
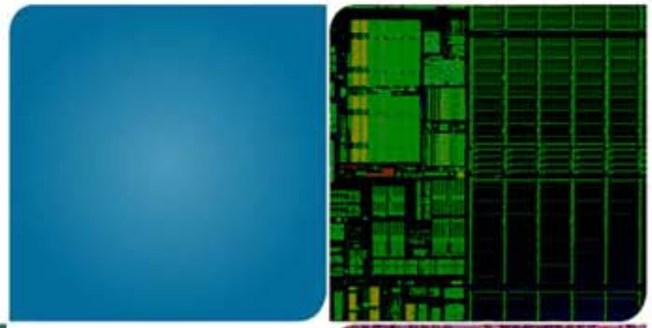

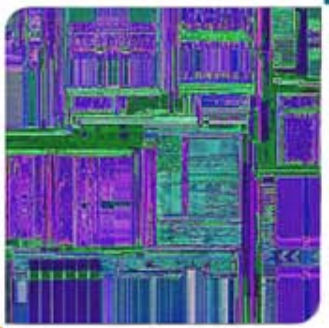
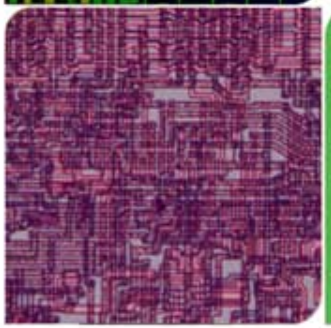


Power Measurements for Computer and Computer Architecture Research

Hyesoon Kim

Georgia Tech

comparch

cercs

Power Measurements

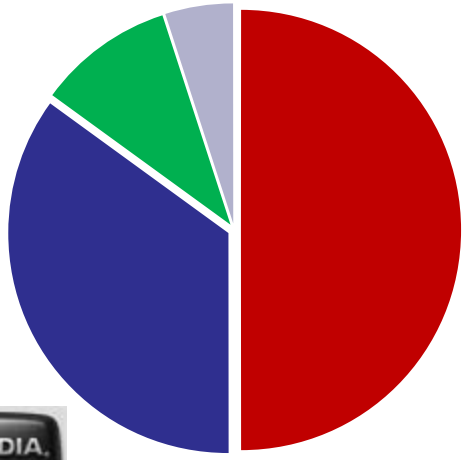

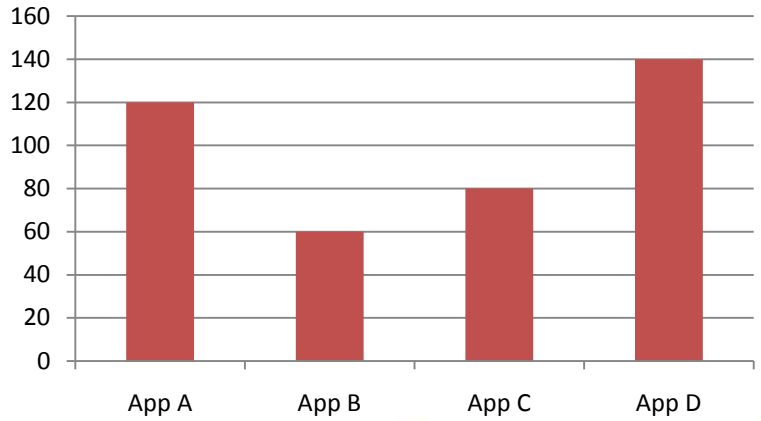

Measurements

- Power measurements
  - System Input power
  - Module power
  - Board level
- Temperature measurements

