Power Supply Noise-Aware Workload Assignments for Homogeneous 3D MPSoCs with Thermal Consideration

Yuanqing Cheng, Aida Todri-Sanial,
Alberto Bosio, Luigui Dilillo, Patrick Girard, Arnaud Virazel
LIRMM, Université Montpellier 2 / CNRS, France

4/11/2014

Outline

- Introduction
- Preliminaries
- Motivation through an example
- Problem formulation & proposed method
- Experimental results
- Conclusions

3D MPSoC Is Here !

- Moore's Law Ends ?

Transistor Count

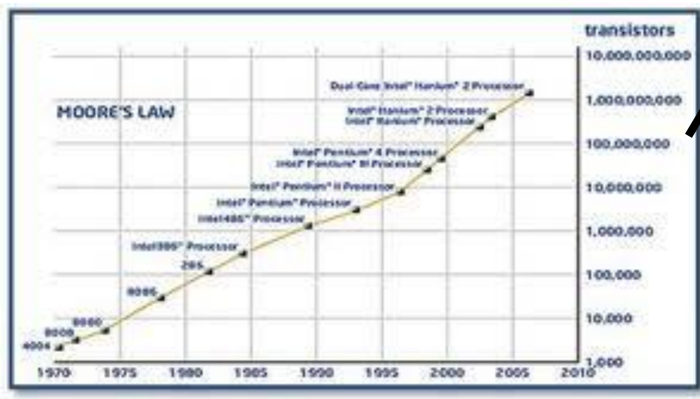

Interconnect Delay

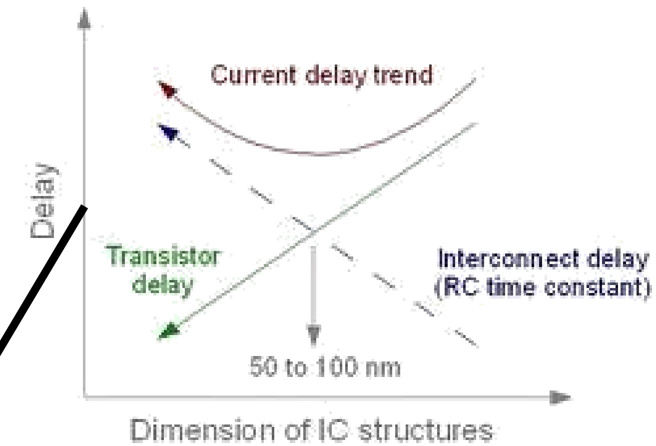

- MPSoC & 3D Technology

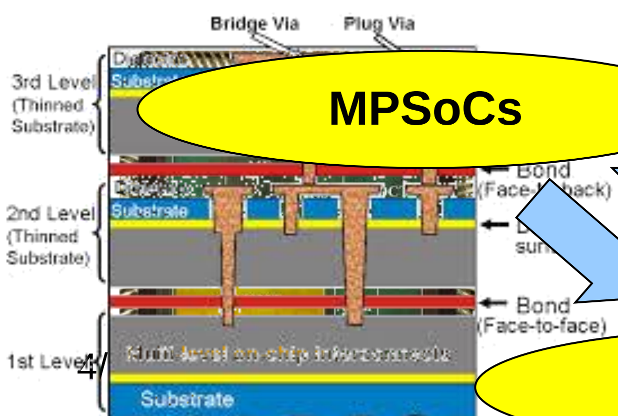

MPSoCs

3D Integration

1. D… …nnect delay thank… …ption of TSVs
2. Ease … integration of disparate tec…ologies
3. Smaller form factor …

3D MPSoCs

Signal Integrity & Thermal Challenges

- Power Supply Noise

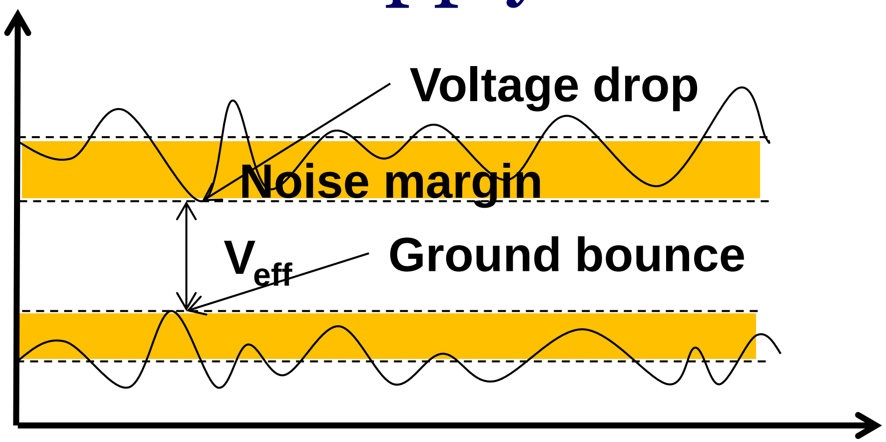

1. Current density increase
2. Supply voltage and transistor scaling
3. Shrinking noise margin
4. Power supply noise propagation among tiers thru. PG grid

Challenge 1. Signal integrity

- Heat dissipation

1. Non-uniform thermal distribution among layers
2. Thermal correlation among layers
3. Need effective and efficient thermal model

