infiniband 40Gb/s

- Largest power-aware compute system in the world
- Over 30 power and thermal sensors per node
- http://scape.cs.vt.edu/

Power-Aware Predictive Models of Hybrid (MPI/OpenMP) Scientific Applications on Multicore Systems

Charles Lively III*, Xingfu Wu*, Valerie Taylor*, Shirley Moore+, Hung-Ching Chang^, Chun-Yi Su^, and Kirk Cameron^

***Department of Computer Science & Engineering, Texas A&M University**

+Electrical Engineering and Computer Science, University of Tennessee-Knoxville

^Department of Computer Science, Virginia Tech

General Methodology

- Explore which application characteristics (via performance counters) affect power consumption of system, CPU, and memory
- Develop accurate models based on hardware counters for predicting power consumption of system components
- Develop different models for each application class (Previous work used same set of performance counters across all applications).
- Validate predictions using actual power measurements

Modeling Methodology

- Training Set: 5 training execution configurations
  - 1x1, 1x2, 1x3, 1x8, and 2x8
- 16 larger execution configurations are predicted.
  - 1x4, 1x5,…3x8, 4x8, 5x8, …..16x8
- 40 performance counter events are captured.
- Performance counter events are normalized per cycle.
- Performance-Tuned Supervised Principal Component Analysis Method is utilized to select combination of performance counters for each application.

Applications

- NAS Multizone Benchmark Suite
  - written in Fortran
  - Uses MPI and OpenMP for communication
  - Block Tri-diagonal algorithm (BT-MZ)
    - represents realistic performance case for exploring discretization meshes in parallel computing
  - Scalar Penta-diagonal algorithm (SP-MZ)
    - representative of a balanced workload
  - Lower-Upper symmetric Gauss-Seidel algorithm (LU-MZ)
    - coarse-grain parallelism of LU-MZ is limited to 16 MPI processes

- Large-Scale Scientific Application
  - Gyrokinetic Toroidal code (GTC)
    - 3D particle- in-cell application
    - Flagship SciDAC fusion microturbulence code
    - written in Fortran90
    - Uses MPI and OpenMP for communication

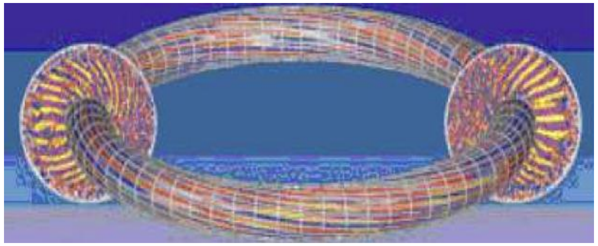

Application-specific Modeling

Multivariate regression coefficients

	Time		System Power		CPU Power		Memory Power	
BT-MZ	Cache_FLD	-1.611	PAPI_L2_TCH	-1.6769	PAPI_L1_TCM	3.5432	PAPI_L1_TCA	0.0763
	PAPI_TOT_INS	0.0967	PAPI_L2_TCA	1.5967	PAPI_L2_TCH	-3.9389	PAPI_L1_DCM	4.0496
	PAPI_L2_TCH	0.2992	PAPI_RES_STL	0.0803	PAPI_RES_STL	0.3967	PAPI_L2_TCH	-1.9443
	PAPI_L2_TCA	1.2152					PAPI_L2_TCA	2.1806
SP-MZ	PAPI_TOT_INS	0.1818	PAPI_L1_ICA	0.355	LD_ST_stall	0.1917	Cache_FLD	0.4563
	PAPI_L1_TCA	0.0744	PAPI_L2_TCH	-1.3452	PAPI_L1_TCM	1.5008	LD_ST_stall	0.0192
	PAPI_L2_TCH	-1.2834	PAPI_L1_TCM	0.9911	PAPI_L2_TCH	-1.6914	PAPI_L2_TCH	-3.5895
	PAPI_L1_TCM	1.1761					PAPI_L2_TCA	3.1151
LU-MZ	Cache_FLD	-0.0006	LD_ST_stall	0.0166	LD_ST_stall	0.0869	PAPI_L1_TCA	0.27923
	PAPI_TOT_INS	0.0011	PAPI_L2_TCH	-0.9886	PAPI_L2_TCH	-8.0003	PAPI_L2_TCH	-3.9574
	PAPI_TLB_DM	3.9085	PAPI_L2_TCA	1.0411	PAPI_L2_TCA	7.9137	PAPI_RES_STL	-0.29141
	PAPI_L2_TCH	-0.0591	PAPI_RES_STL	0.025				
GTC	PAPI_TOT_INS	0.0006	PAPI_RES_STL	1.5689	PAPI_RES_STL	0.9261	PAPI_TOT_IN	0.169617
	PAPI_L2_TCH	-1.8976	PAPI_L2_TCH	-3.2505	PAPI_TOT_IN	0.2663	PAPI_L2_TCH	-2.881
	PAPI_L2_TCA	1.9351	PAPI_L1_TCA	1.6916	PAPI_L1_TCA	0.0816	PAPI_L2_ICM	2.7119
	PAPI_BR_INS	-0.0381			PAPI_L2_TCH	-1.2640		