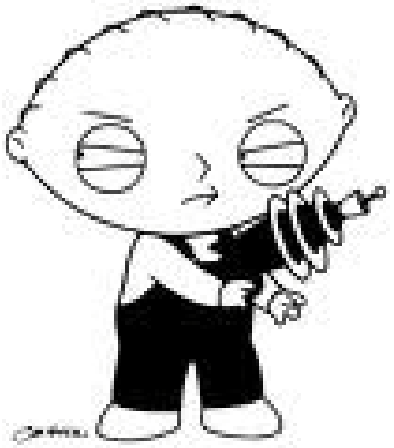

System Power Measurement

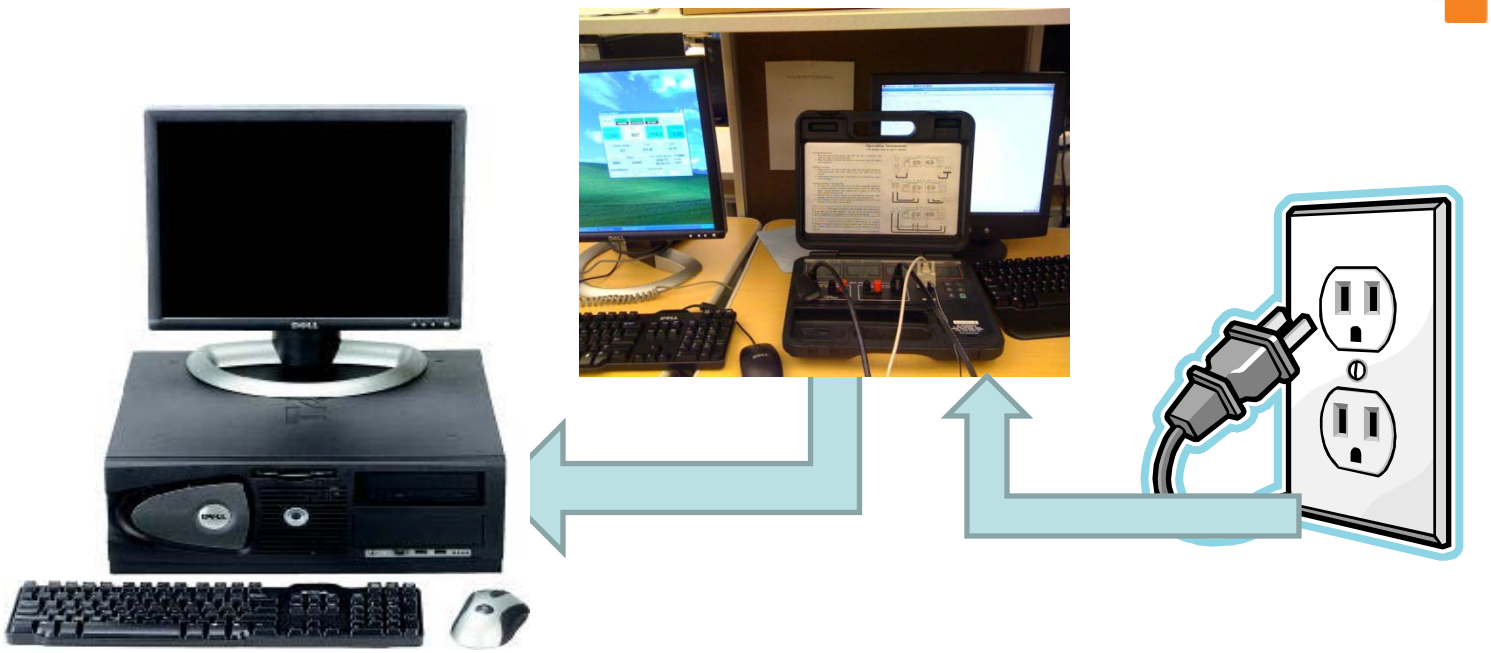

Power = Dynamic power + static power

Dynamic power $\propto$ dynamic activity

Static power $\propto$ Area, Temperature

Module Power: Memory

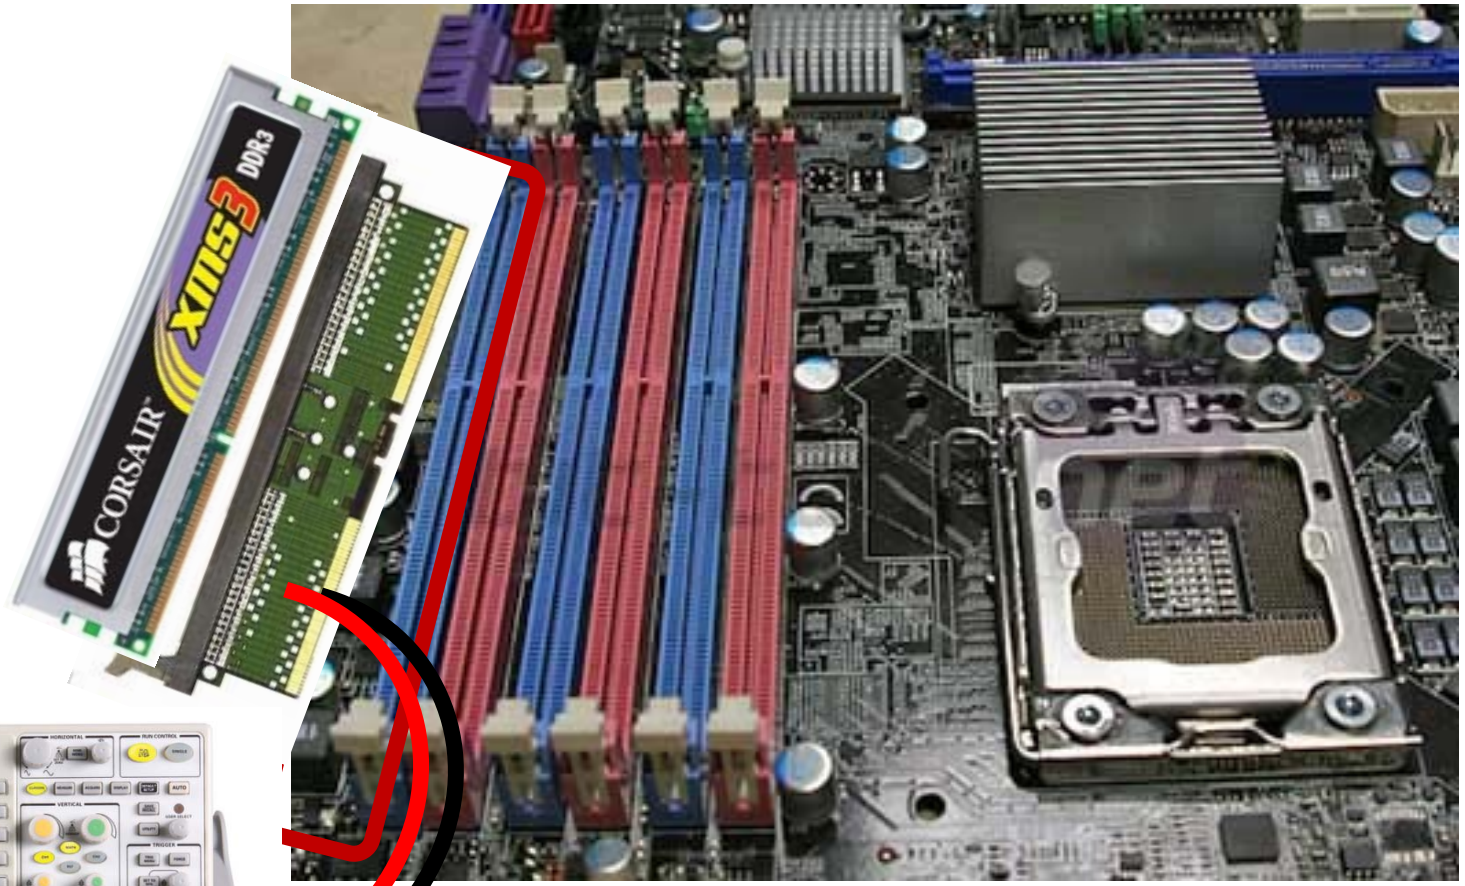

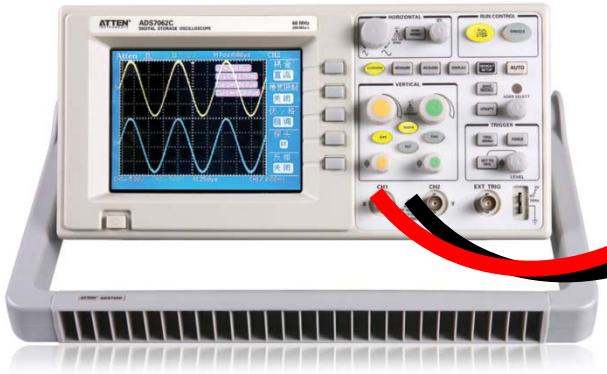

Module Power: GPU

Module Power: CPU

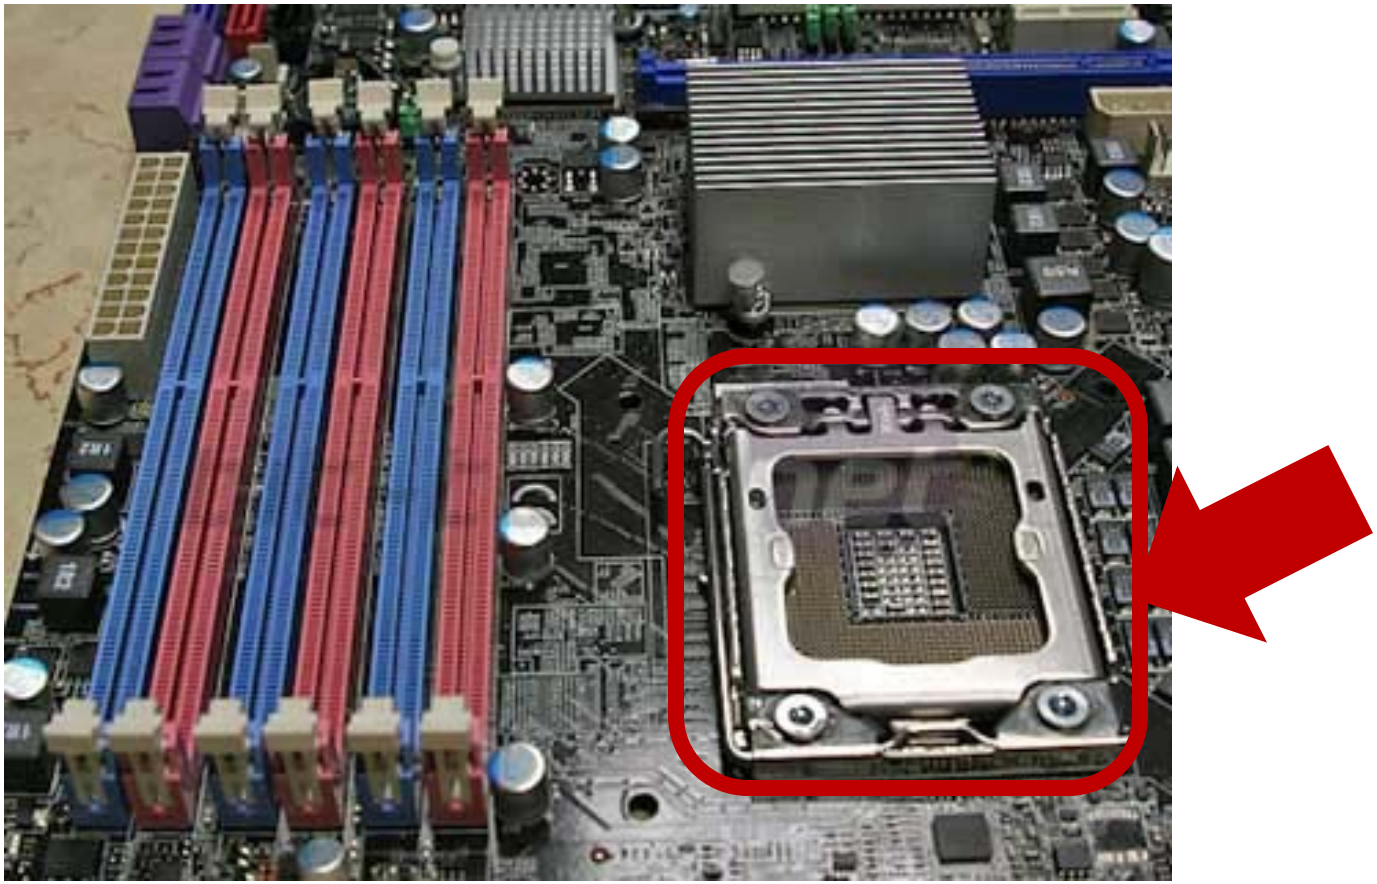

Board Level Power Measurements

CASE STUDY: GPU POWER MODEL [ISCA'10]

Overview of GPU Architecture