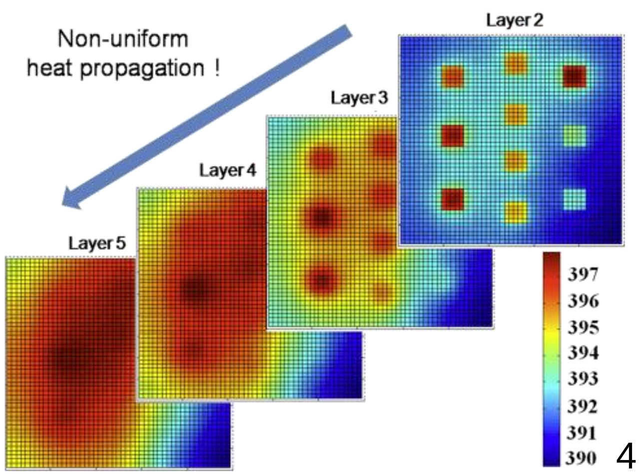

Related Work

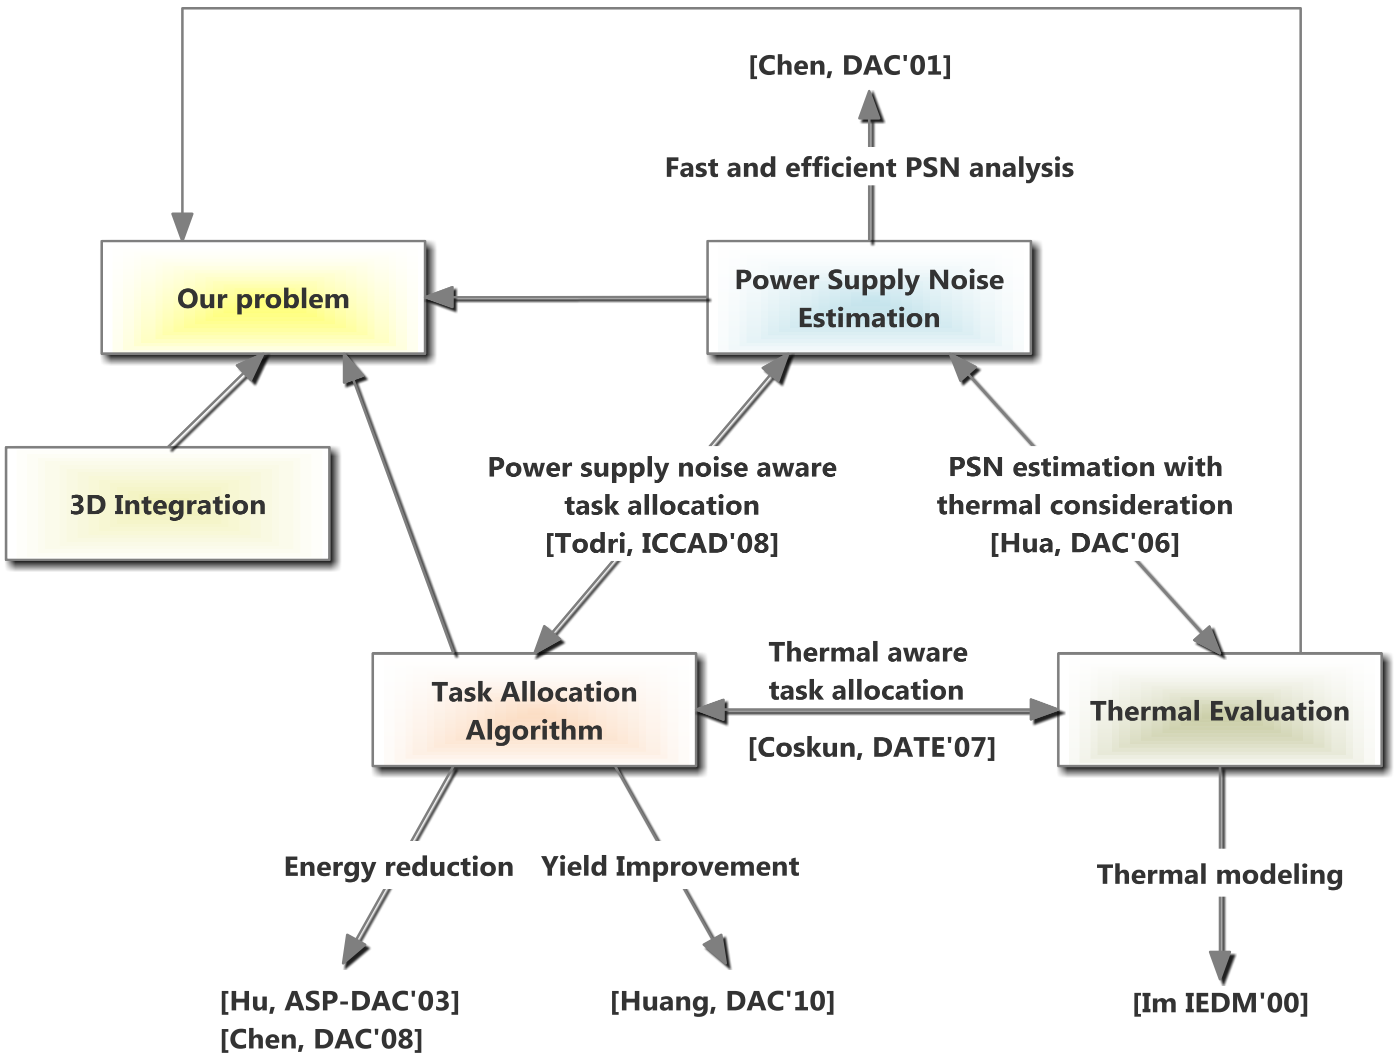

Overview of Our Research Work

Workloads

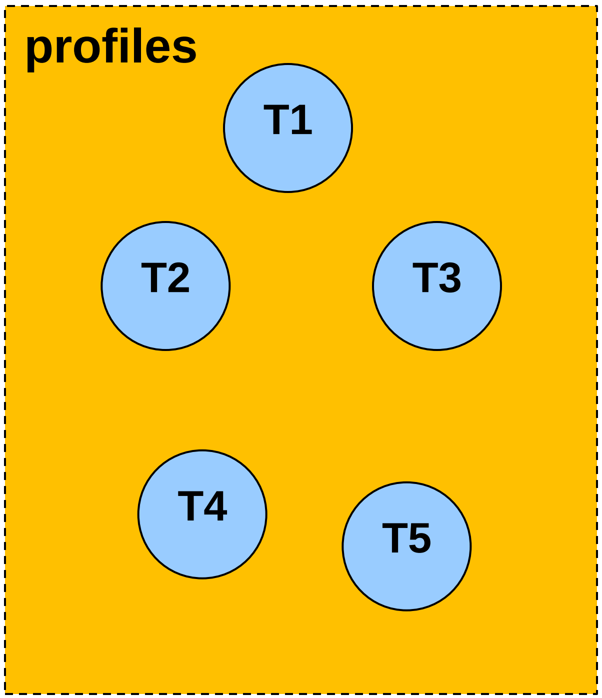

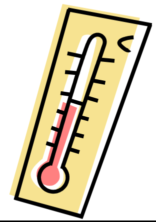

Thermal Cons.

3D Homogenesous MPSoC

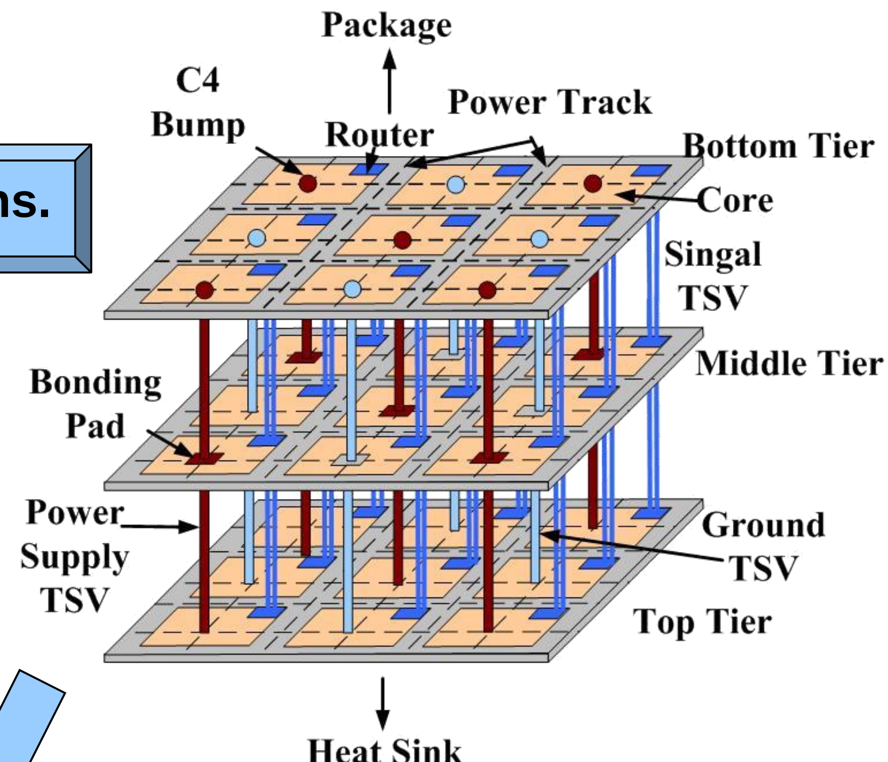

PSN minimization,
Performance maximization

Outline

- Introduction
- Preliminaries
- Motivation through An Example
- Problem formulation & proposed method
- Experimental Results
- Conclusions

Workload Assignment Challenges on 3D MPSoCs

- Challenges
  - Workload Characterization
  - Power Supply Noise Coupling
  - Thermal Coupling
  - Efficient and effective method is required

Workload Characterization

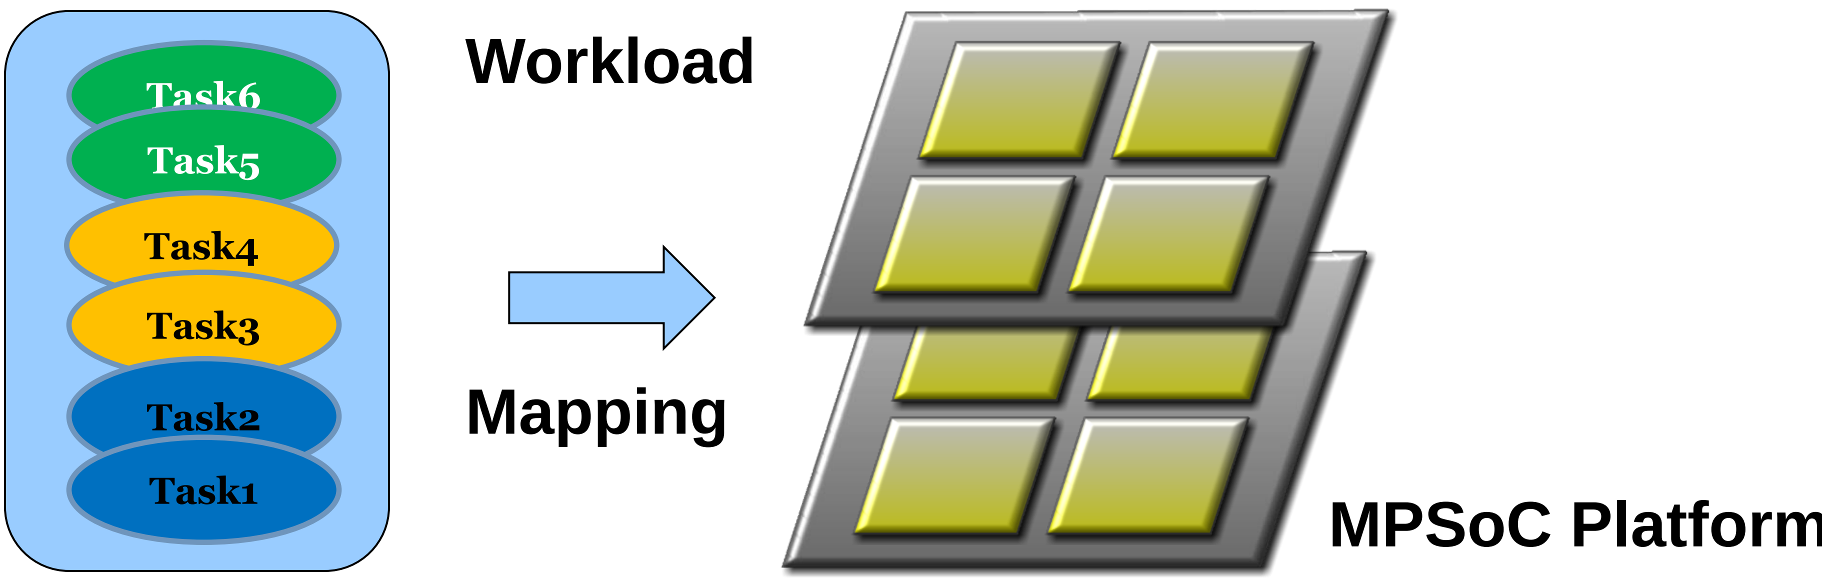

1. Different switching activities → different PSNs
2. Different running powers → different thermal dissipations
3. PSN coupling through TSVs

Workload Characterization (Cont.)

- Workload characteristics
  - Switching activity
    - Switching period
    - Rise / fall time

