Temperature-Constrained Power Management Scheme for 3D MPSoC

Arnica Aggarwal, Sumeet S. Kumar, Amir Zjajo, Rene van Leuken

Delft University of Technology
The Netherlands.

14th May 2012

INTRODUCTION

- Better system integration and performance by scaling down device size
- Further scaling is becoming challenging due to limitations of:

  - interconnect performance - leakage power - process variation

- 3D integration offers increased system integration and performance

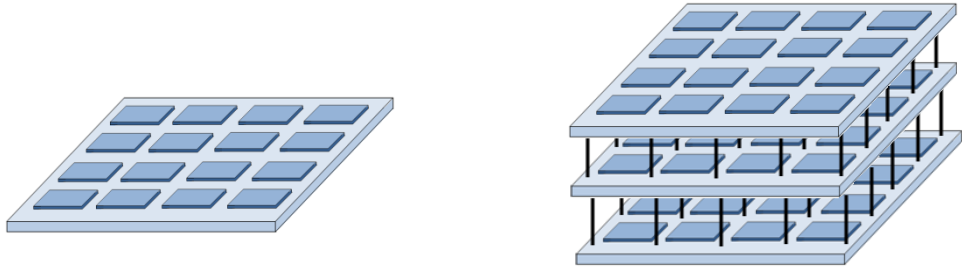

- Increased power density

MOTIVATION

- Temperature constraints for power management is important
- DVFS is a widely used power management technique

2D MPSoC	3D MPSoC
• Monitor PE workload ***only***	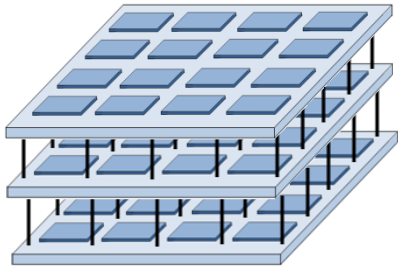
• PE temp. monitored ***independently***	

***PE**: Processing Element*

MOTIVATION

- Temperature constraints for power management is important
- DVFS is a widely used power management technique

2D MPSoC	3D MPSoC
• Monitor PE workload ***only***	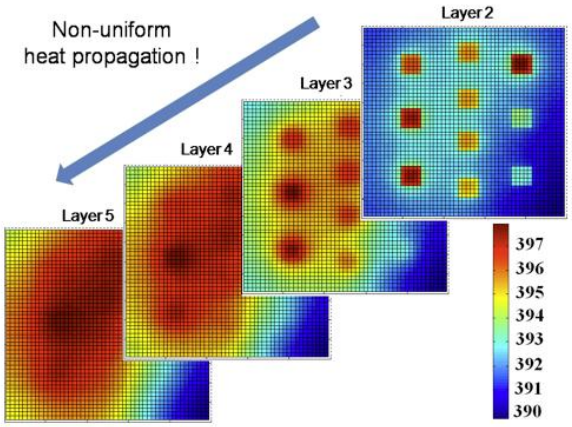
• PE temp. monitored ***independently***	

Need for new approaches to deal with the differences

***PE**: Processing Element*

OUTLINE

- System modeling
- Power management algorithm
- Experimental setup
- Results
- Conclusions

SYSTEM MODELING (I)