Overall Prediction Accuracy

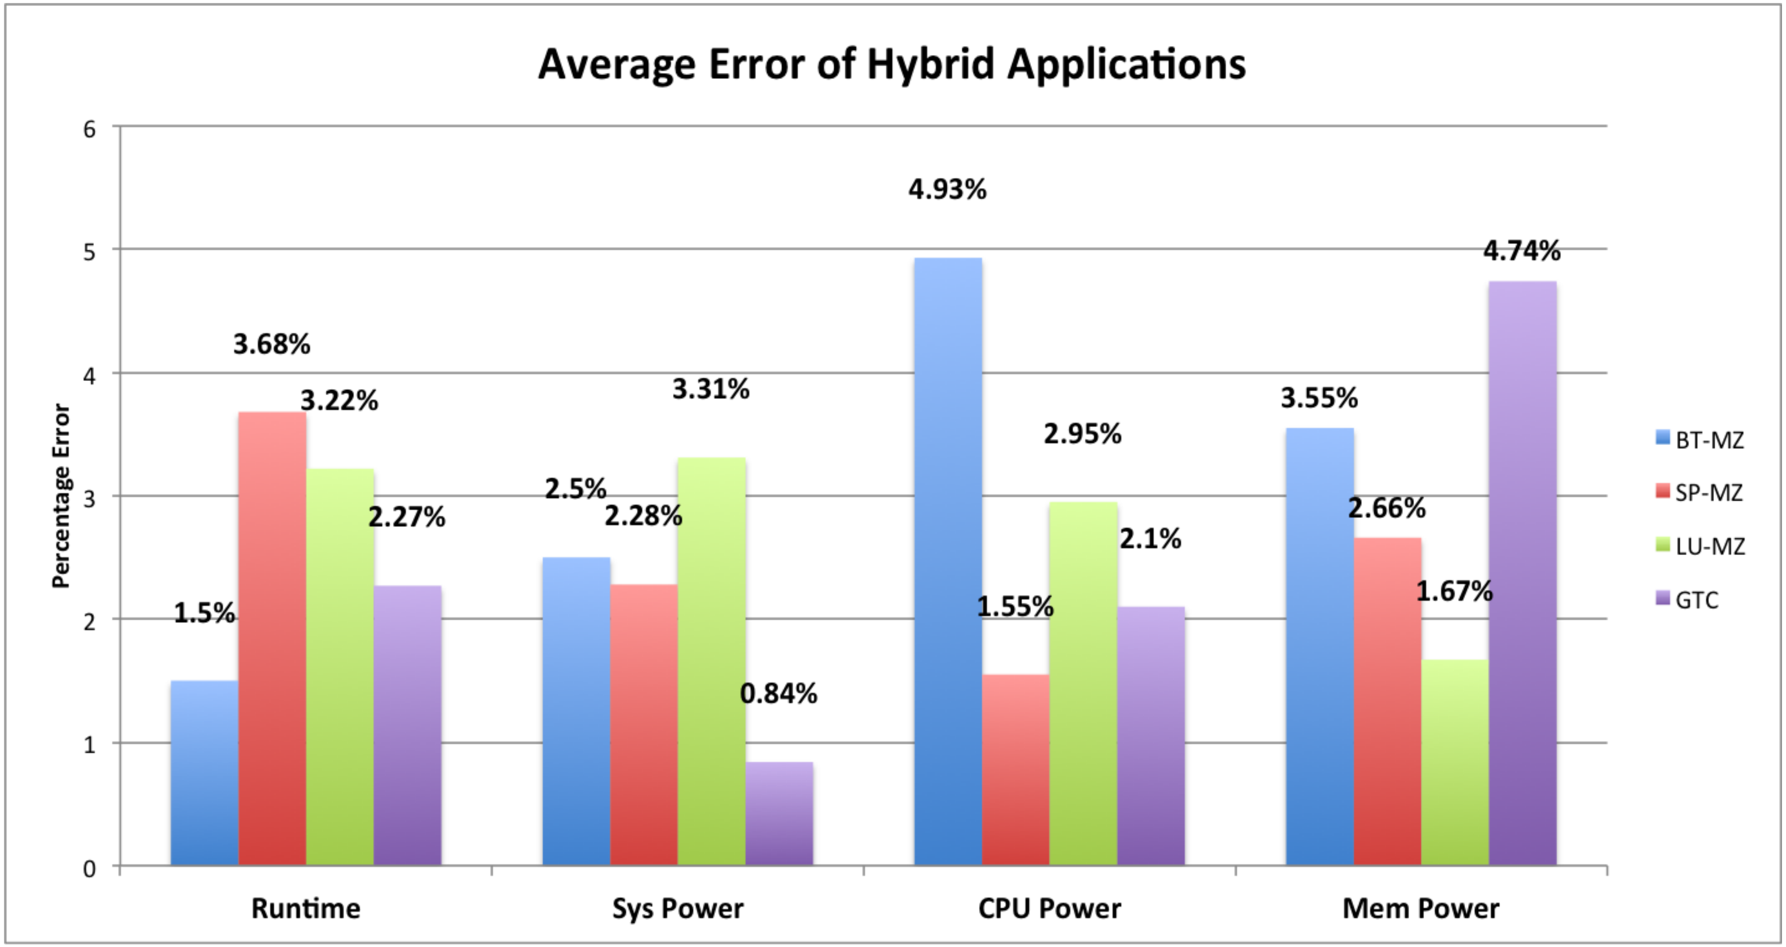

Conclusions

- Predictive performance models for hybrid MPI+OpenMP scientific applications.
  - Execution time
  - System power consumption
  - CPU power consumption
  - Memory power consumption
- 95+% accuracy across four hybrid (MPI+OpenMP) scientific applications
- Future work
  - Explore use of microbenchmarks and application classes to derive application-centric models
  - Finer-granularity analysis of large-scale hybrid scientific applications

Contact

For more information

www.mumi-tool.org

This work is supported by NSF grants CNS-0911023, CNS-0910899, CNS-0910784, CNS-0905187.