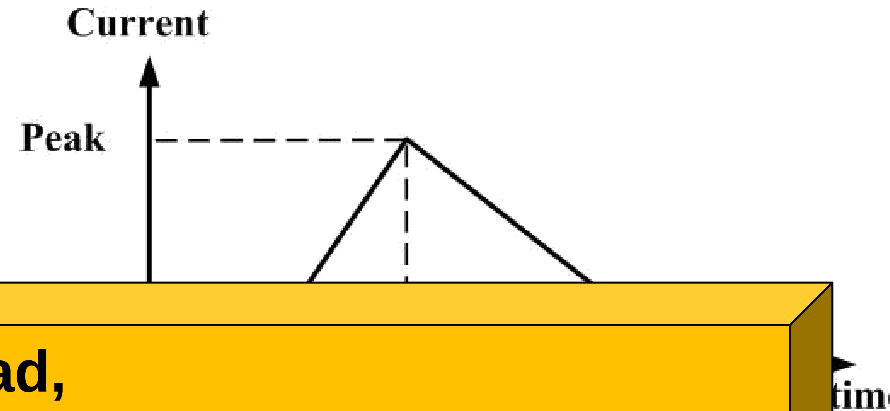

Workload,
{leakage, peak current, rise/fall time, period, average power}

Thermal behavior

Power Supply Noise

Power Grid Model of 3D MPSoCs

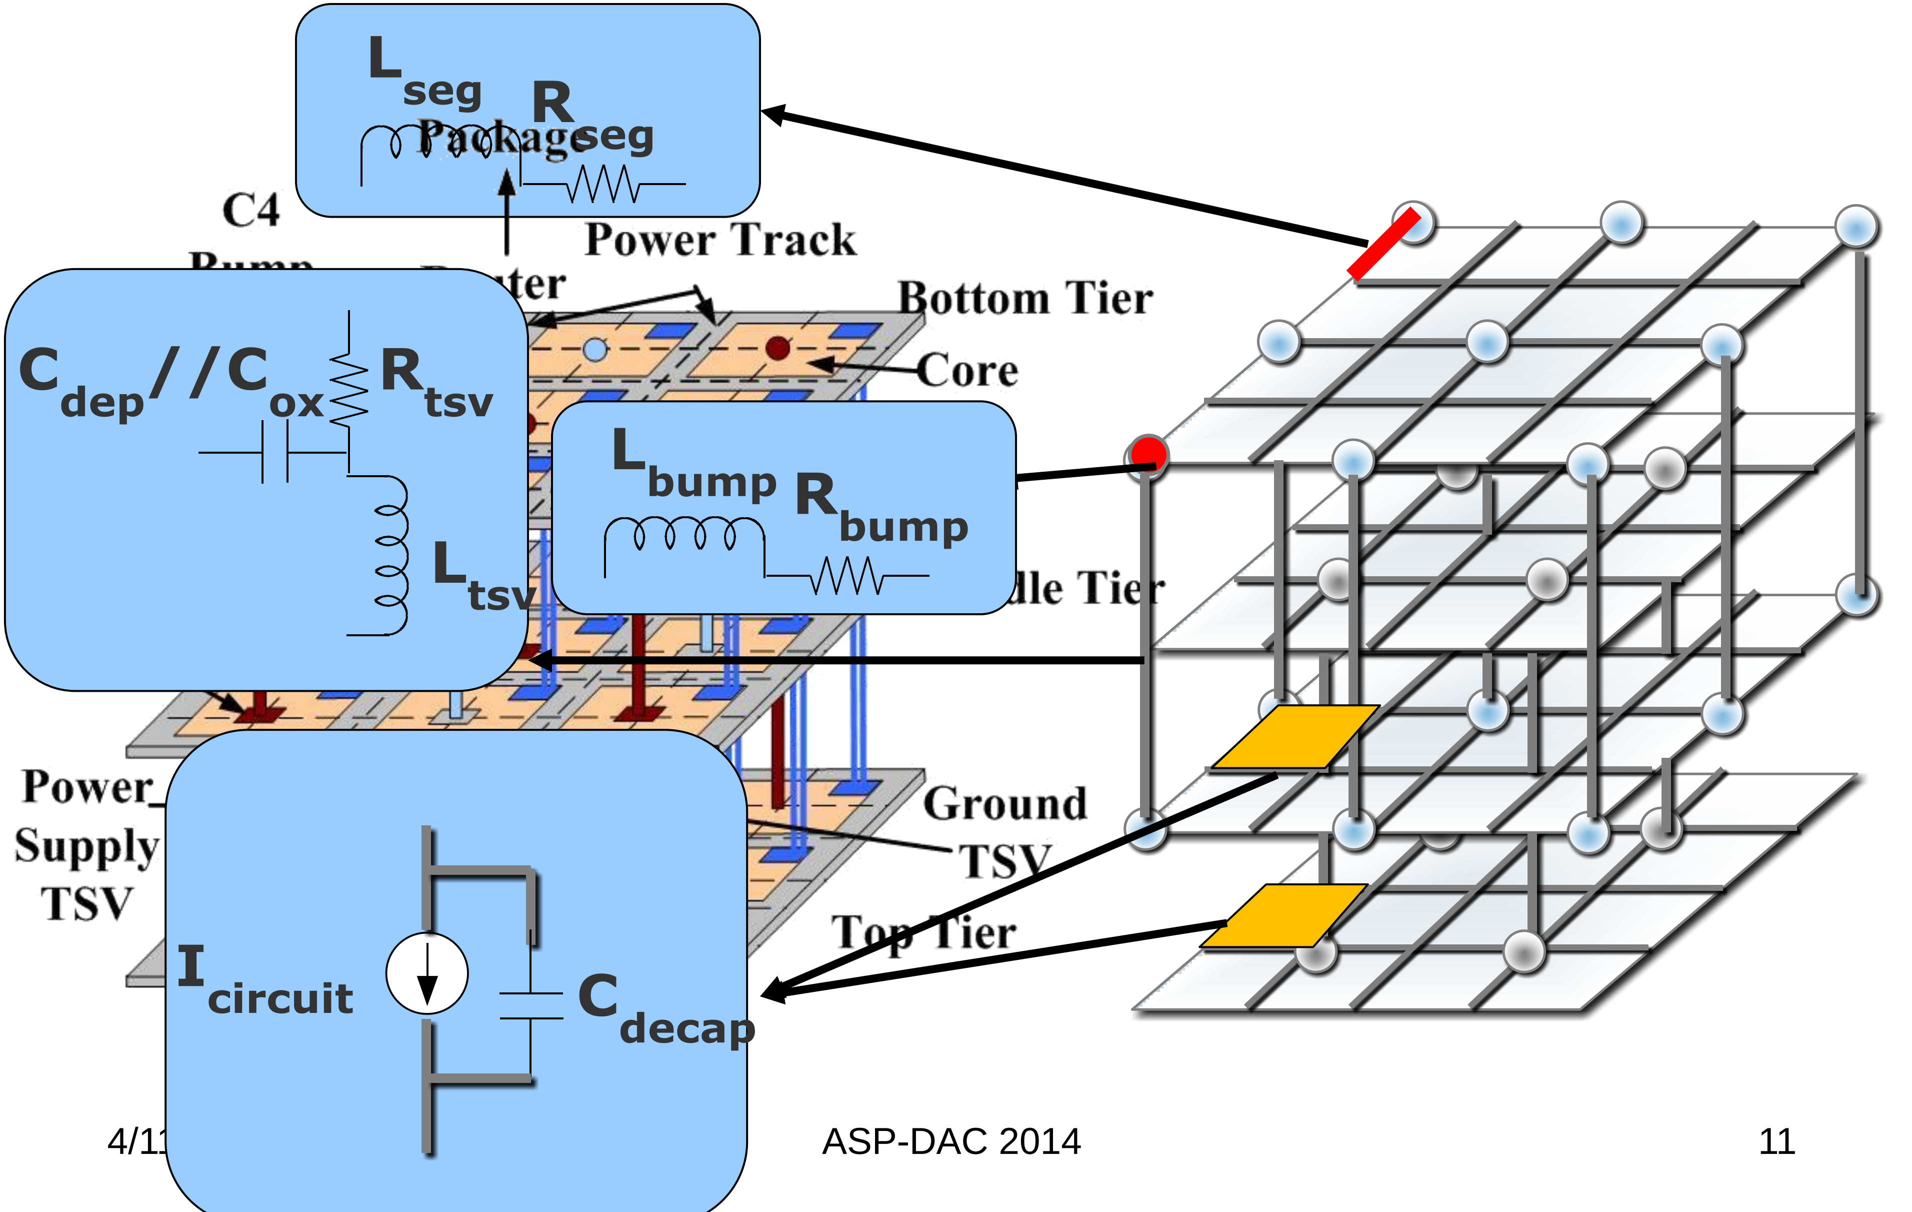

Physical Parameters for 3D MPSoC Power Grid

Interconnect		TSV	
Segment length	100µm	Diameter	10µm
Metal width	15µm	Aspect ratio	1:8
Metal space	15µm	Resistance	25mΩ
Segment resistance	122mΩ	Inductance	25pH
Segment inductance	60pH		
Segment Capacitance	152.7fF		

C4 Bump		Core	
Resistance	9.52mΩ	Size	1.6mm×1.6mm
Inductance	12.65pH	Mesh grid	16×16 / core

Thermal Modeling

- PDE based theoretical analysis

$$\rho c_p \frac{\partial T(\boldsymbol{r},t)}{\partial t} = k_t \nabla^2 T(\boldsymbol{r},t) + g(\boldsymbol{r},t)$$ **(Transient)**

$g(r)$ **state)**

High computing complexity

- Finite … grid based differential methods
  - Hotspot …
- Thermal estimation by power information
  - Easy to use for online thermal evaluation

Thermal Evaluation by Stack Power

- Concept of "stack power" for 3D MPSoCs

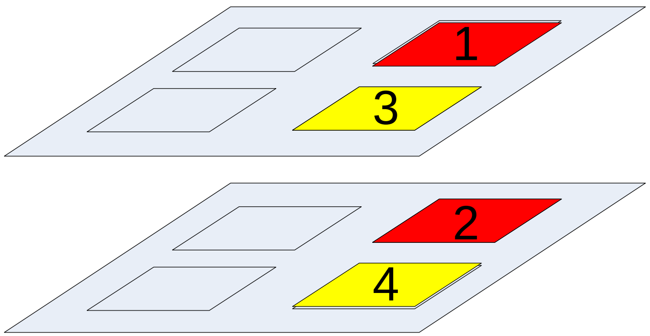

$$Power(stack_1) = Power(core_1) + Power(core_2)$$

$$Power(stack_2) = Power(core_3) + Power(core_4)$$

- Power gradient

$$Power\_Gradient(stack_1, stack_2) = \left| Power(stack_1) - Power(stack_2) \right|$$

$$\forall i, j \in S_{stack}$$

$$Power\_Gradient(stack_i, stack_j) = \left| Power(stack_i) - Power(stack_j) \right| \leq P_{\max}$$

Outline

- Introduction
- Preliminaries
- Motivation through An Example
- Problem formulation & proposed method
- Experimental Results
- Conclusions