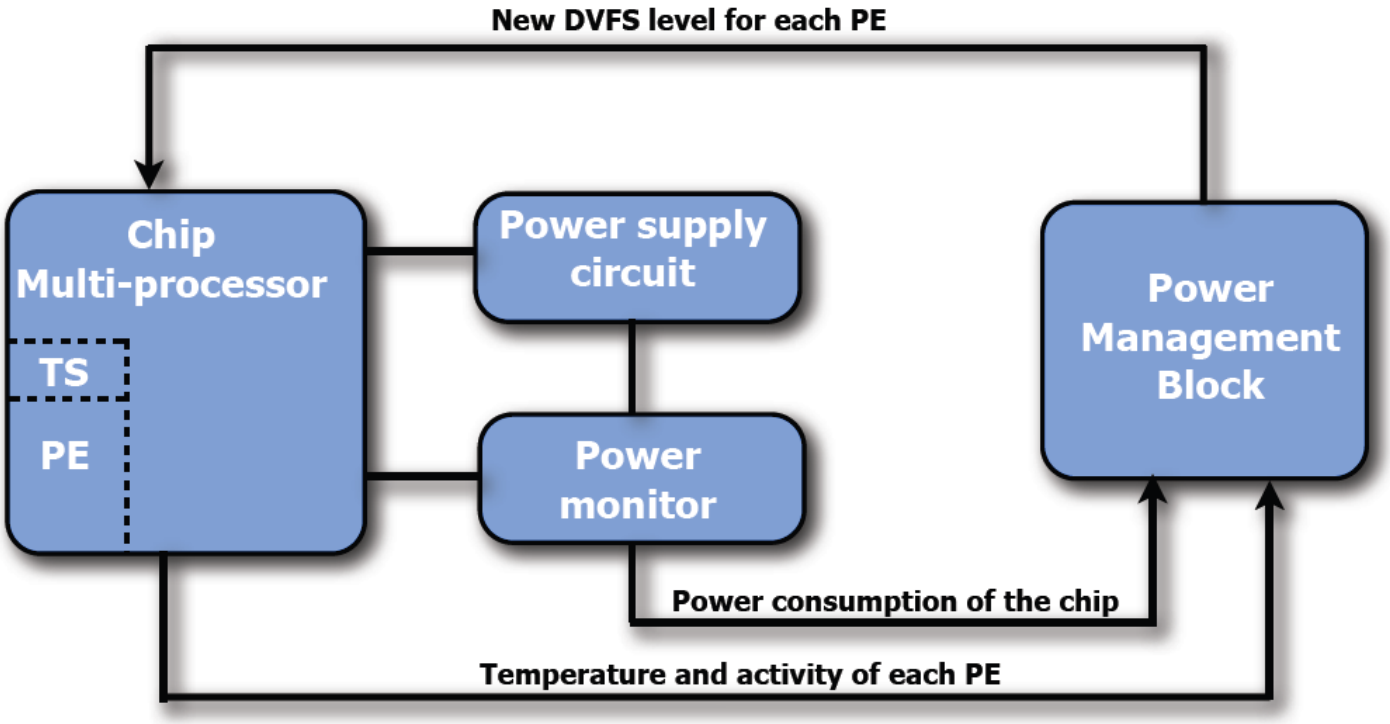

***TS**: Temperature Sensor*
***PE**: Processing Element*

SYSTEM MODELING (II)

- Each PE has fixed V/F levels: VF1, VF2, VF3, VF4, VF5, VF6
- Relation between Power and DVFS level can be given by:

$$P = \boldsymbol{A}.(V^2F) + \boldsymbol{B}$$

$$\Rightarrow \quad \boldsymbol{\Delta P = A \,.\, \Delta(V^2F)}$$

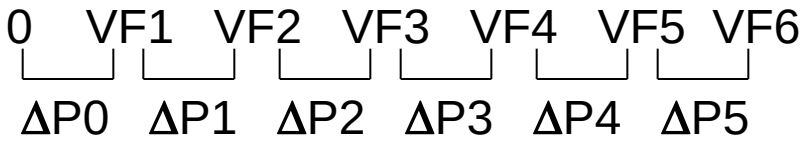

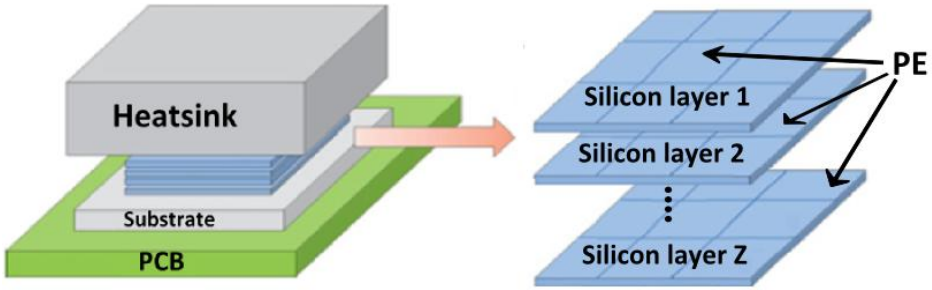

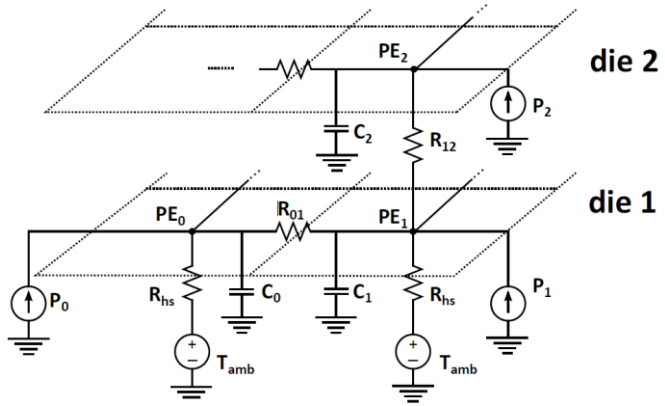

Thermal Model Equations

$$T_1 = T_{amb} + (x.P_1 + P_2) * R_{hs}$$
$$T_2 = T_1 + (-y.P_1 + P_2) * R_{hs}$$

x+y+… = 1

Direct relation between temperature and power

$$\Delta T_i = R_{eff\_ij} . \Delta P_j$$

$$R_{eff} = \begin{pmatrix} R_{11} & R_{12} & R_{13} & R_{14} \\ R_{21} & R_{22} & R_{23} & R_{24} \\ R_{31} & R_{32} & R_{33} & R_{34} \\ R_{41} & R_{42} & R_{43} & R_{44} \end{pmatrix}$$

$$\Delta P = A . \Delta(V^2F) \quad => \quad \Delta T = R_{eff} * A. \Delta(V^2F)$$

POWER MANAGEMENT ALGORITHM

New approach : Proposed approach

2D approach : approach used in 2D MPSoC

PE temp. monitored independently

POWER MANAGEMENT ALGORITHM

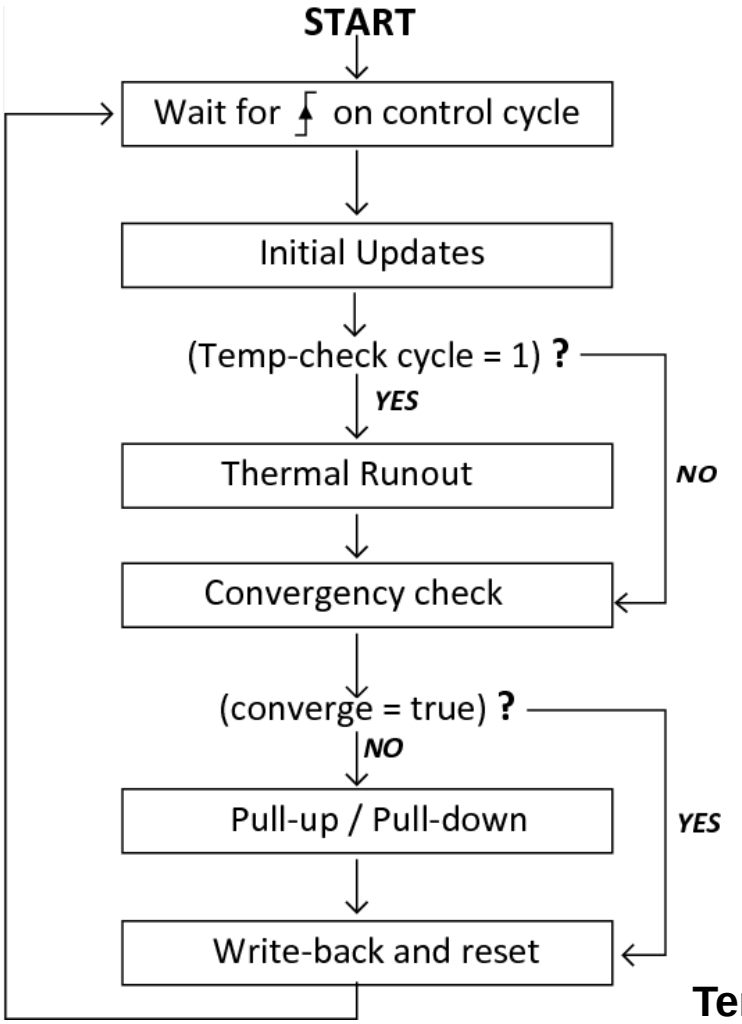

Control period: PMB decides new V-F levels
Temp-check Period: current temp. available to PMB