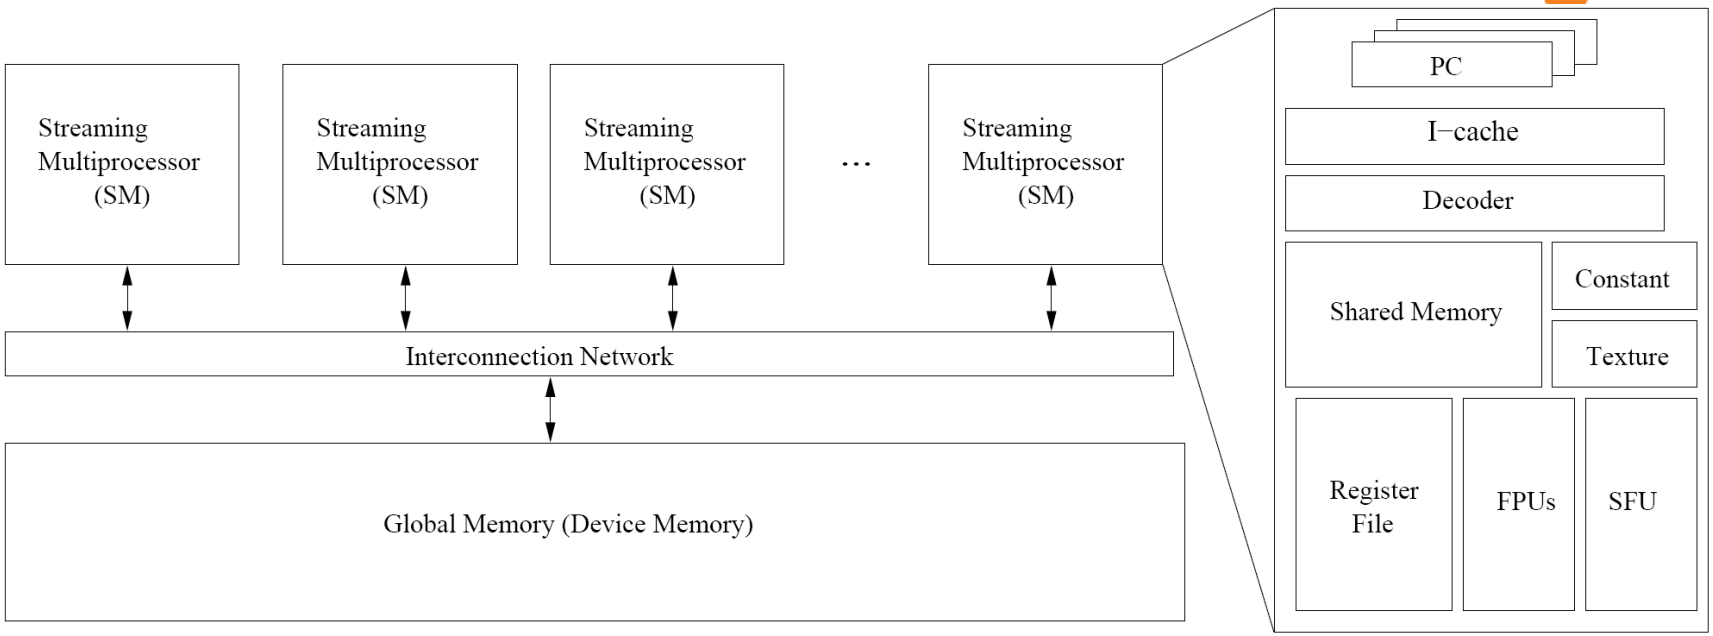

- Software-managed cache
- SIMD Execution Unit inside SM

Power Model

$$Total\ System\ Power = Runtime\_power + IdlePower = \sum_{i=0}^{n} RP\_Component_i + IdlePower$$
$$= RP\_SMs + RP\_Memory + IdlePower$$

- Power contributions: SMs, Memory and Idle-power

$$RP\_SMs = Num\_SMs \times \sum_{i=0}^{n} SM\_Component_i$$

$$\sum_{i=0} SM\_Component_i = RP\_Reg + RP\_Fp + RP\_Int + RP\_Sfu + RP\_Alu + RP\_Shared$$
$$+ RP\_Const\_SM + RP\_TextureCache + RP\_ConstCache$$