Motivation - An illustrative Example

- Workload characterization

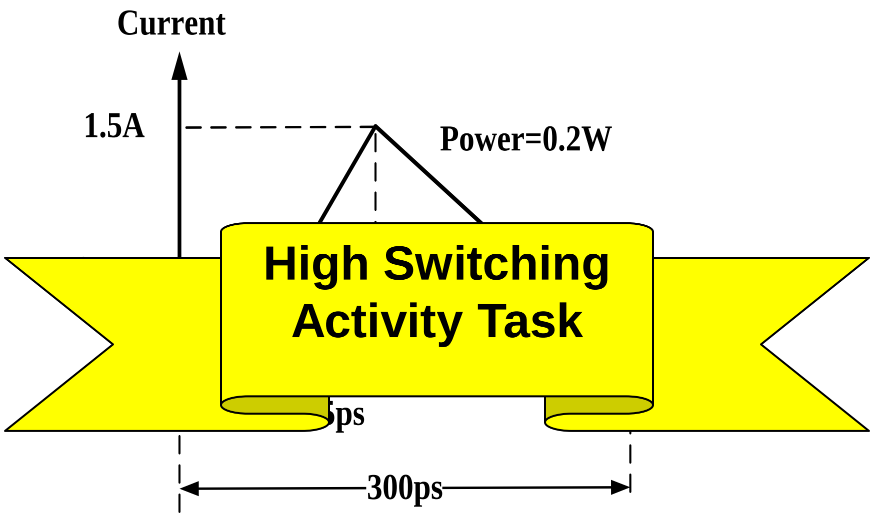

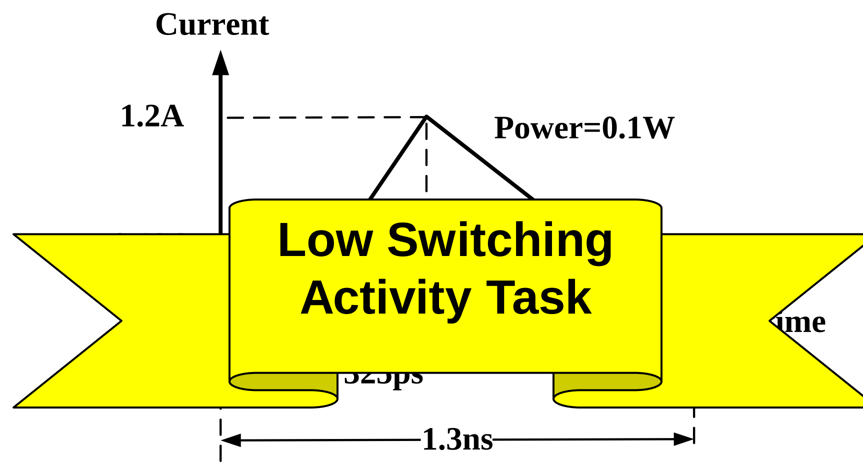

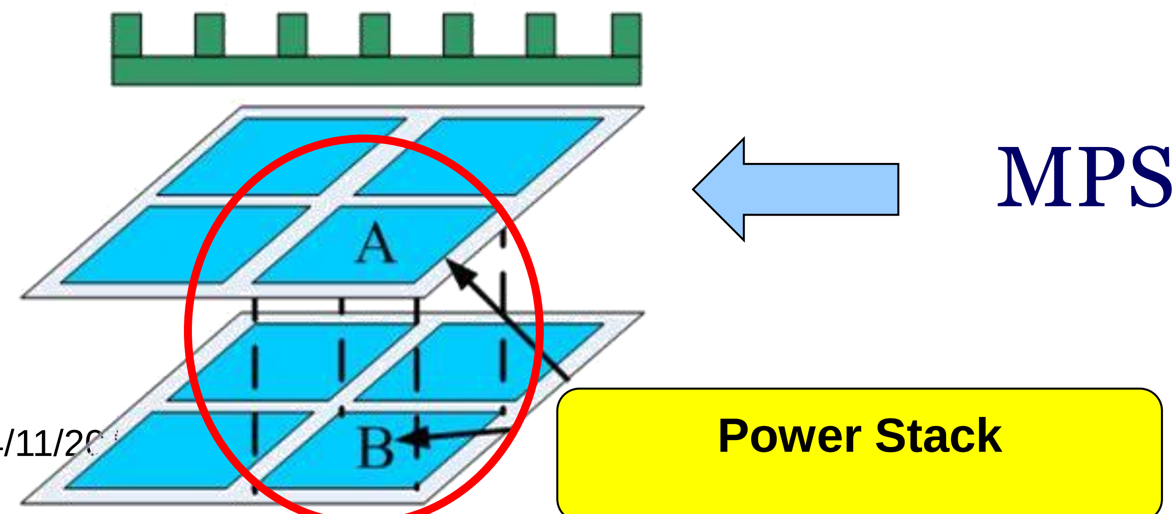

MPSoC Target

Motivation - An illustrative Example

- Different scheduling schemes with the same stack power gradient

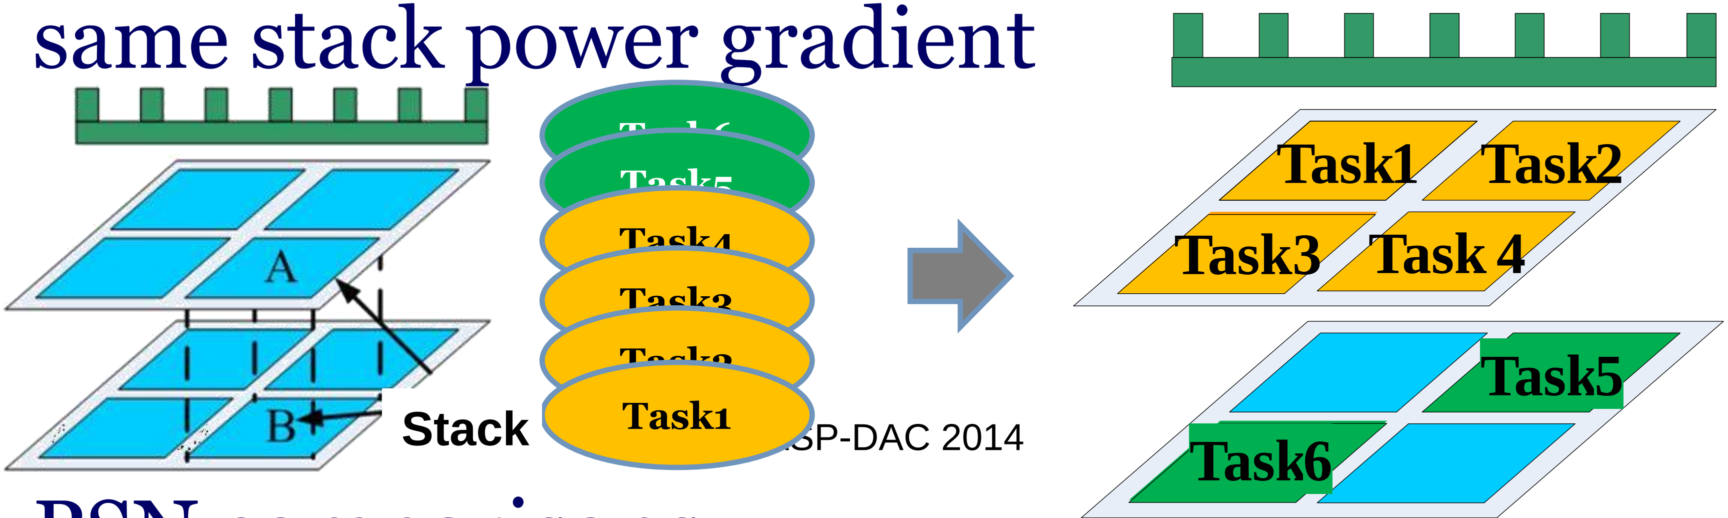

- PSN comparisons

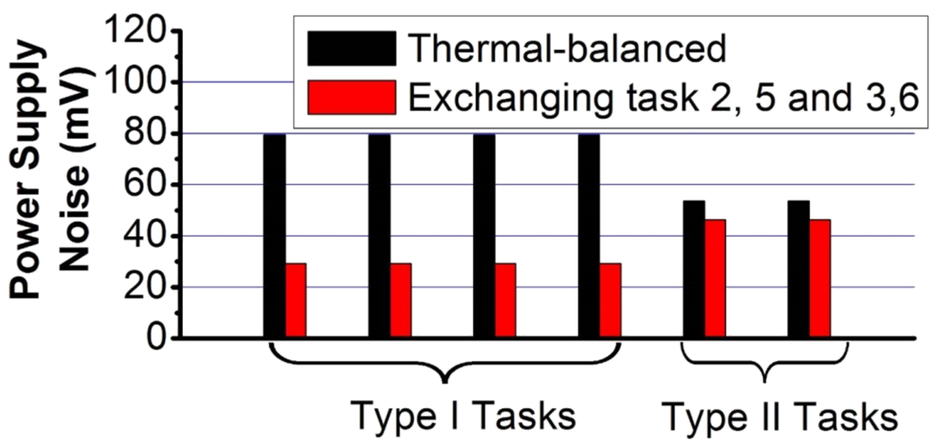

Power Supply Noise & Performance Calculation

- Voltage drop & ground bounce

$$V_{pnoise} = \int_{t_s}^{t_e} \max\{V_{dd} - V_p, 0\} dt / (t_e - t_s) \qquad V_{gnoise} = \int_{t_s}^{t_e} \max\{V_{ss}, 0\} dt / (t_e - t_s)$$

$$V_{noise} = V_{pnoise} + V_{gnoise}$$

- Timing Variations due to PSN

$$\frac{D}{D_0} = 1 + k_1 \frac{\Delta V}{V_{dd} - V_t} + k_2 \left(\frac{\Delta V}{V_{dd} - V_t}\right)^2$$

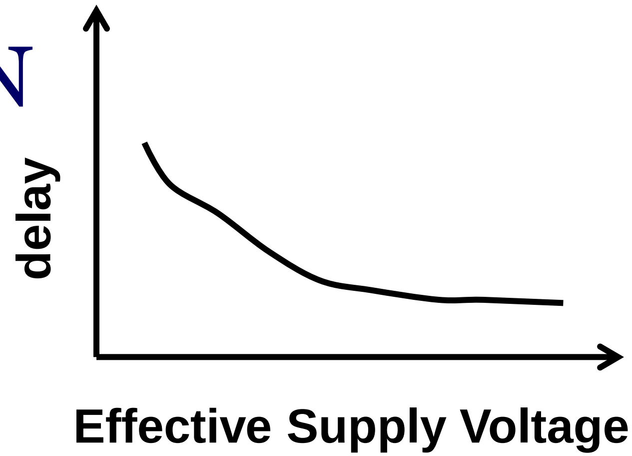

- Performance Evaluation

$$MIPS_i = \frac{IPC_i \times f_i}{10^6} \qquad Performance = \sum_{i=1}^{p} \mathrm{MIPS}_i$$

Outline

- Introduction
- Preliminaries
- Motivation through An Example
- Problem formulation & proposed method
- Experimental Results
- Conclusions

Problem Formulation

- Given,
  - 3D MPSoC architecture **AR(m,n,l)**
    - **m**×**n** PEs within one layer
    - **l** stacking layers
  - Workload profiles