POWER MANAGEMENT ALGORITHM

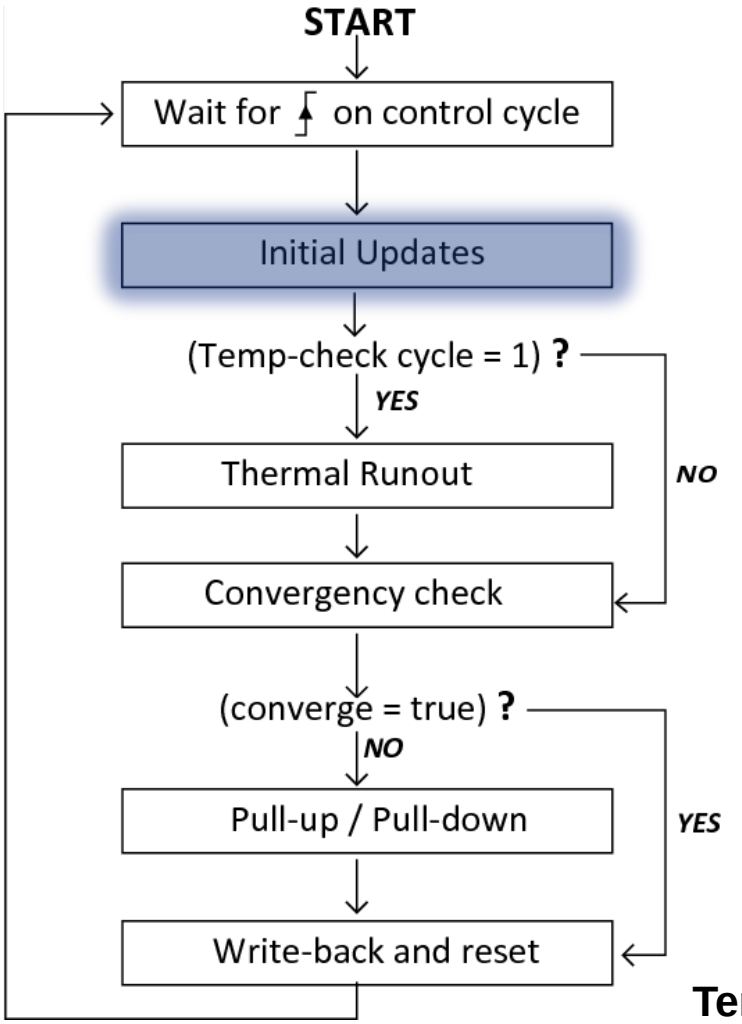

Control period: PMB decides new V-F levels
Temp-check Period: current temp. available to PMB

