A UT/ORNL PARTNERSHIP
NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES
NICS

Power-aware Computing on GPGPUs

Kiran Kumar Kasichayanula – 2, Haihang You – 1, Shirley Moore – 2, Stanimire Tomov – 2, Heike Jagode – 2, Matt Johnson – 2

1. NICS Scientific Computing Group, 2. ICL, EECS Department, University of Tennessee

ICL UT

Power-aware Computing on GPGPUs

- Tracing of power measurement exactly when the kernel is running
- NVML (NVIDIA Management Library) library to measure the real-time power consumption of several fundamental BLAS libraries and LAPACK routines.
- Activity-based model (AMG) to estimate activity factors and power for micro-architectures on GPGPUs

NICS

Methodology

- NVML supports function calls such as nvmlDeviceGetPowerUsage to measure power in watts.
- NVML supports function calls such as nvmlDeviceGetTemperature to measure temperature in Celsius.
- Use of pthreads where we run MAGMA kernels on one thread and power measurement on another.
- NVML power usage function calls fully supported on C2075 and partially supported on c2050 and other GPUs.
- We show the trace of power consumption exactly when MAGMA kernels are running using the latest feature provided by CUDA 4.0, which allows access to the same GPU by different threads.

NICS

Traces of Real-time Power Consumption

- Power trace of MAGMA
  - **http://icl.cs.utk.edu/magma/**
- Implementation of LU factorization (magma_dgetrf) for solving a dense linear system of equations
- MAGMA upper Hessenberg reduction (magma_dgehrd) for solving the general eigenvalue problem are shown.