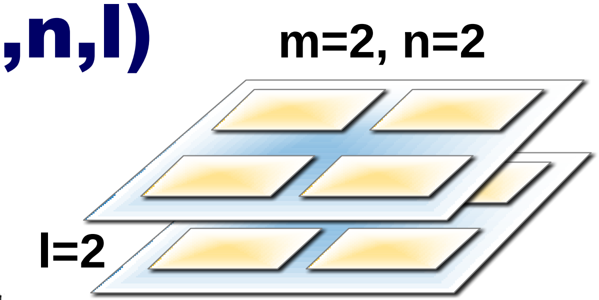

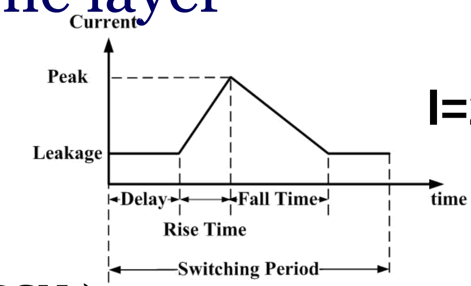

- Maximize: $\sum_{i=1}^{p} MIPS_i(PSN_i)$
  - Where $PSN_i = \max\{PSN_j\}, Node(j) \in PE_i$
  - $Gv + C\frac{dv}{dt} = i, PSN_j = \int_{t_s}^{t_e} \max\{v_{dd} - v_i, 0\}dt / (t_e - t_s)$
- Constraint $Power\_Gradient(stack_i, stack_j) = \left|Power(stack_i) - Power(stack_j)\right| \le P_{\max}$

Power Supply Noise Estimation Method

- Challenge
  - High computing overhead
- Main idea
  - Estimate relative PSN magnitude
  - Superposition PSN impact of individual task

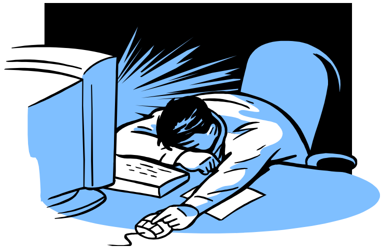

An Illustrative Example

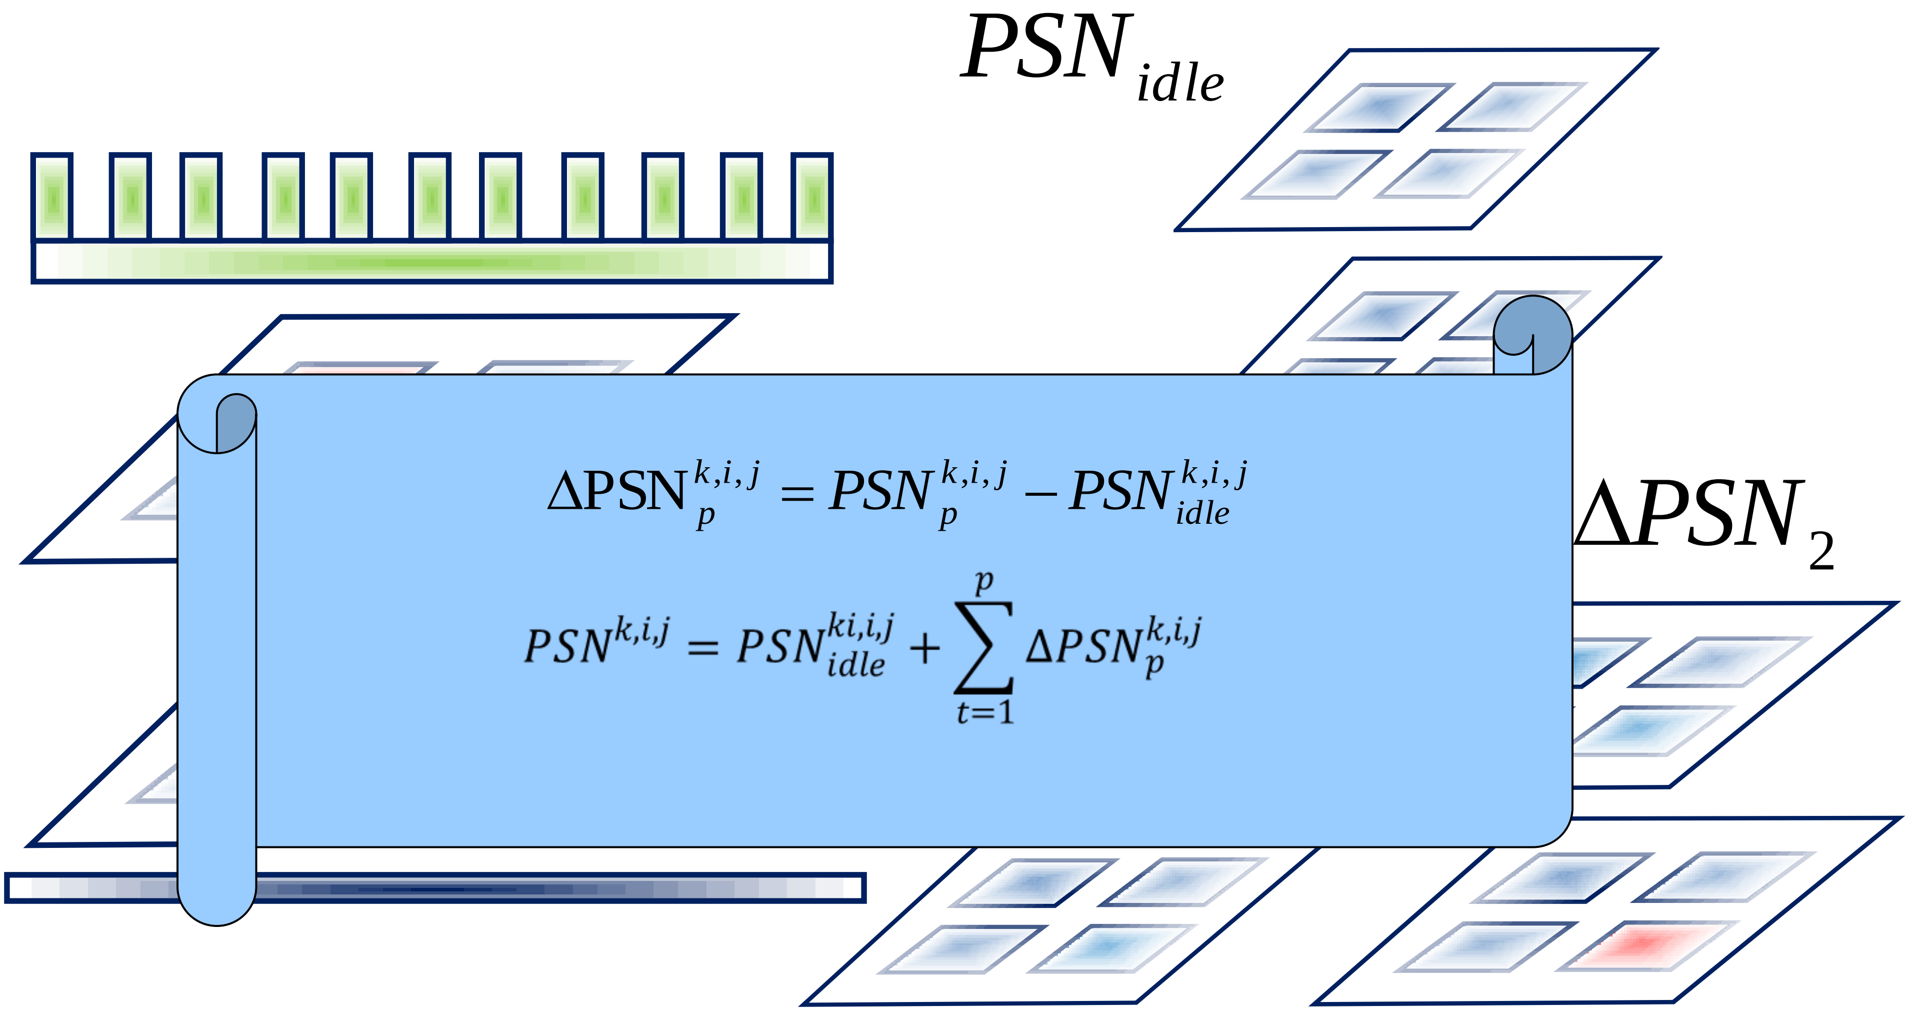

$$\Delta \mathrm{PSN}_{p}^{k,i,j} = PSN_{p}^{k,i,j} - PSN_{idle}^{k,i,j}$$

$$PSN^{k,i,j} = PSN_{idle}^{ki,i,j} + \sum_{t=1}^{p} \Delta PSN_{p}^{k,i,j}$$

Power Supply Noise Estimation Method (Cont.)

- Procedure
  - Store individual workload PSN impact into a table for online reference
  - Estimate relative PSN magnitudes based on individual workload analysis