POWER MANAGEMENT ALGORITHM

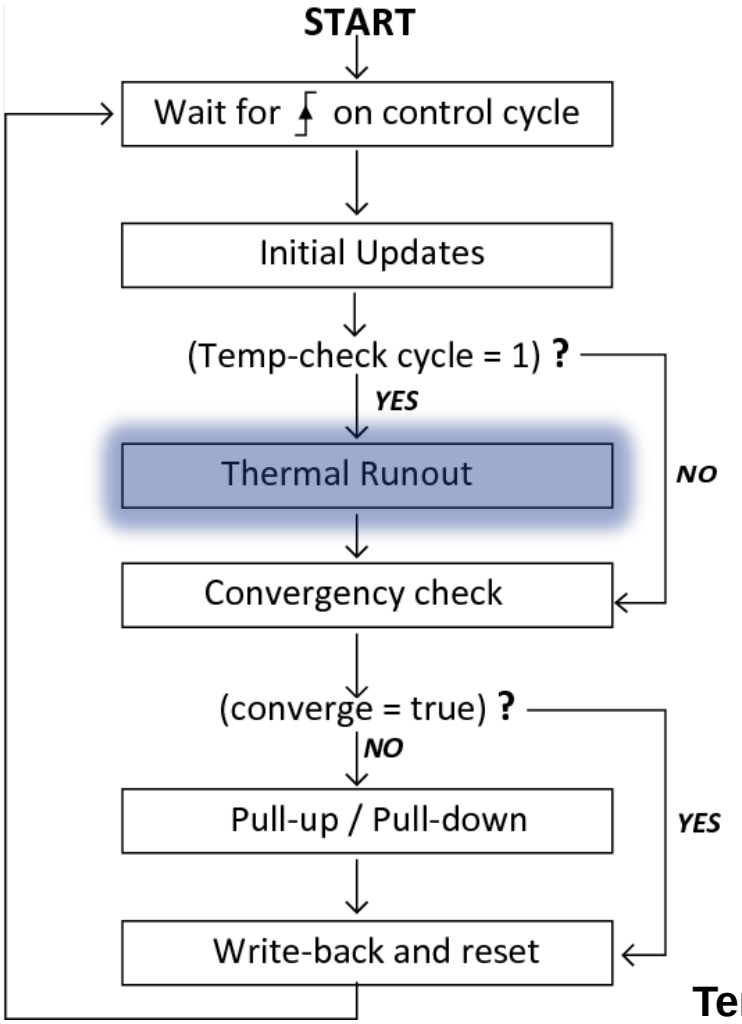

Control period: PMB decides new V-F levels
Temp-check Period: current temp. available to PMB

POWER MANAGEMENT ALGO. (II)

Thermal Runout

$$\boldsymbol{R}_{\text{eff}} = \begin{pmatrix} R_{11} & R_{12} & R_{13} & R_{14} \\ R_{21} & R_{22} & R_{23} & R_{24} \\ R_{31} & R_{32} & R_{33} & R_{34} \\ R_{41} & R_{42} & R_{43} & R_{44} \end{pmatrix}$$

- If temperature crosses safety limit (e.g. 328 K for critical temp of 330 K)
  - Operating V-F levels are adjusted
- 2D approach: V-F of victim PE is scaled down
- New approach:

  Weight of each PE = ***(a** * (1-utilization)) + (**b** * R_{eff})*

(the row of R_{eff} denoting effect on victim is considered)

Utilization : PE activity in previous control period

POWER MANAGEMENT ALGO. (III)

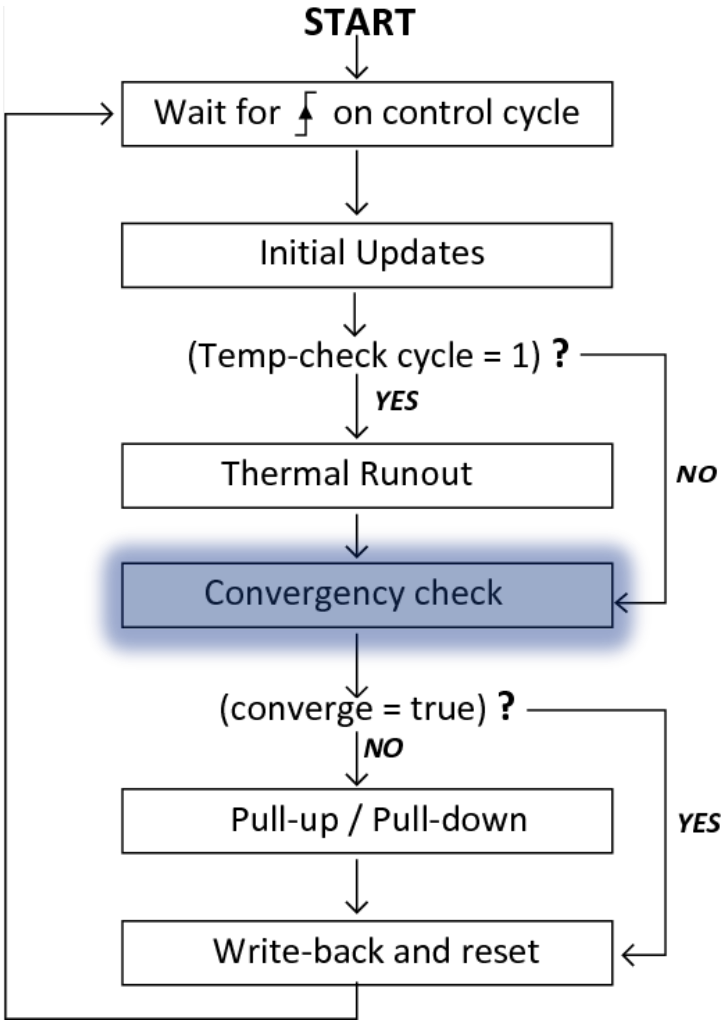

Power budget: Constrains peak power of the chip

POWER MANAGEMENT ALGO. (III)