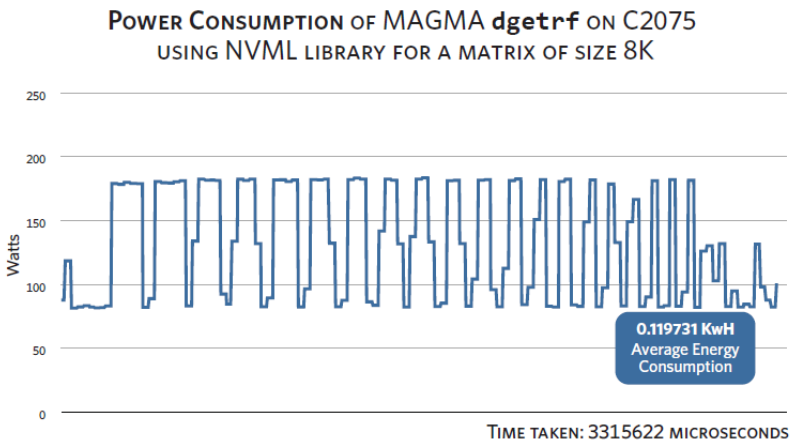

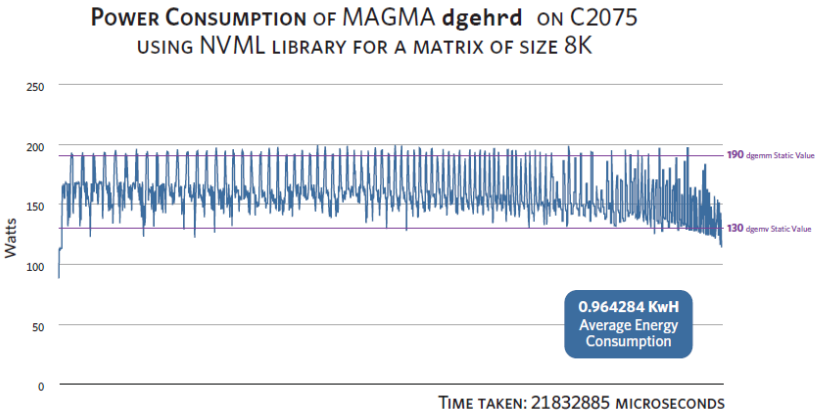

NICS

HPC @ 1/10th the cost & 1/20th the energy

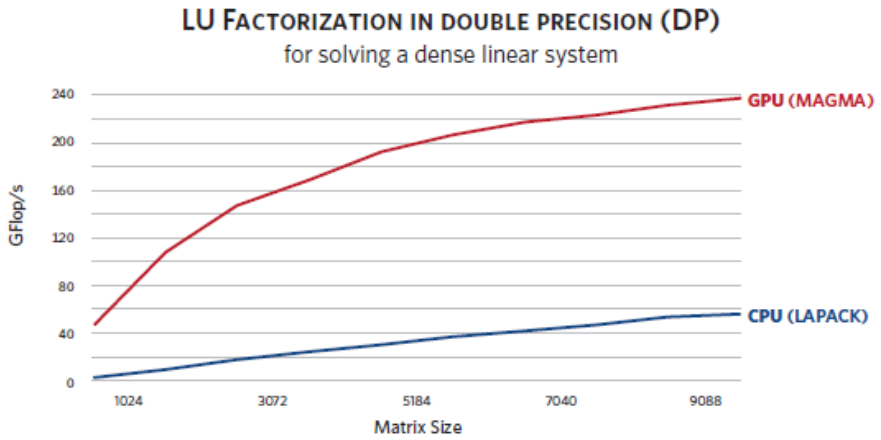

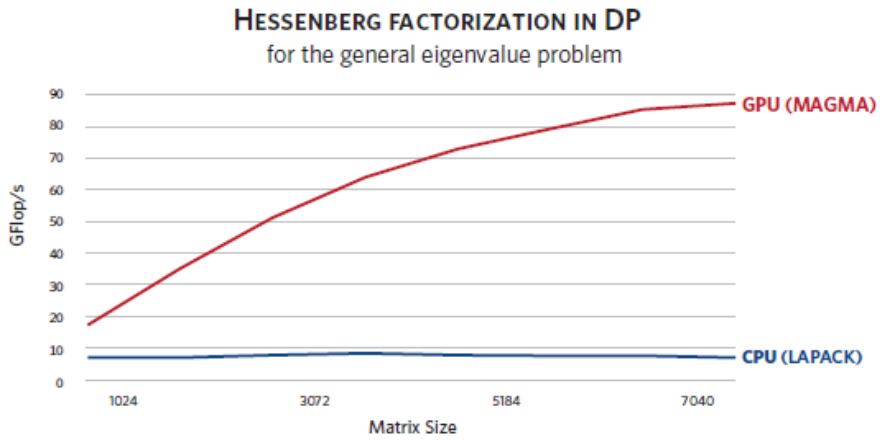

GPU	Fermi C2075 [448 CUDA Cores @ 1.15 GHz]	
DP Peak	515 + 40 GFlop/s	
System Cost	~$3,000	
Power	~220 w	
CPU	AMD ISTANBUL [8 sockets x 6 cores (48 cores) @ 2.8 GHz]	
DP Peak	538 GFlop/s	
System Cost	~$30,000	
Power*	~1,022 w	

* Computation consumed power rate (total system rate minus idle rate), measured with KILL A WATT PS, Model P430

- MAGMA's LU factorization is almost entirely based on GEMM which is a level 3 BLAS routine.
- MAGMA's HESSENBERG factorization has 80% FP ops from DGEMM which is a level 3 BLAS routine and 20% FP ops from DGEMV which is a level 2 BLAS routine.

NICS

Temperature trace for MAGMA dgehrd

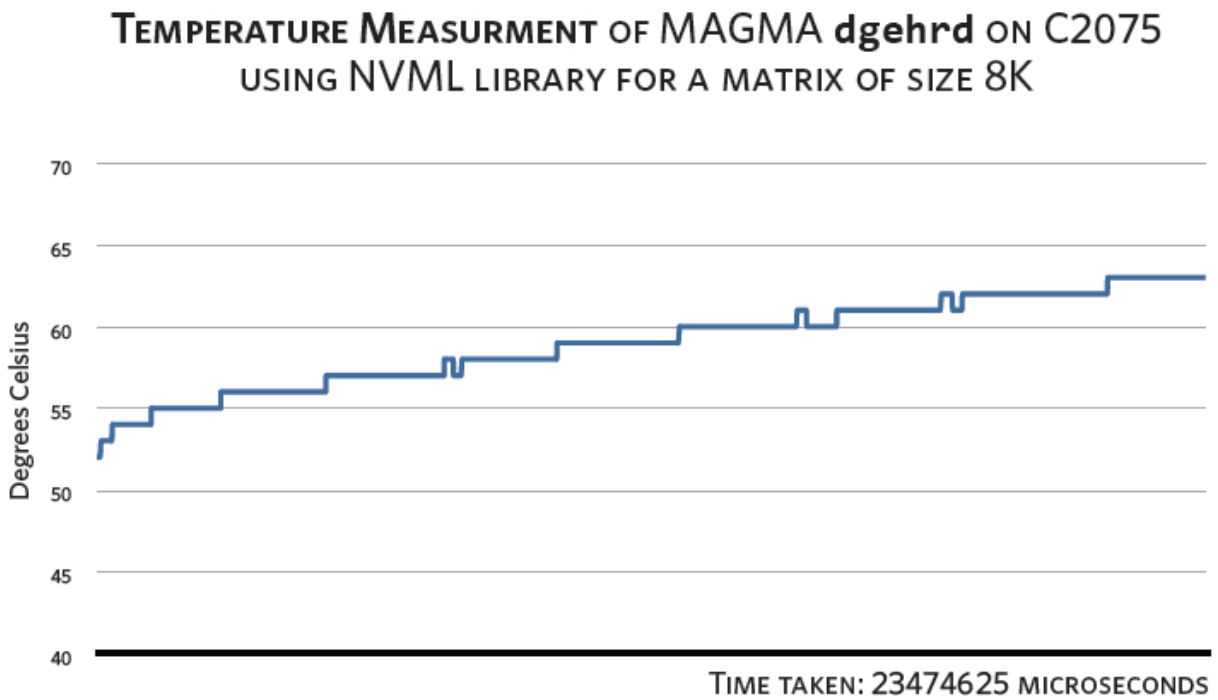

- The graph shows the real-time temperature while running MAGMA dgehrd.
- We would like to correlate temperature with power consumption and failures.

NICS

Future Work

- Development of activity-based models for performance and power consumption prediction
- Deeper understanding of performance and power tradeoffs
- Use of libraries like PAPI, TAU, CUPTI for per component analysis
- Identification of relevant hardware counter events and addition to power and performance prediction model.
- Energy saving parameters are being identified and added to the MAGMA library.

NICS

Acknowledgments

- Work at ICL/UTK was partially supported by NSF grants 0910899 and by NVIDIA.

Contact

Haihang You
Computational Scientist
hyou@utk.edu

Kiran Kumar Kasichayanula
Research Assistant
kkasicha@utk.edu

Shirley Moore
Research Assoc. Prof.
shirley@eecs.utk.edu