- Streaming multiprocessor (SM) decomposition

$$RP\_Memory = \sum_{i=0}^{n} Memory\_component_i$$
$$= RP\_GlobalMem$$

- Memory: Global memory system

Hong, Kim, ISCA'10

Training Benchmarks To Find Parameters

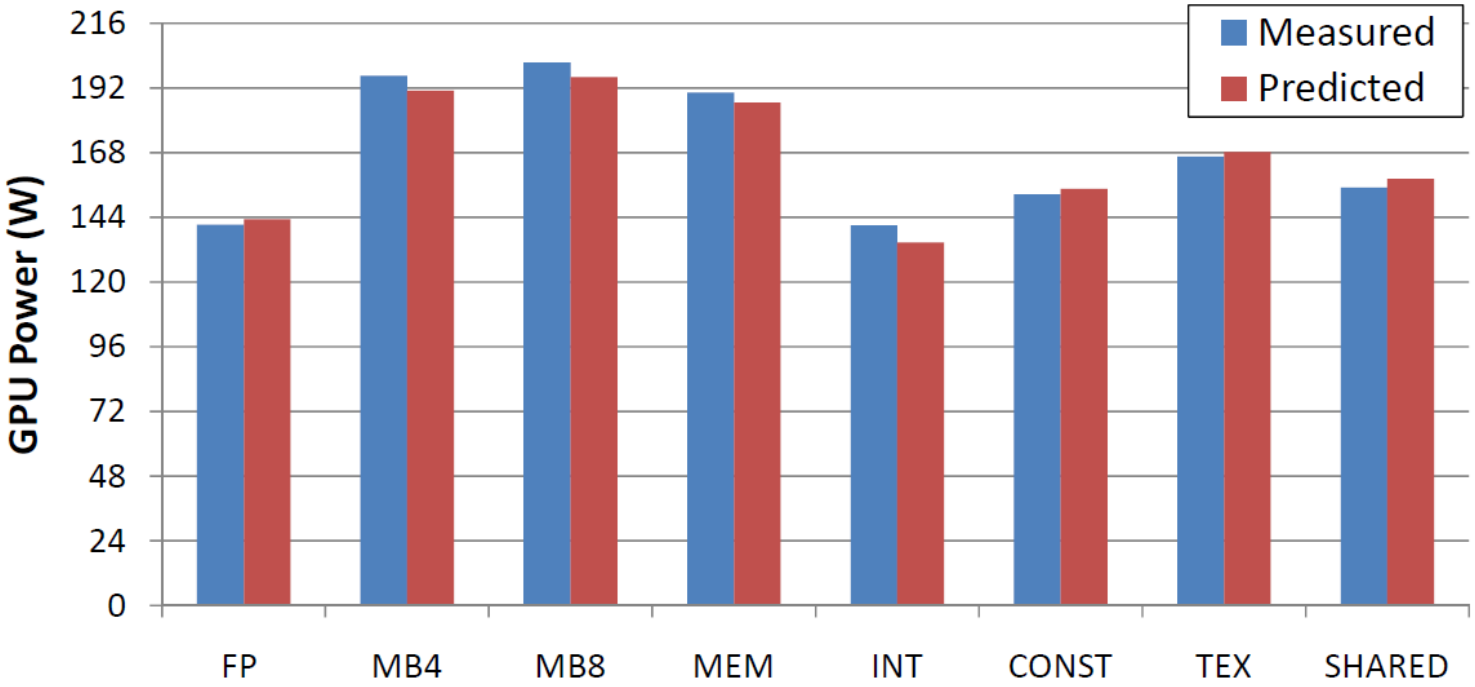

- Training benchmarks stress particular units of processors, memory

Hong, Kim, ISCA'10

Power Model

Table 2. Power parameters

Names	MaxPower	OnChip	SpecializedLinear	Term Location
FP	0.2	Yes	Yes	*RP_SM*
REG	0.3	Yes	Yes	*RP_SM*
ALU	0.2	Yes	No	*RP_SM*
SFU	0.5	Yes	No	*RP_SM*
INT	1	Yes	Yes	*RP_SM*
FDS (Fetch/Decode/Schedule)	0.5	Yes	Yes	*RP_SM*
Shared memory	1	Yes	No	*RP_SM*
Texture cache	0.9	Yes	Yes	*RP_SM*
Constant cache	0.4	Yes	Yes	*RP_SM*
Global memory	52	No	Yes	*RP_Memory*
Const_SM	0.213	Yes	No	*RP_SM*

- Power parameters used
- Units inside SM are per SM granularity
- Initial increase from idle to low accessrate causes large power consumption for some units