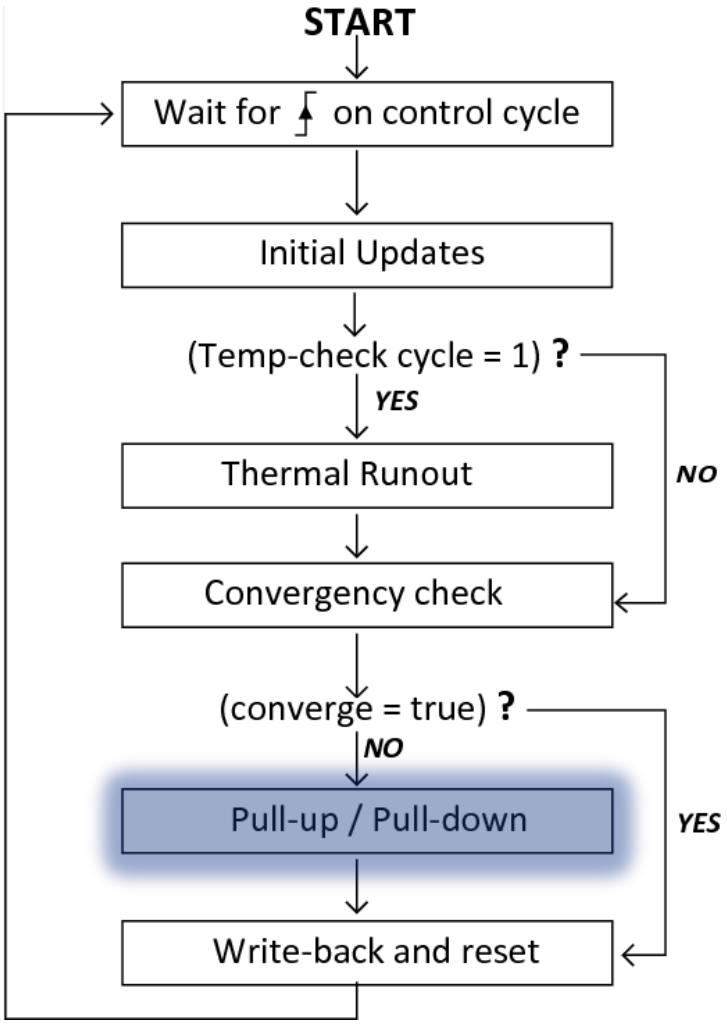

Power budget: Constrains peak power of the chip

POWER MANAGEMENT ALGO. (IV)

pull-up / pull-down

- 2D approach: V-F scaled according PE utilization ***ONLY***
- New approach: V-F scaled according to a weighted equation

  *(**c** * utilization) + (**d** * temp margin) + (**e** * height) + (**f** * area)*

- V-F of a PE is scaled up only if
  - (***T + ΔT)*** does not cross safety limit
  - Was not scaled down in ***Thermal Runout*** stage

POWER MANAGEMENT ALGO. (V)

