UCSD Jacobs | Computer Science and Engineering

Acknowledgement: Sun Microsystems, CNS, NSF GreenLight, UC Micro, Cisco, GSRC, DARPA

THERMAL MANAGEMENT IN HETEROGENEOUS MULTIPROCESSOR SOCS

Shervin Sharifi Ayse K. Coskun Tajana Rosing

Computer Science and Engineering Department, University of California San Diego (UCSD)

THERMAL CHALLENGES

- **Thermal Hot Spots**
  - High leakage power
  - Slower devices
  - Degraded reliability
  - Higher interconnect resistivity
- **Thermal Cycles**
  - Higher permanent failure rate
- **Spatial Gradients**
  - Timing failures
  - Increased interconnect delay and IR drop

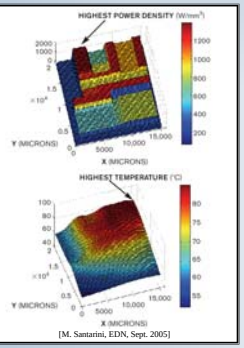

[M. Santarini, EDN, Sept. 2005]

HETEROGENEOUS MPSOCS

- **MPSoCs vs. Single core CPUs**
  - Higher performance per Watt
  - Lower design complexity
- **Heterogeneous MPSoCs**
  - Integrate cores with different architectures & various power/performance characteristics
  - Allow customized power/performance for the current requirements of application
- **Challenges**
  - Efficient tuning of power/performance under thermal constraints
  - Large thermal variations due to inherent disparities in power densities across the chip

FLOW

- **Offline Phase → Performed *once* at design time**
  - Calculates voltage/frequency configurations for thermal balancing
  - Characterizes tasks to pick the preferred core for each task
- **Online Phase → Performed *repeatedly* at runtime**
  - One of the V/F configurations calculated in offline phase is selected to meet the current performance demand
  - Tasks are balanced among cores to maximize the performance under the current V/F setting

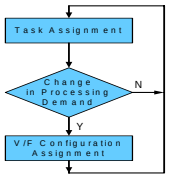

OFFLINE PHASE

A. Voltage/Frequency Configuration Calculation

1. For a target temperature and the given layout and core characteristics compute the V/F configurations that balance the temperature
2. Estimate the *processing capability* for the V/f configurations computed in step (1)
3. Generate a table containing the processing capability values and corresponding V/F configurations

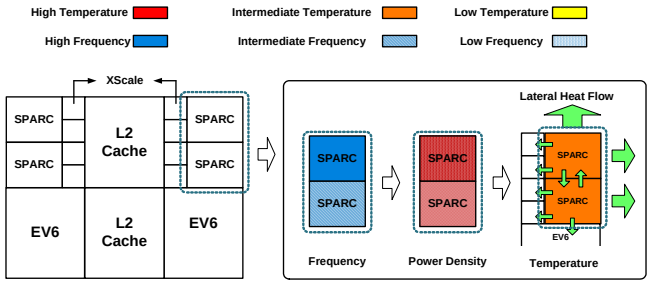

B. Task Characterization

- Determines which core type is the best fit to run a task on based on runtime changes when run on different cores

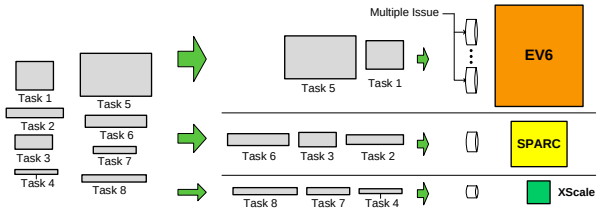

ONLINE PHASE

1. Voltage/Frequency Selection

- Processing demand is calculated based on the task arrival rates and lengths
- The lowest processing capability which meets the current demand is selected from the table generated in offline phase (A)
- The corresponding V/F configuration is assigned to the MPSoC

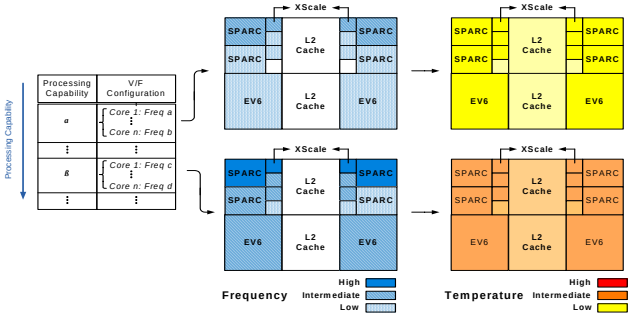

2. Task assignment

- Each core type (e.g. EV6, SPARC, Xscale) has a corresponding queue
- Tasks suitable for a certain type of core are directed to the appropriate queue
- Tasks are moved among queues if there is imbalance

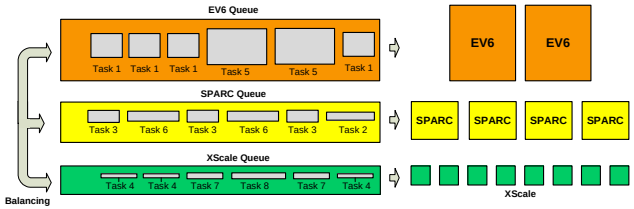

EXPERIMENTAL RESULTS

- **Benchmarks from the MiBench suite :**
  Automotive, networking and telecommunication applications (fft, disjkstra, bitcount, patricia, etc.)
- **Thermal Simulator:**
  HotSpot [Skadron, ISCA'03] with a cheap package
- **Workloads:**
  A mix of the MiBench benchmarks [Guthaus et al. WWC'01]
  * Some benchmarks have low performance penalty when running on small cores → "Frequency Tolerant Benchmarks"

Benchmark mixes:

A — Contains equal number of frequency tolerant and in-tolerant benchmarks

B — Contains mostly frequency tolerant benchmarks

C

D

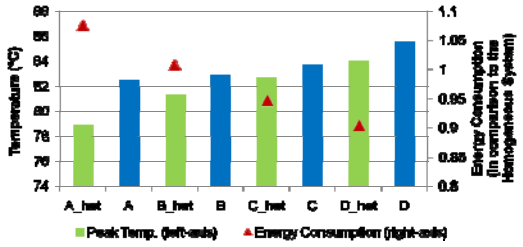

A, B, C, D: *Workload running on a homogeneous MPSoC of **4 EV6 cores***
_het:* *Workload on a heterogeneous MPSoC of **2 EV6s, 4 SPARCs and 8 XScales*

On a heterogeneous MPSoC of EV6, SPARC and XScale cores, our thermal management technique achieves:
- **Up to 10% additional energy savings**
- **Up to 4°C decrease in peak temperature**
- **Spatial gradients lower than 5°C**

in comparison to a homogeneous EV6-based MPSoC with the same die area, while maintaining the performance constraints.