$$PSN^{k,i,j} = PSN_{idle}^{ki,i,j} + \sum_{t=1}^{p} \Delta PSN_{p}^{k,i,j}$$

- Advantage
  - Easy to compute for online PSN estimation

Validation

- One representative test case

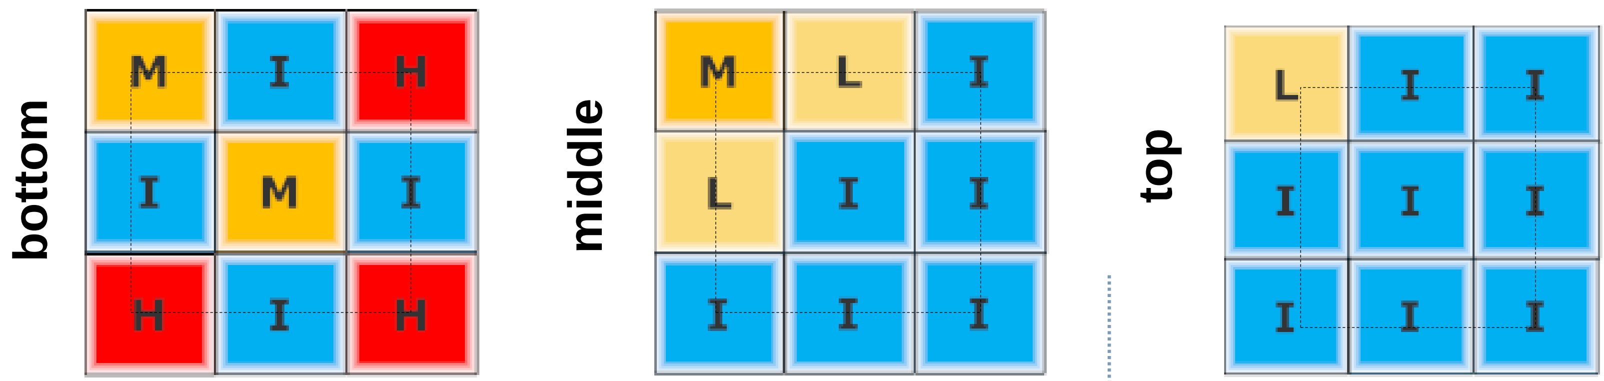

- Workload profiles

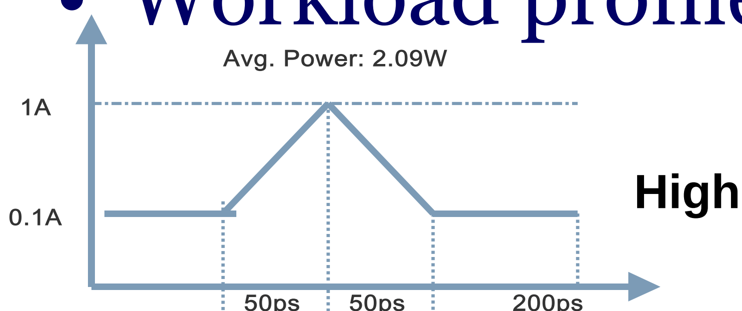

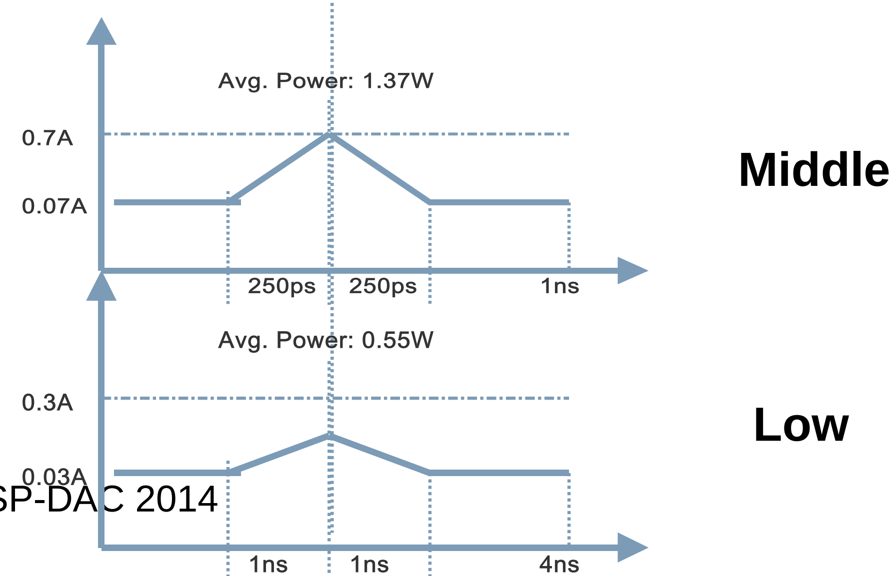

Validation (Cont.)

- PSN distribution comparisons

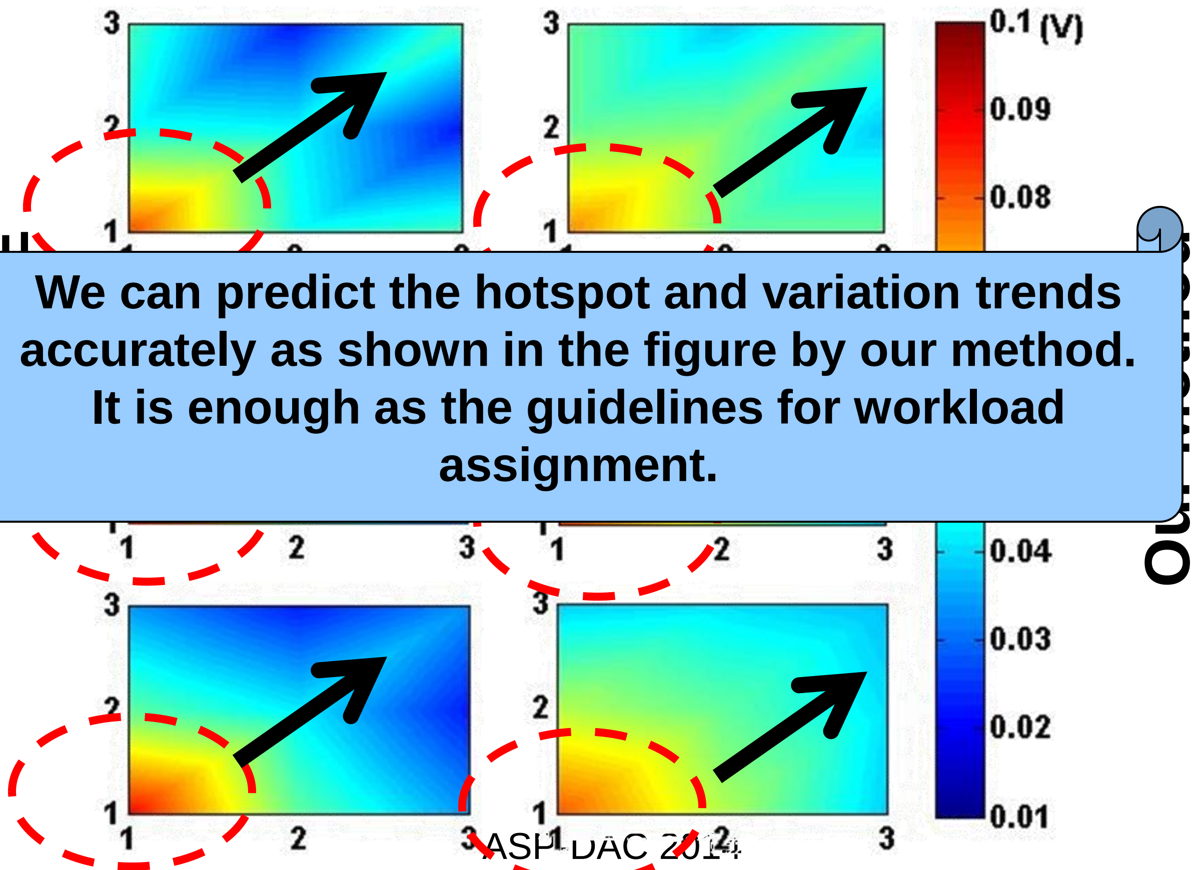

We can predict the hotspot and variation trends accurately as shown in the figure by our method. It is enough as the guidelines for workload assignment.

