Design Space Exploration and Dynamic Thermal Management of Multi-core Processors

Sarma Vrudhula
Department of Computer Science and Engineering
Arizona State University

(vrudhula@asu.edu)

Thermal Management (TM) for Multi-core

- Single or "few" cores:
  - DTM (clock gating, fetch throttling, DFVS, thread migration . .) employed infrequently, when temperature exceeds preset threshold
  - Package & cooling solution designed to handle worst-case power
- With "many" cores
  - large temporal variations in # of threads executing, hence, in power consumption
  - Not feasible to design package to remove worst-case power, but closer to average power
  - DTM will be used more frequently and will have a much greater impact on performance when compared to single cores
  - Need to account for coupled relation between power, speed and temperature

Speed, Power & Temperature

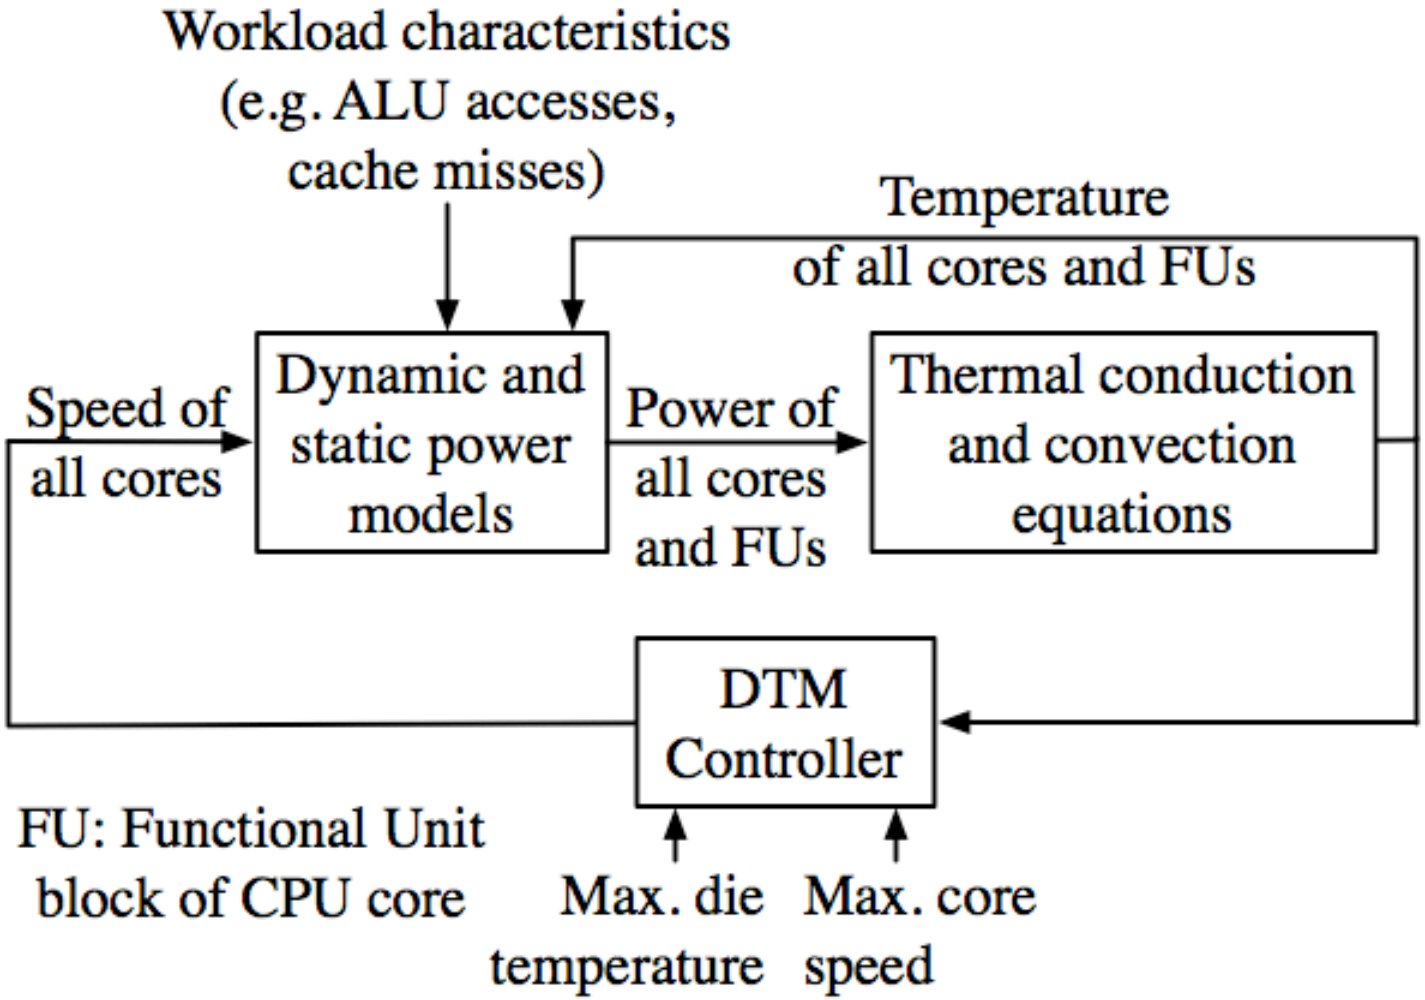

Hotspot thermal model

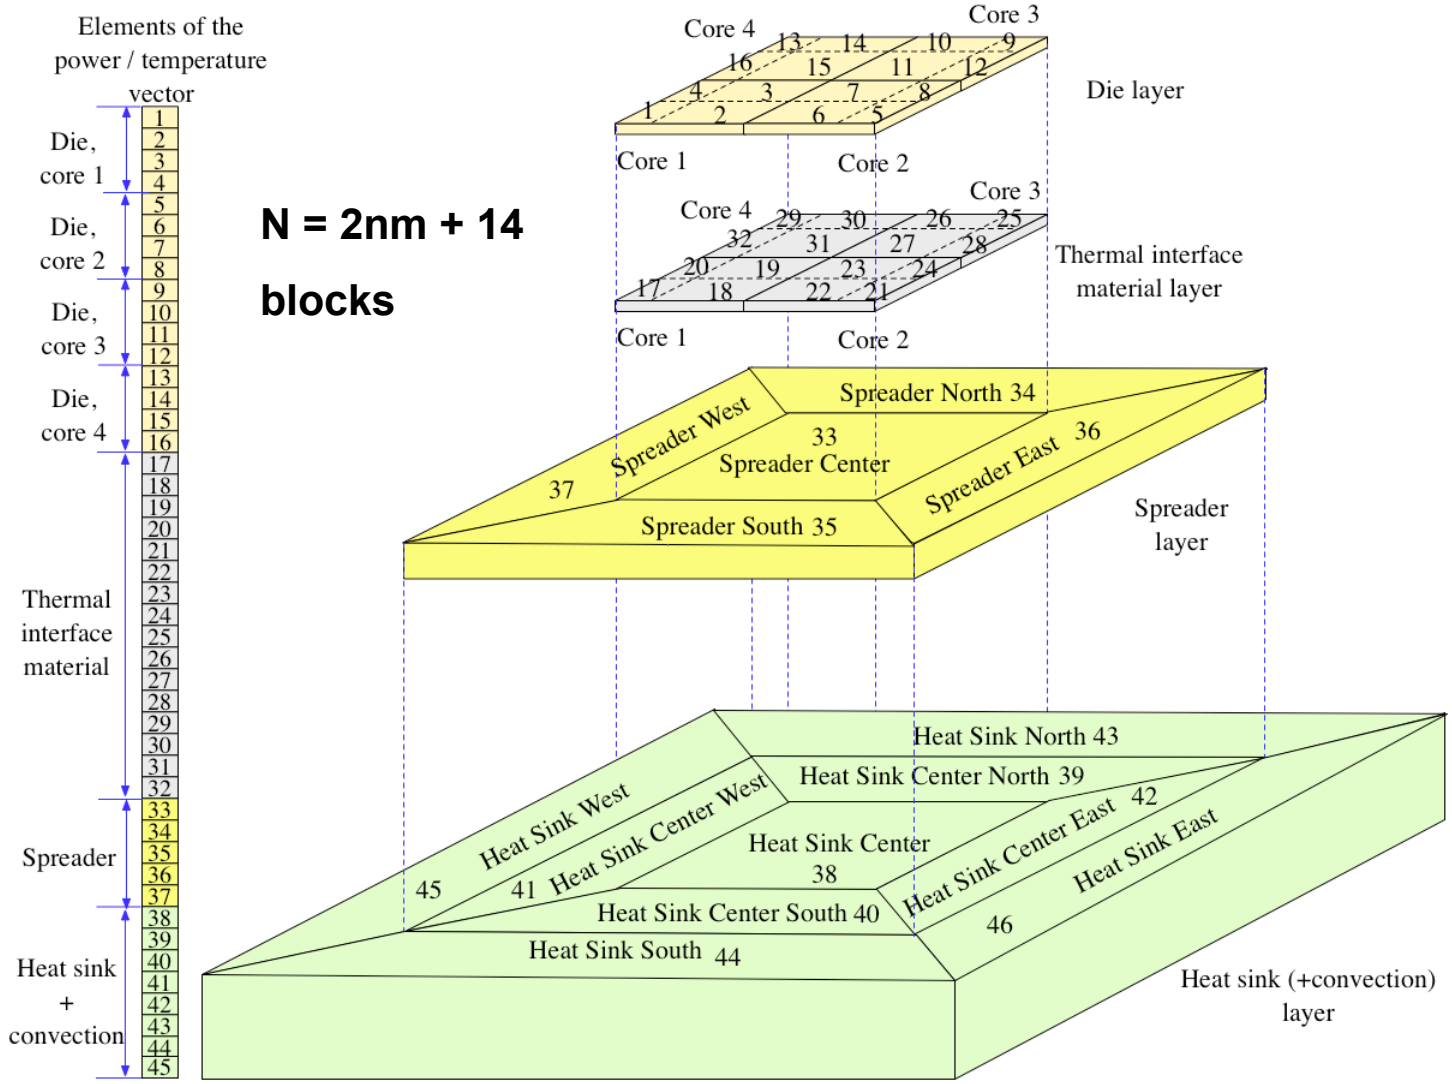

Hotspot RC Circuit Model

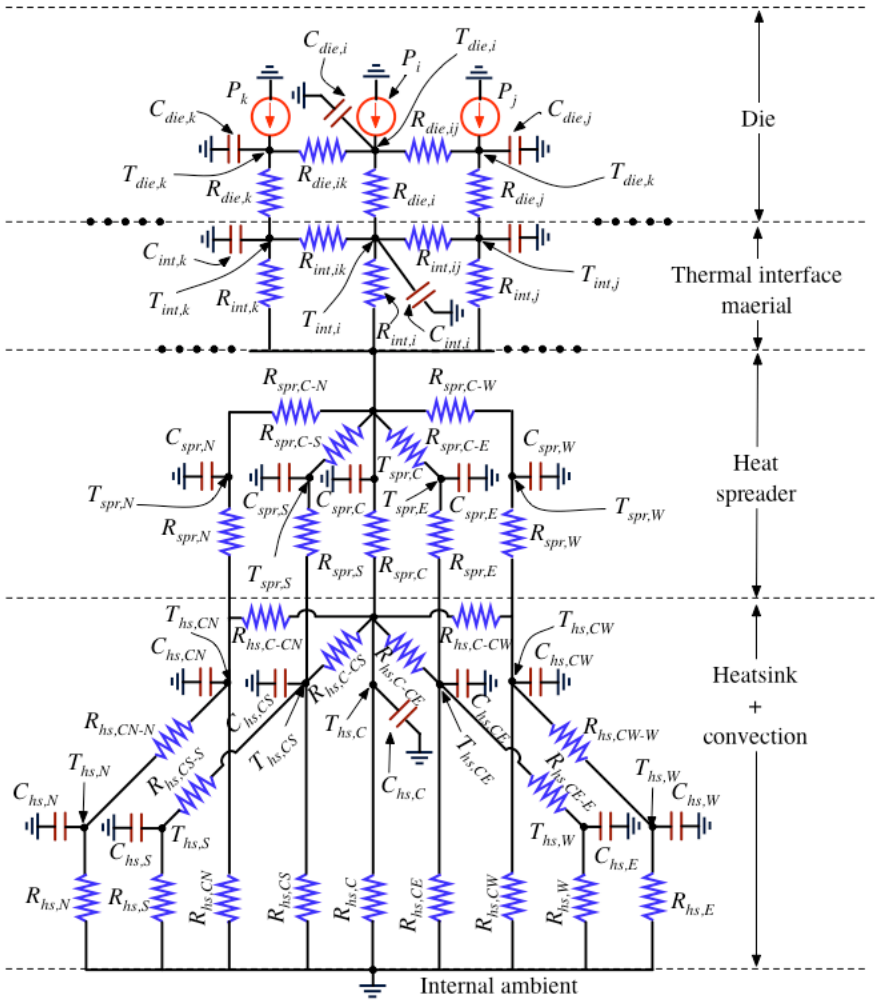

Subscripts for variables in the spreader and heat-sink layer
C-Center, N-North, S-South, E-East, W-West,
CN-Center North, CS-Center South, CE-Center East, CW-Center-West

Design Space Exploration

- What is the maximum number of cores that can be activated with and without throttling in steady-state?
- What is the steady-state throughput and power as a function of the number of active cores ?

Answers in terms of

- Static and dynamic power consumption of hottest block and full chip
- thermal resistances of hottest block
- leakage sensitivity to temperature
- ambient and chip threshold temperature