Predictions

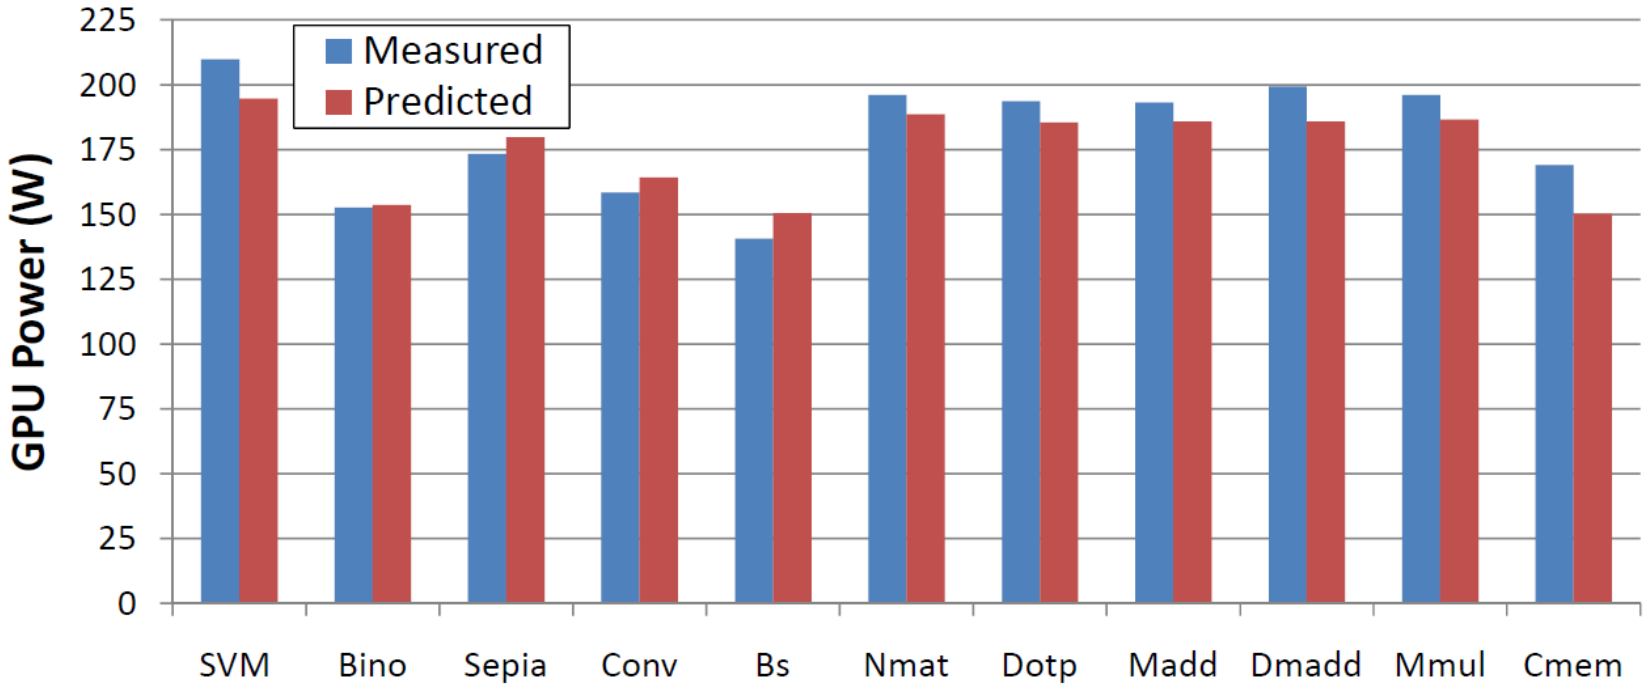

Hong, Kim, ISCA'10

Power Consumption Breakdown

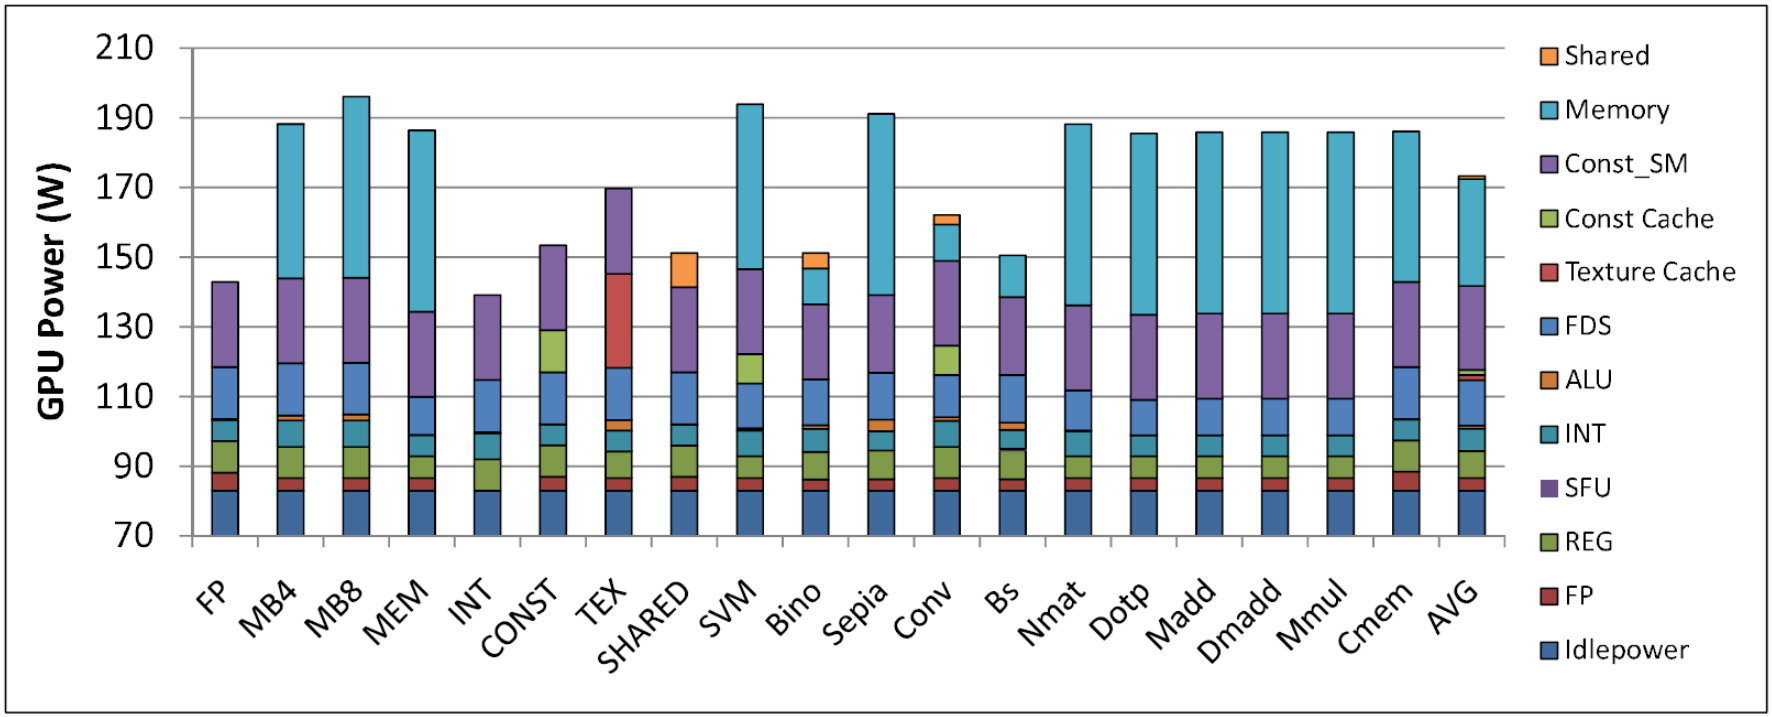

Hong, Kim, ISCA'10

Usage 1: Find Optimal Number of Cores

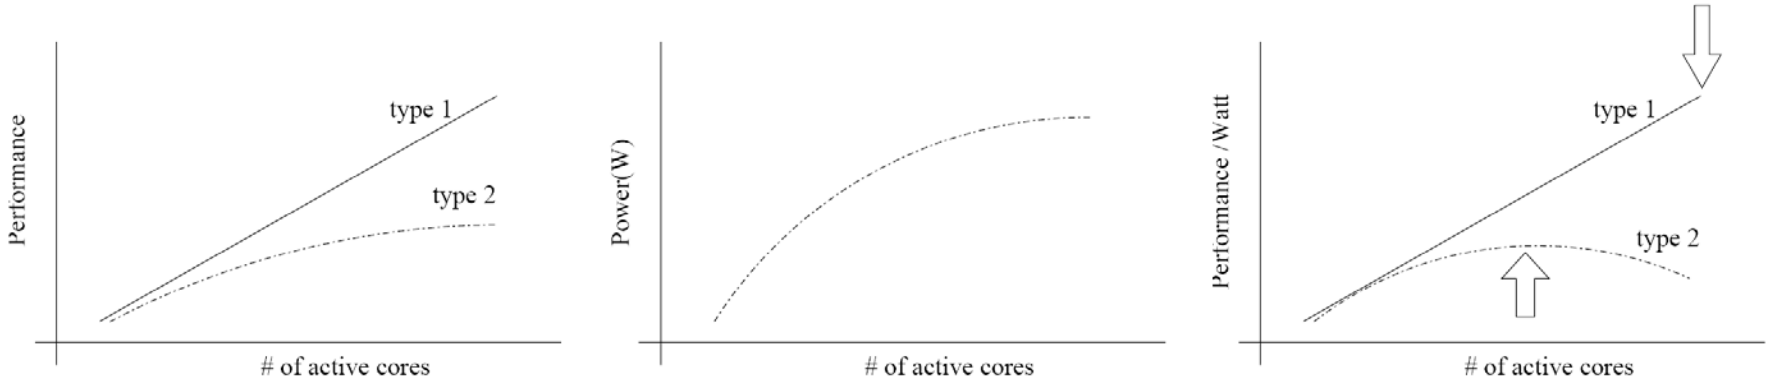

- Performance, Power and Efficiency vs. Number of Active Cores