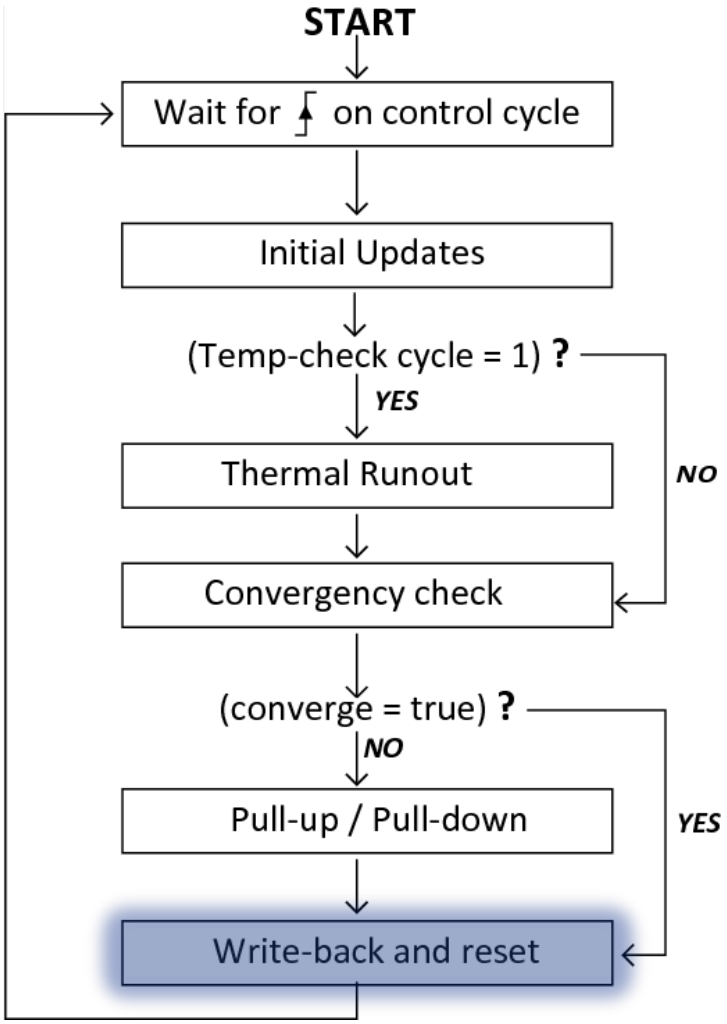

EXPERIMENTAL SETUP (I)