The Proposed Heuristic Algorithm

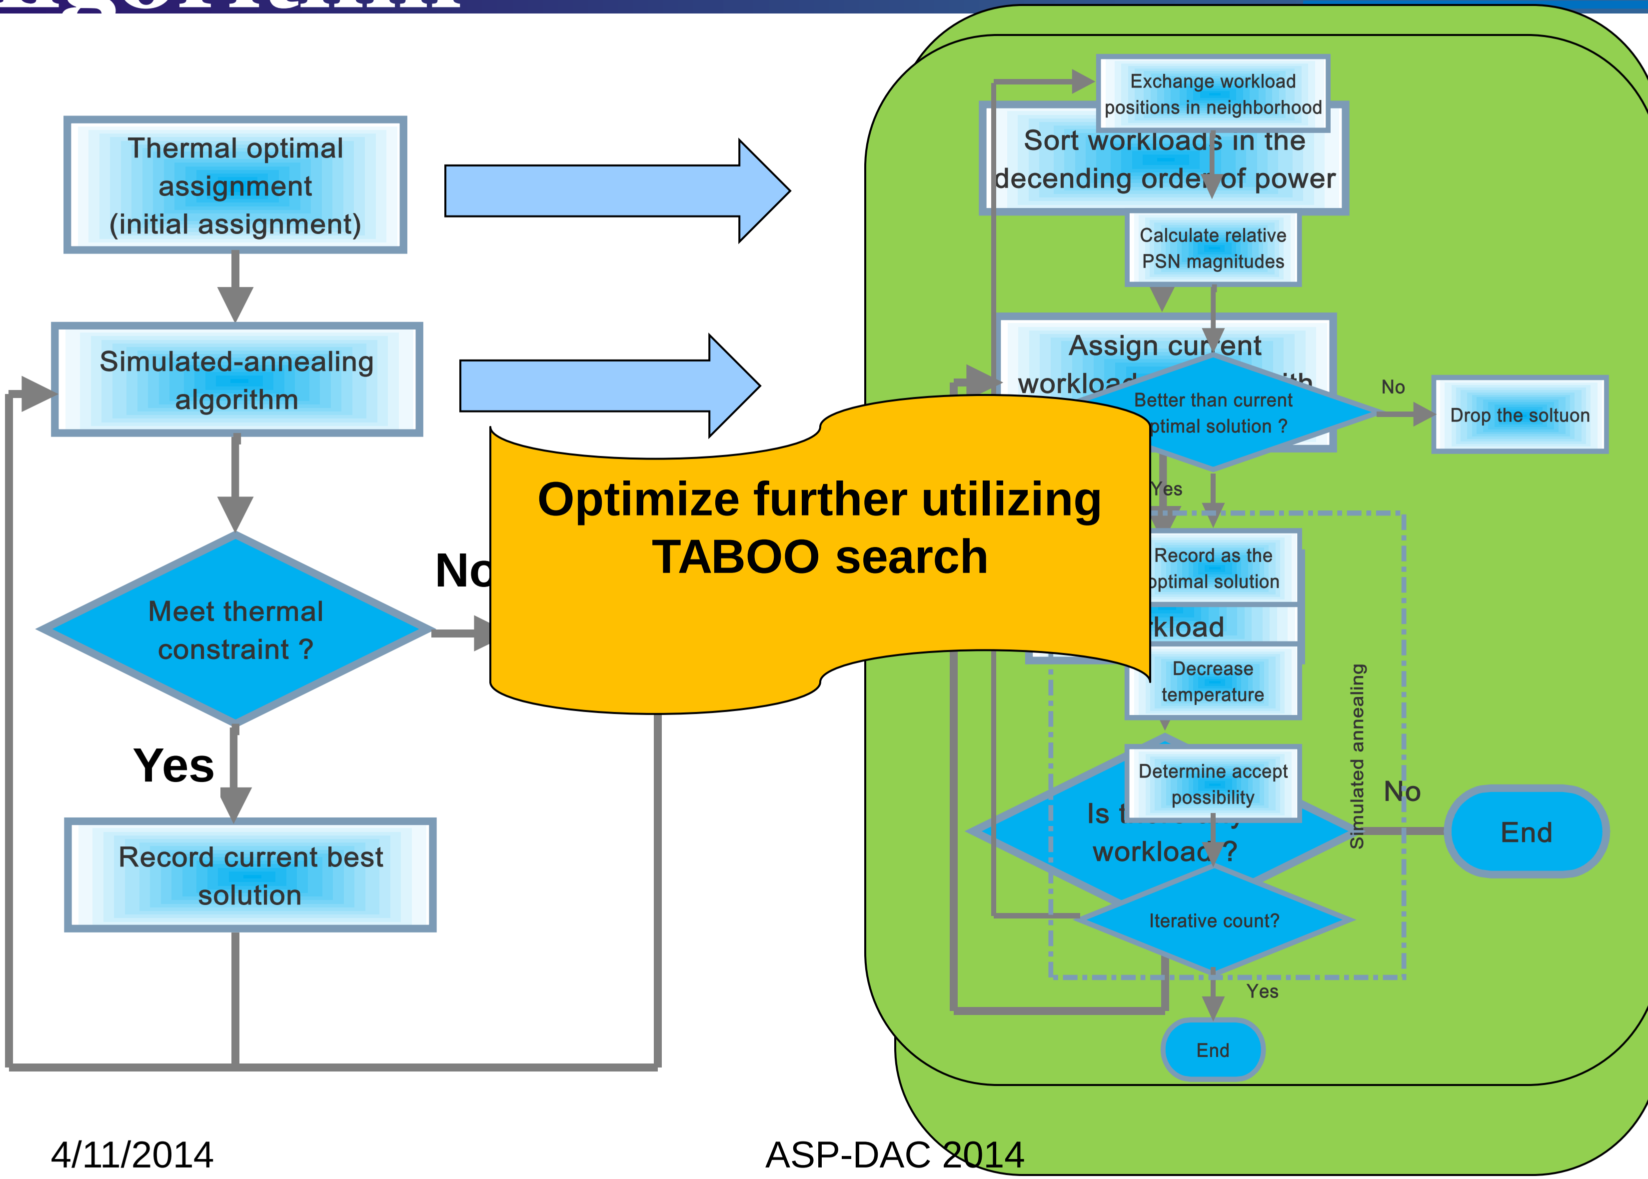

Outline

- Introduction
- Preliminaries
- Motivation through An Example
- Problem formulation & proposed method
- Experimental Results
- Conclusions