- Type 1 **utilizes** all the cores (Performance increases linearly)
- Performance for type 2 **saturates** (Due to bandwidth limitations)

Hong, Kim, ISCA'10

Results

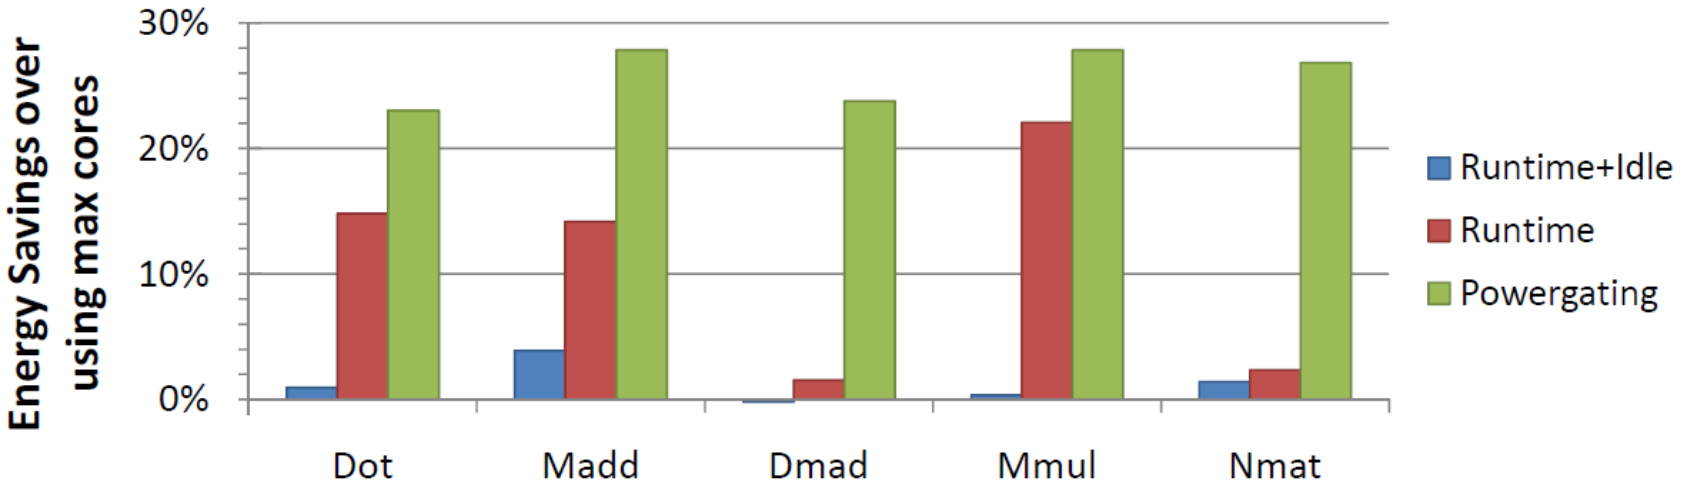

- Energy savings using optimal number of cores based on IPP system (NVIDIA GTX 280 and power gating GPUs)