DSE requires requires fast and scalable (simple algebraic expressions) answers

Design Space Exploration$_2$

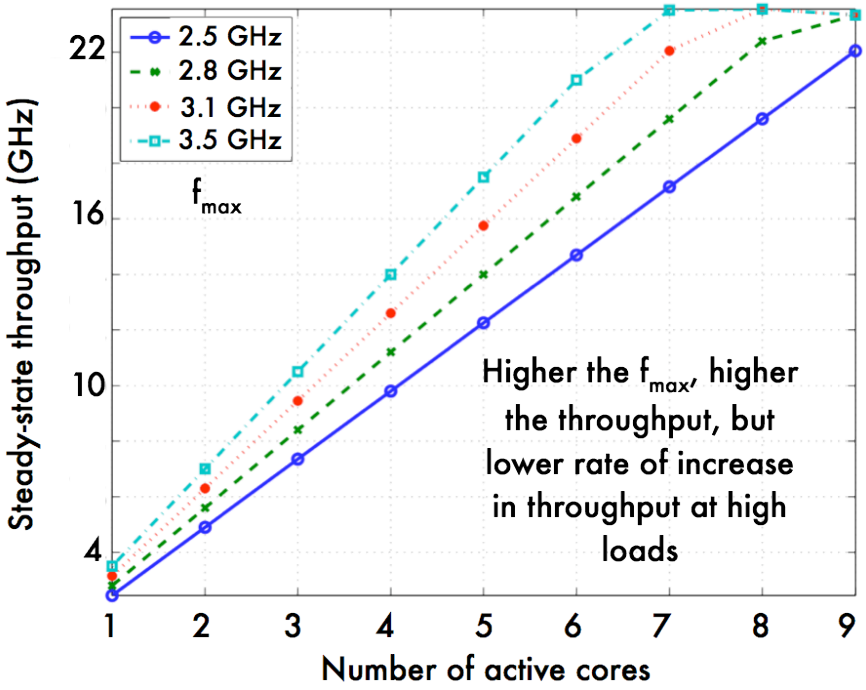

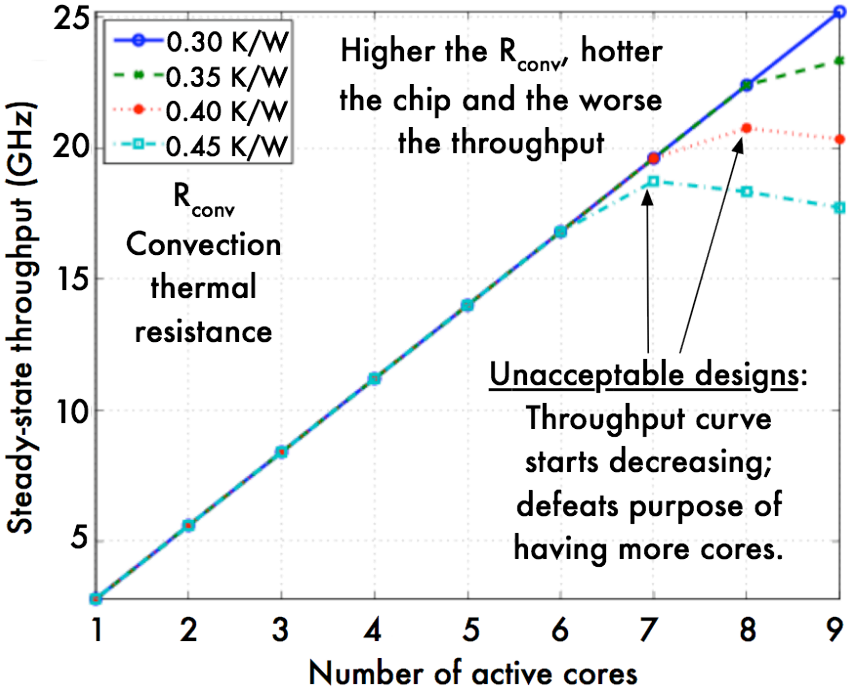

Dynamic Thermal Management

Steady-state problem formulation

$$\max_{\mathbf{s}_{ss}} \quad \mathbf{w}' \, \mathbf{s}_{ss}, \quad \text{(t'put = weighted sum of speeds)}$$

$$\text{s.t.} \quad d\mathbf{T}_{ss}/dt = 0 = \hat{\mathbf{A}} \, \mathbf{T}_{ss} + \mathbf{B}\left(\mathbf{P}_d(\mathbf{s}_{ss}) + \mathbf{P}_{l0}\right),$$

(steady-state thermal equation)

$$T_{i,ss} \leq T_{\max} \; \forall \; i \in \{1, \ldots, nm\},$$

(thermal constraint)

$$0 \leq \mathbf{s}_{ss} \leq 1.$$

(speed boundaries for each core)

LP: n decision variables, 2nm+14 constraints

Transient speed control problem

$$\max_{\mathbf{s}(t)} \quad \frac{1}{t_f}\int_0^{t_f} \mathbf{w}'\,\mathbf{s}(t)\,dt, \quad \text{(time-averaged throughput)}$$

$$\text{s.t.} \quad d\mathbf{T}(t)/dt = \hat{\mathbf{A}}\,\mathbf{T}(t) + \mathbf{B}\left(\mathbf{P}_d(\mathbf{s}(t)) + \mathbf{P}_{l0}\right), \mathbf{T}(0) = \mathbf{T}_0$$

(transient thermal equation - linear LDT)

$$T_i(t) \le T_{\max}\ \forall\, i \in \{1,\ldots,nm\}, 0 \le t \le t_f,$$

$$0 \le \mathbf{s}(t) \le 1\ \forall\ 0 \le t \le t_f.$$

Optimal control problem: n control variables (speeds), 2nm+14 state variables (temperature). Too big for analytical solution. Very expensive to solve numerically.

DTM Extensions

Stochastic Workloads:

- to model inexact nature of program execution and nondeterministic memory access patterns
- can be included in the existing framework by modeling the power consumption as a random process

Dynamic Task Allocation:

- Need fast (within OS scheduling interval) schemes to migrating threads between heated cores

Process Variations:

- Requires modeling the circuit & thermal parameters as realizations of spatial stochastic process
- Inter-core variations for the thermal parameters (Rs & Cs)
- Circuit parameters (Tox, Leff, Vt) have both intra-core and inter-core variations. Major impact on core leakage

DTM Extensions

- **Real-time Scheduling:**
  - Compute optimal speed profiles when tasks have hard deadlines, requiring minimization of makespan
- **Modeling of Interconnects:**
  - Interconnect power for 16 cores can be more than the combined power of 2 cores
  - Variations in interconnect will exhibit significant variations in power
- **Micro-architecture Components**
  - Incorporating dynamic and leakage parameteric models of micro-architectural components