Experiment Setup

- Target platform
  - 3 ×3 ×3 homogeneous 3D MPSoC
- Workload classification

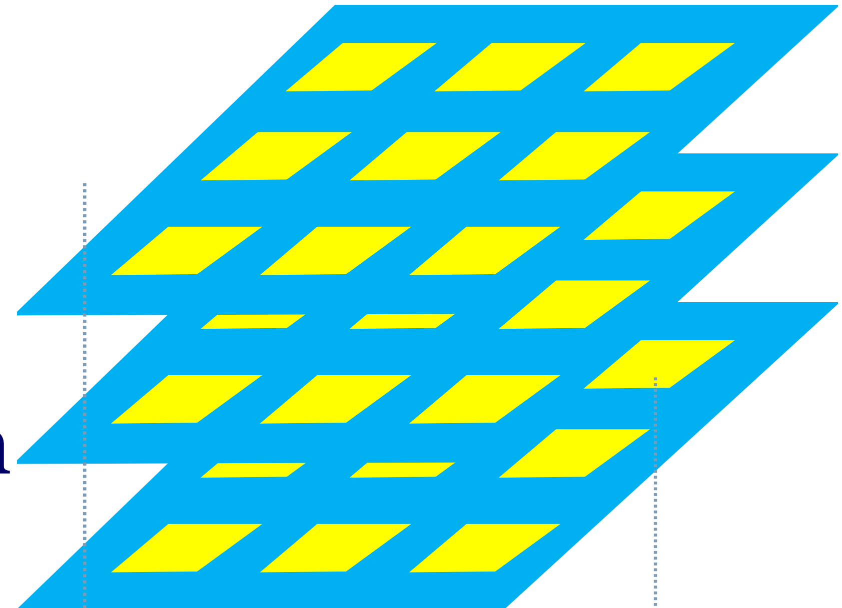

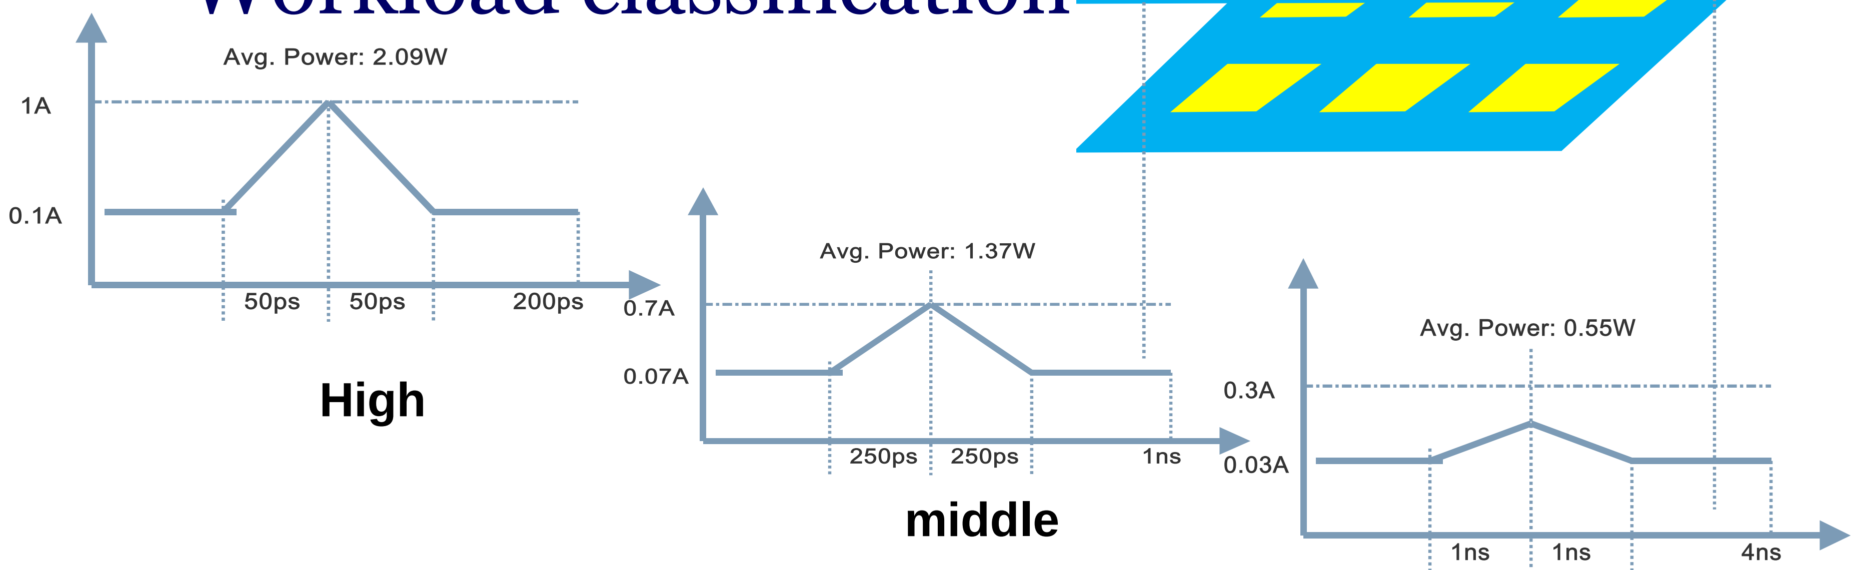

Experimental Results
Case I

- Case I: 3 high, 3 middle 3 low activity workloads

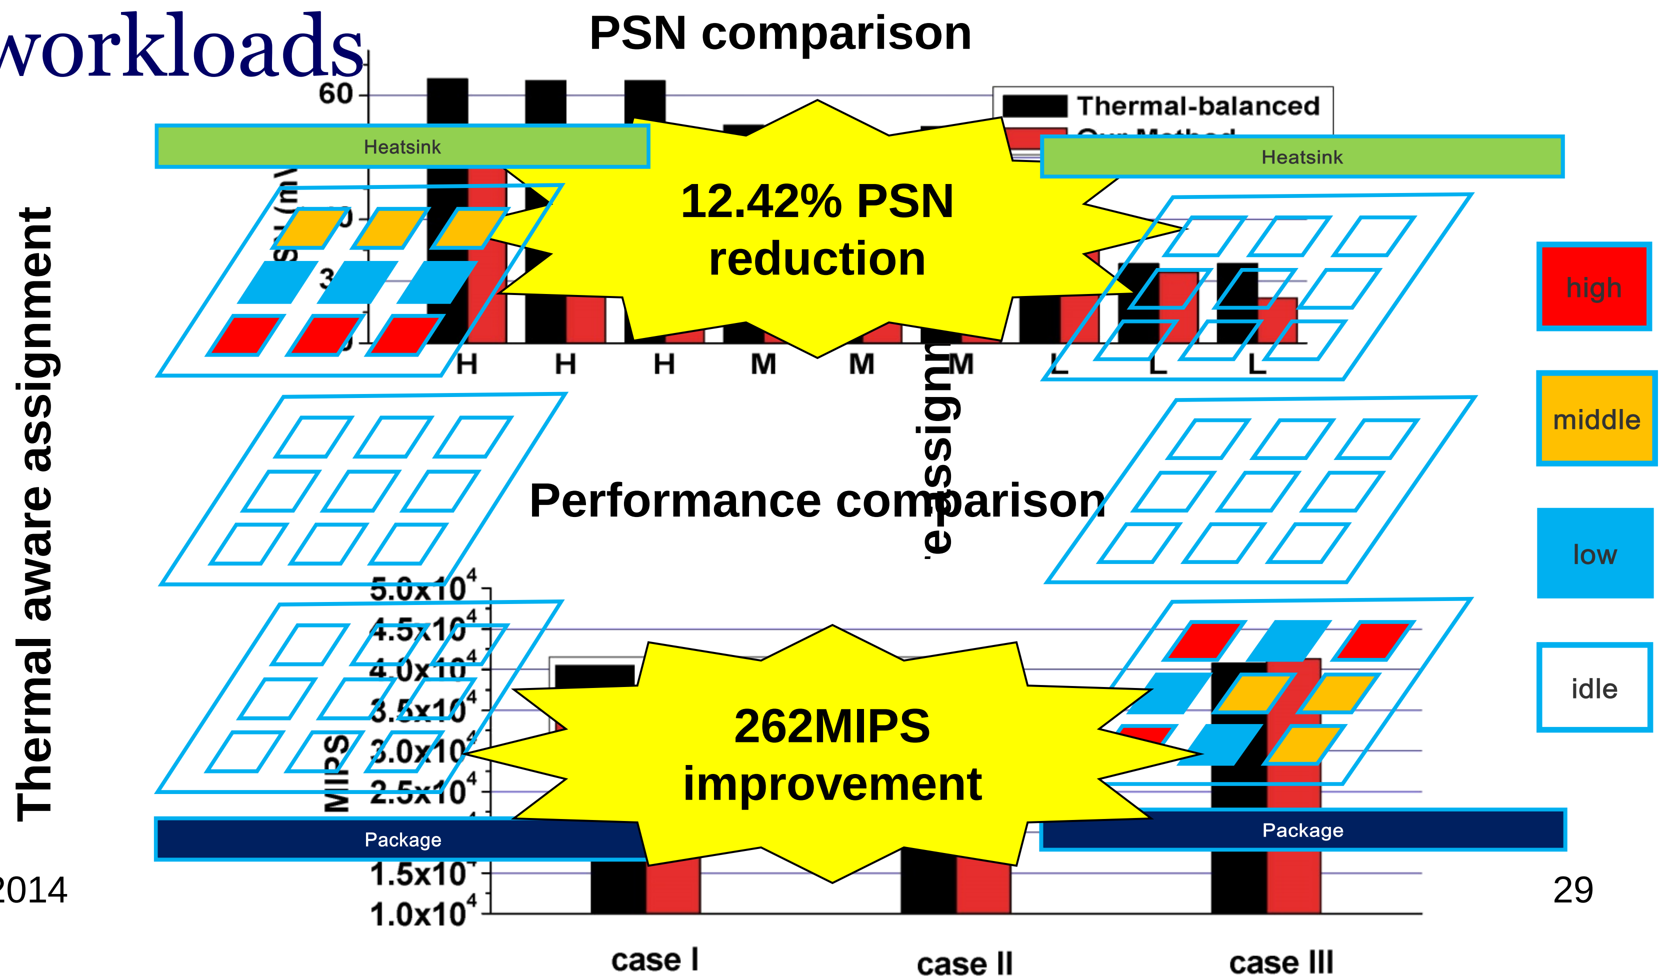

Experimental Results Case II

- Case II
  - 5 high, 2 middle & 2 low activity workloads
  - Similar results as case I

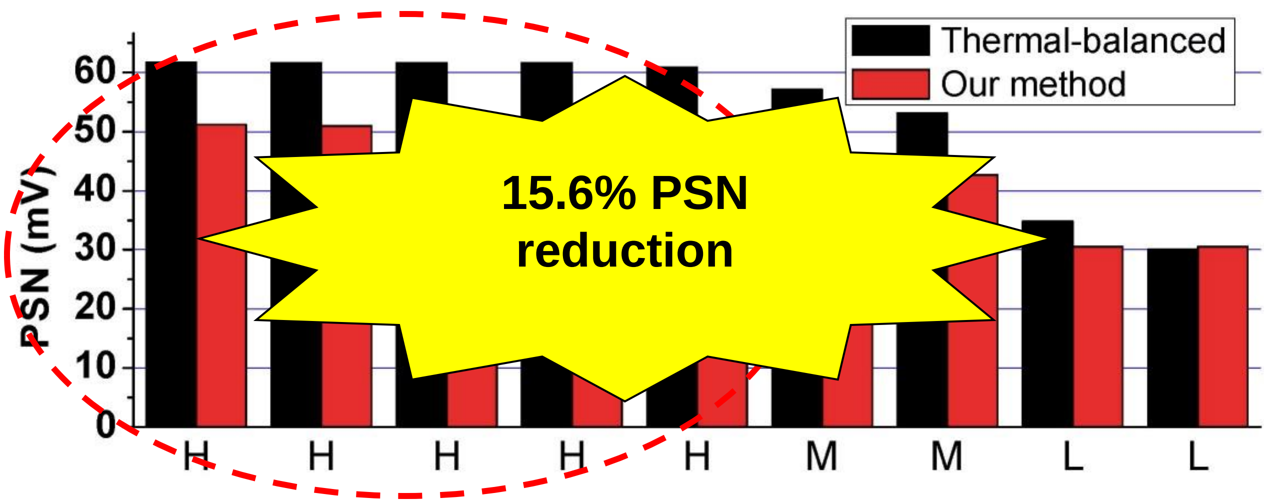

Experimental Results
Case III

- Case III
  - 8 high, 4 middle & 4 low activity workloads

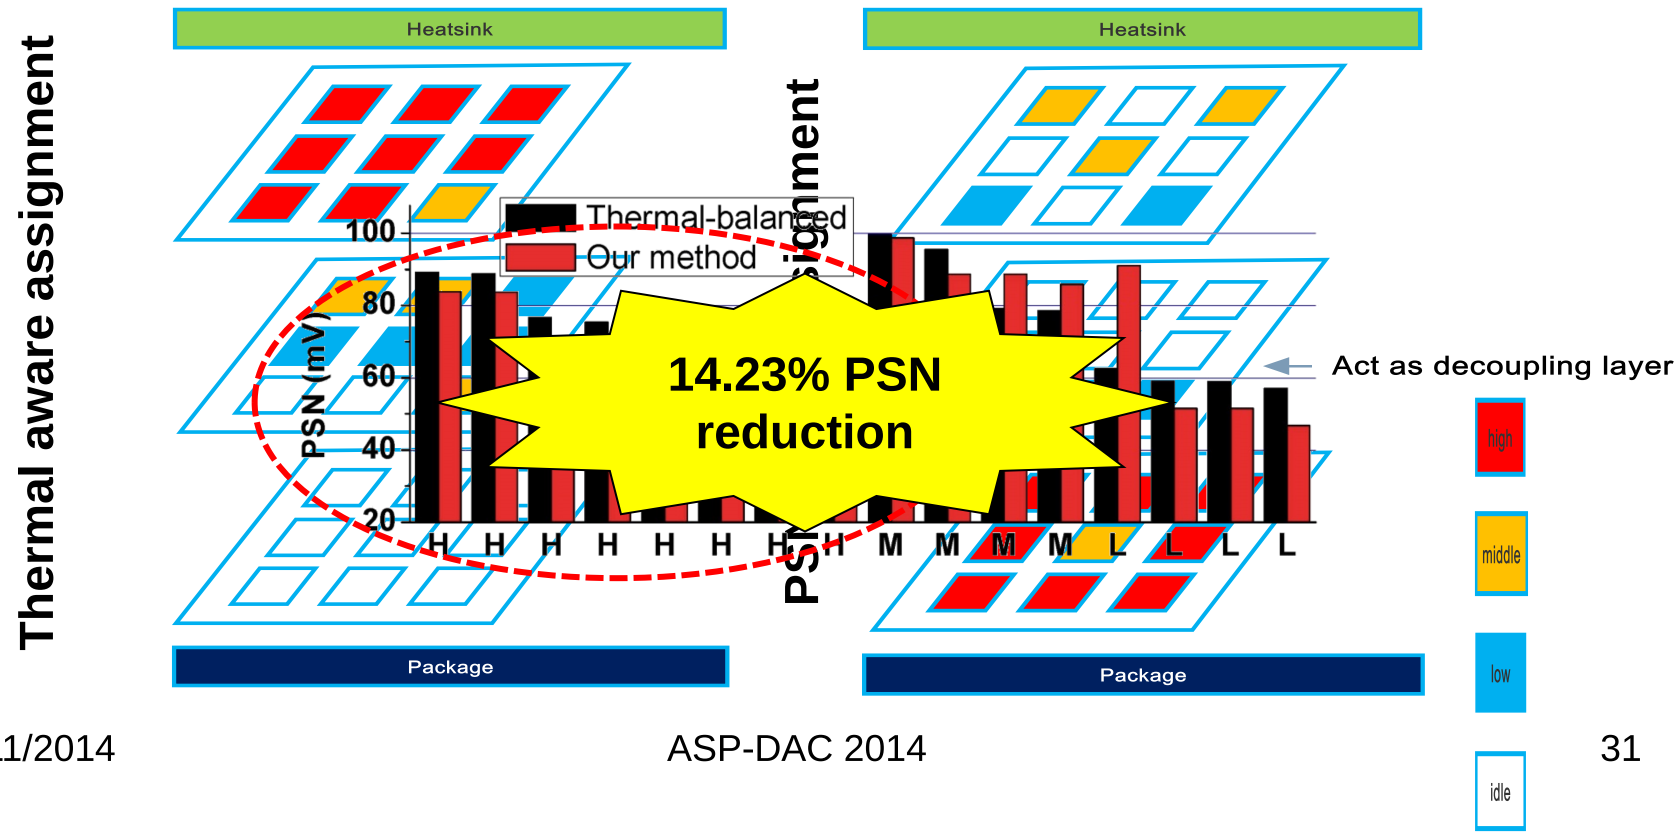

Outline

- Introduction
- Preliminaries
- Motivation through An Example
- Problem formulation & proposed method
- Experimental Results
- Conclusions

Conclusions

- We explored the PSN minimization problem with the thermal constraint from the workload assignment angle
- We proposed an efficient PSN estimation method to guide the workload assignment effectively
- Experimental results indicate that the proposed method can reduce PSN significantly without violating thermal constraint

Thanks for your attention!