Usage 2: Dynamic Mapping

- Distribute work between CPU and GPU for power and performance

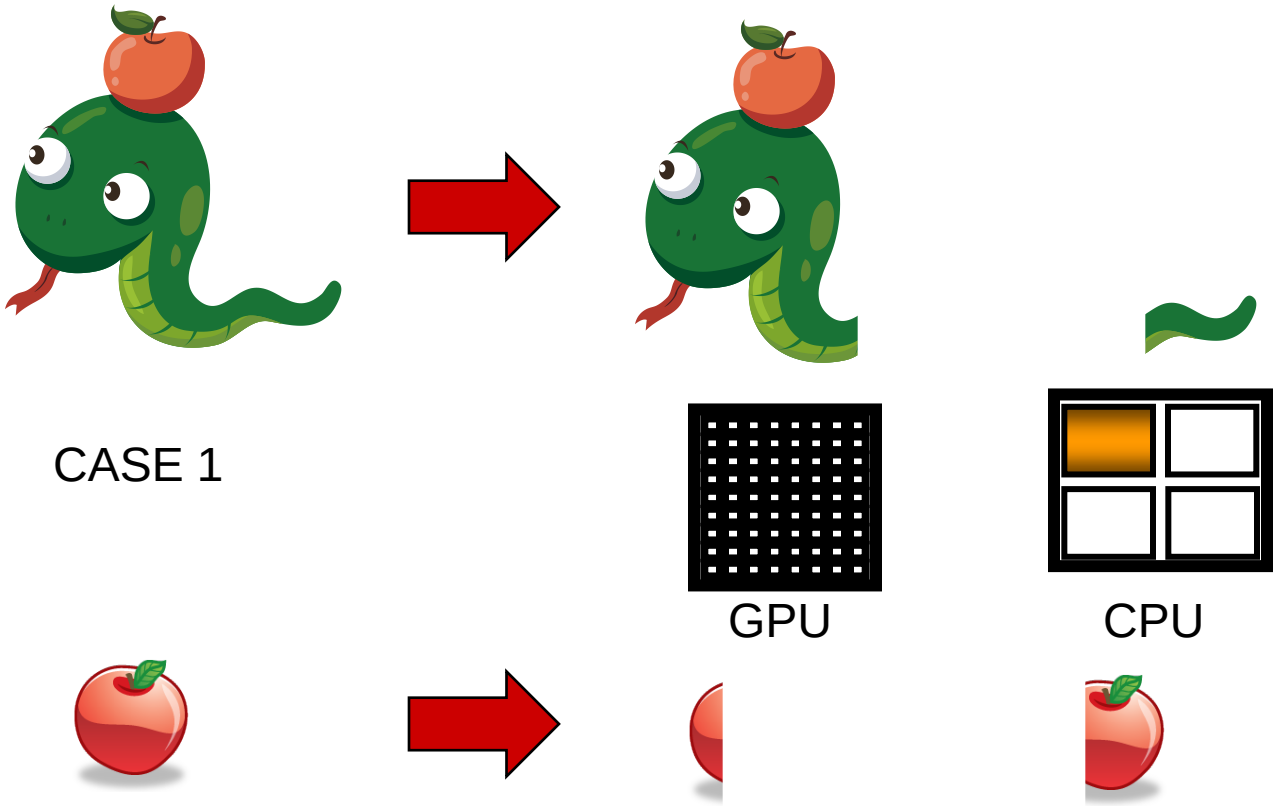

Luk & Hong & Kim MICRO'09

Energy Consumption

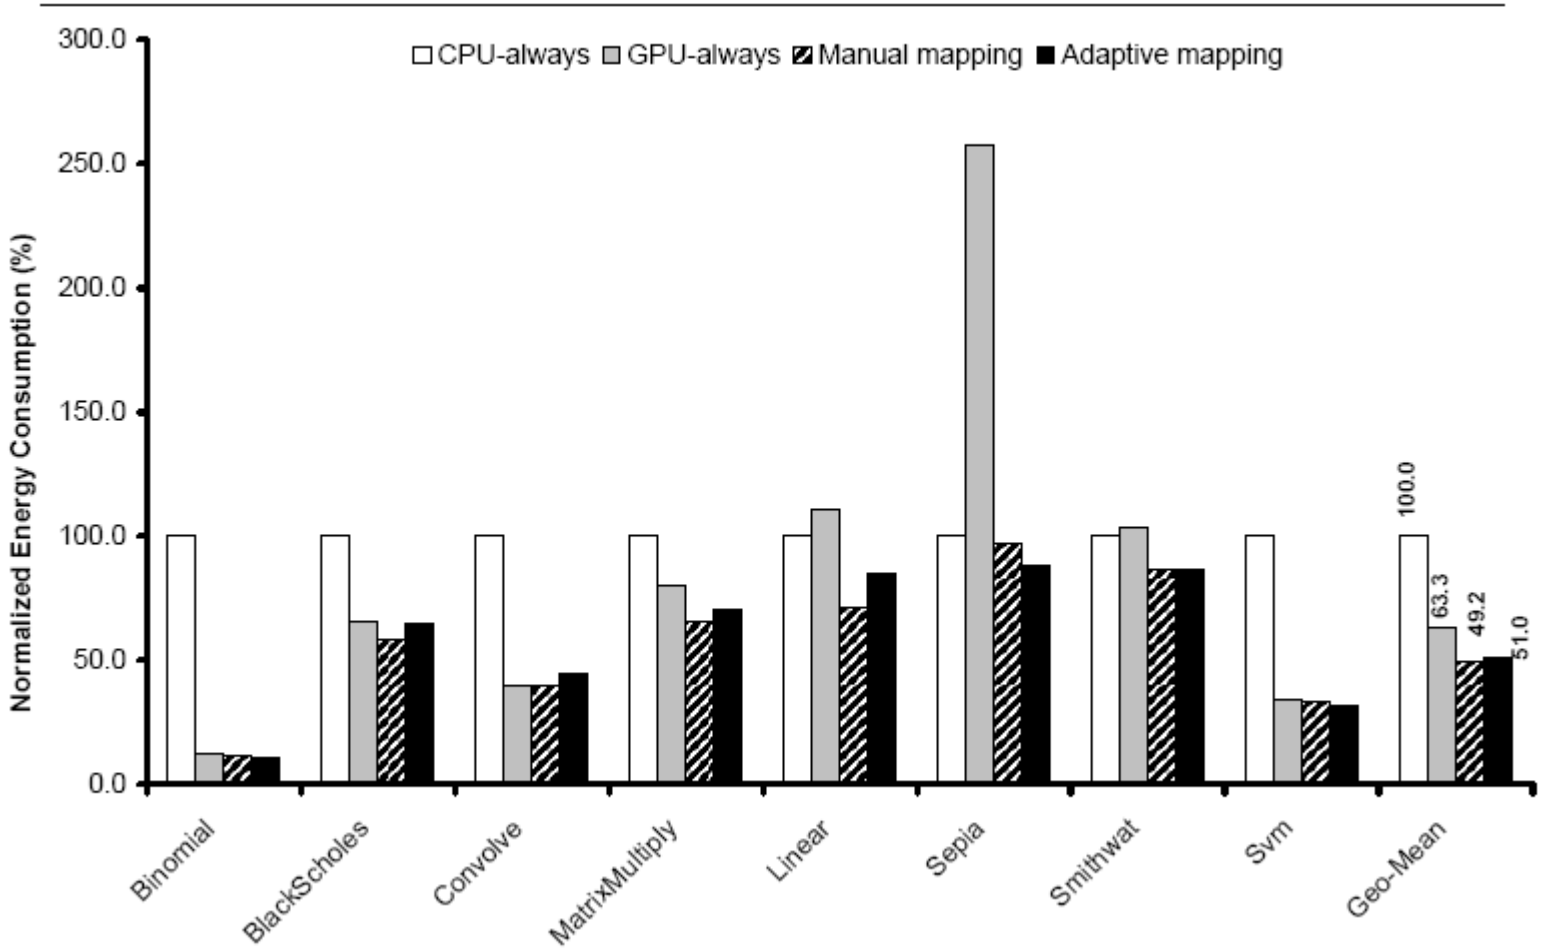

Luk & Hong & Kim MICRO'09

TEMPERATURE MEASUREMENTS

Using Thermal Camera

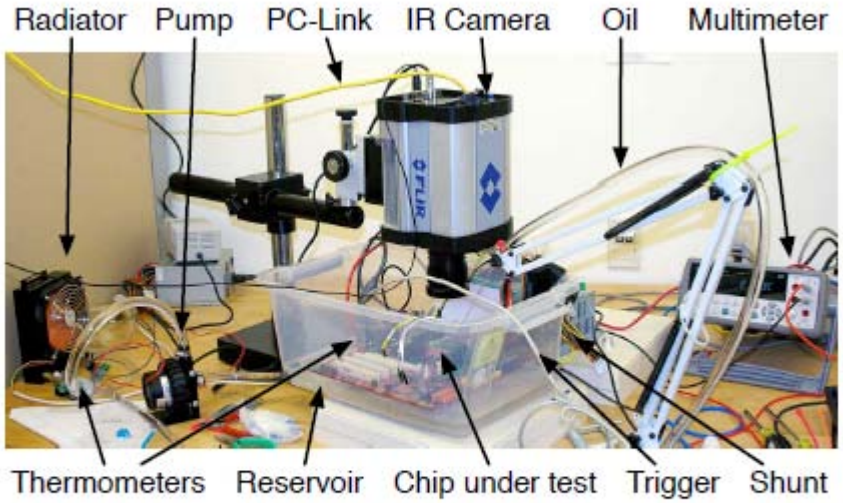

Figure 1: Measuring setup

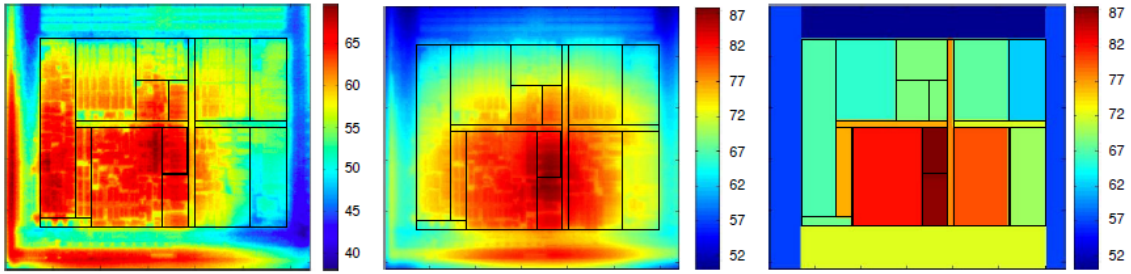

Figure 3: Full-thermal image with overlapped floorplan (a); hottest captured image (b); and its average temperature per block (c).

http://users.soe.ucsc.edu/~renau/rpapers.html

Using Thermocouples

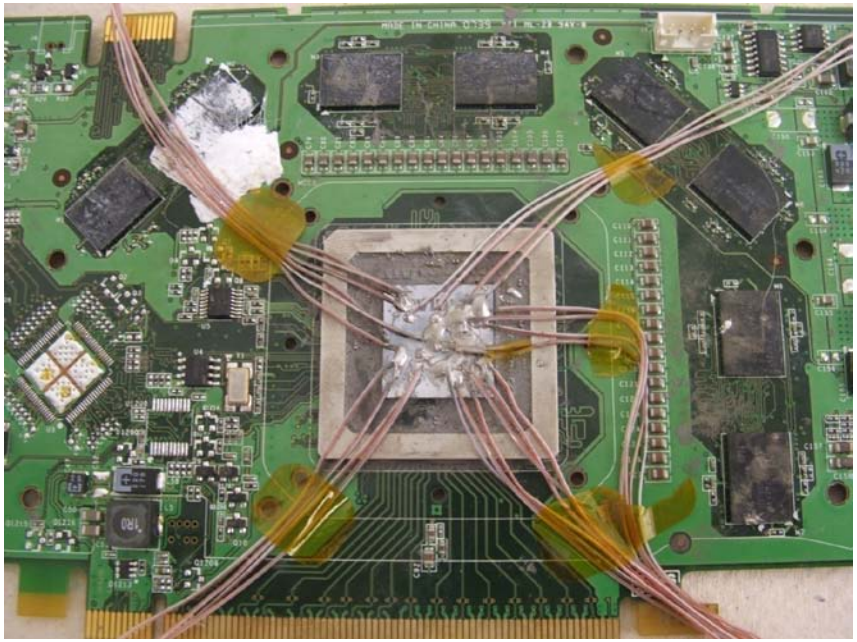

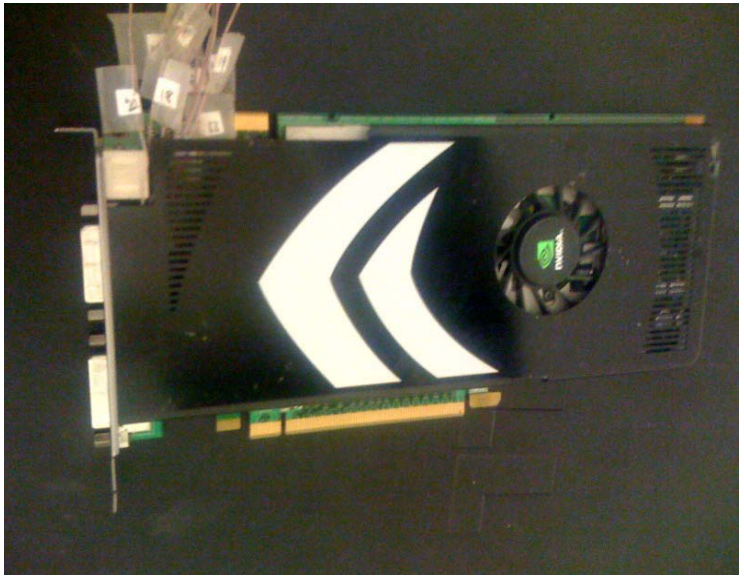

With Thermocouples, the original case can be assembled
so we can measure power & temp behavior with the original heat sink

On-going Researches Topcis

- Thread scheduling policies
- DVFS managements

H	C	C	H
C	C	C	C
C	C	C	C
H	C	C	H

H	H	C	C
H	H	C	C
C	C	C	C
C	C	C	C

H: running hot jobs, C: cold jobs or low voltage, turn-off power

- Multi-core lay out design to reduce thermal gradient, effective thermal transfers

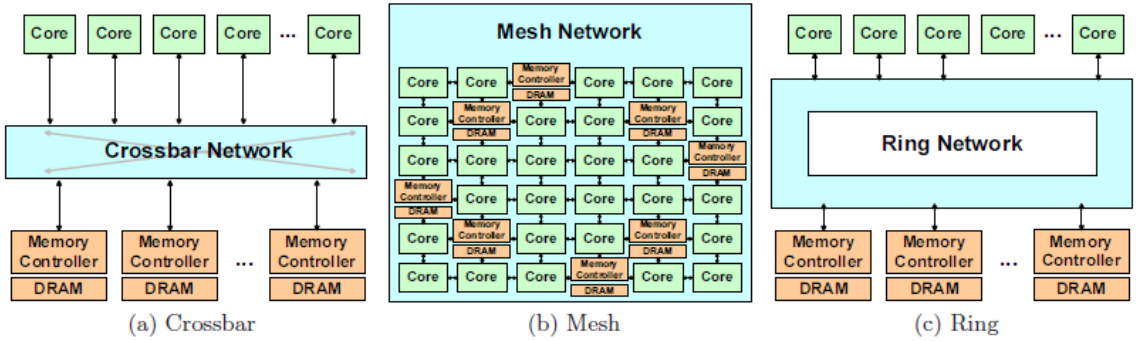

(a) Crossbar (b) Mesh (c) Ring

- Leakage power model (combine with themperature model)

Figures are from http://www.ece.ubc.ca/~aamodt/papers/gyuan.micro2009.pdf

Conclusions

- Power & temperature measurements opens an new architecture research
  - Architecture design (not only performance but also power)
  - Application scheduling, resource managements
  - Dynamic compilation system
- Experimental measurements and simulations together can extend more wider range of research

Georgia Tech comparch

Dream Lab

- Power measurement systems
  - Current –probed oscilloscopes
  - IR measurement systems
  - Current measurement controllers
- Probe stations
- Temperature measurements
  - Thermal image camera, oil cooling system
  - Multiple of thermocouple measurement systems
- Facility
  - Humidity, temperature controlled room
  - High current power supply
- Custom boards, custom riser cards