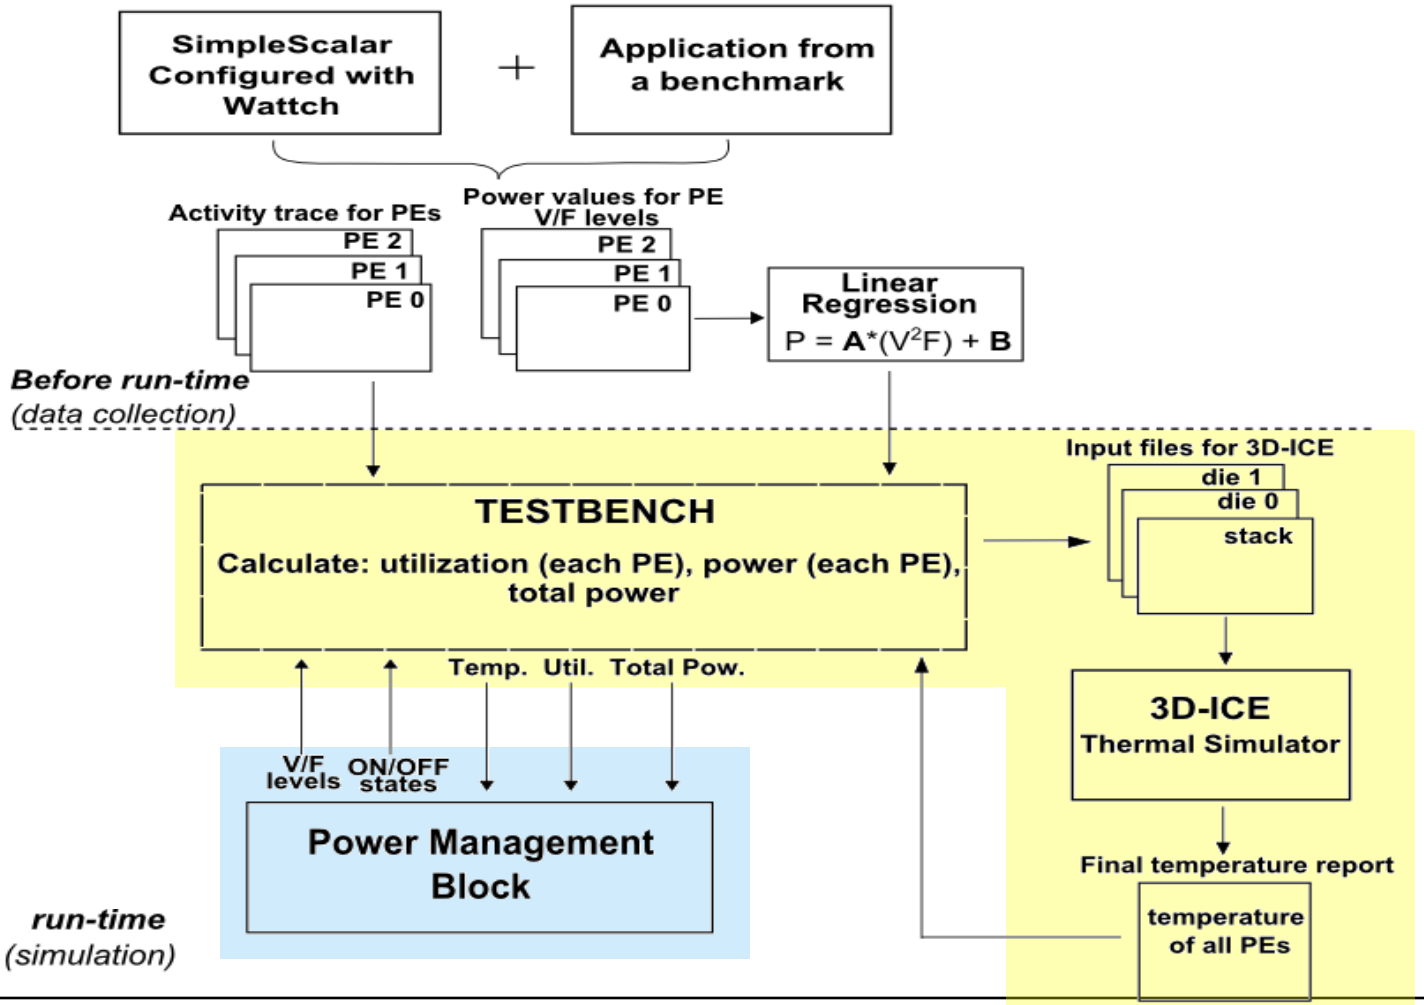

EXPERIMENTAL SETUP (II)

- Control period : 50us (60,000 cycles at max freq)
- Temperature input : every 1ms
- Temperature constraint : 320 K on each PE (strict)
- Safety margin of : 2K

- 6 DVFS levels per PE are taken
  - Frequency: 700 – 1200MHz
  - Voltage : 0.8V – 1.1V

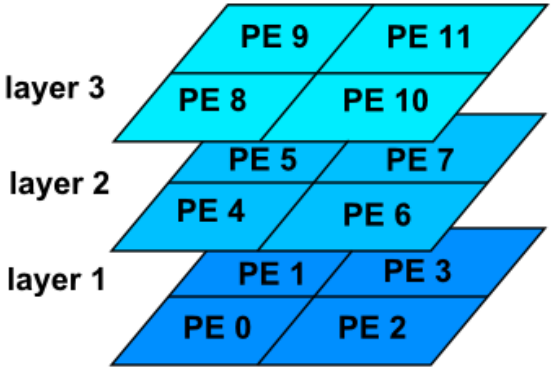

- Deep sleep mode: ***only*** clock is gated
- Basicmath application from MiBench Benchmark
- Two sets of simulation setup

RESULTS (I)

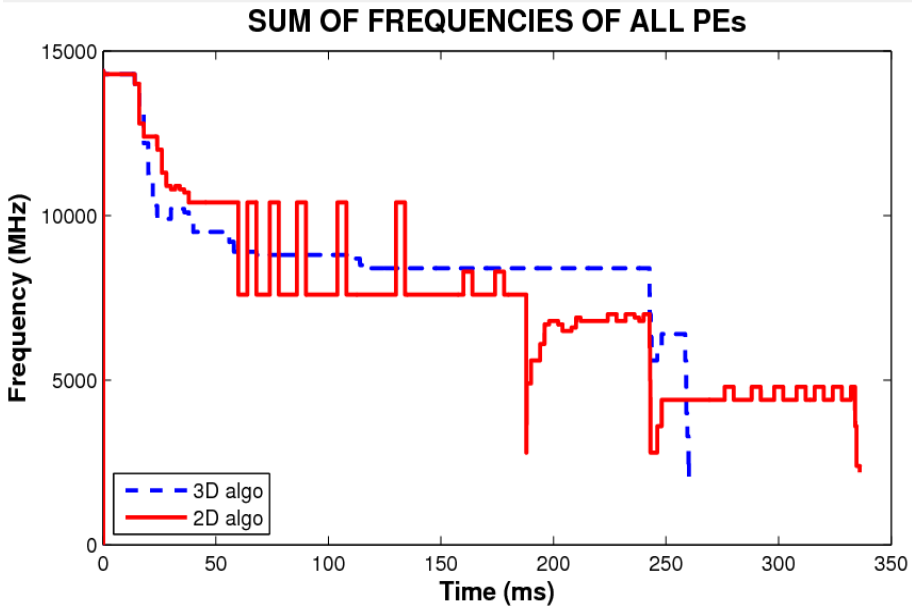
SUM OF FREQUENCIES OF ALL PEs
Frequency (MHz)
Time (ms)
3D algo
2D algo
0
5000
10000
15000
0
50
100
150
200
250
300
350

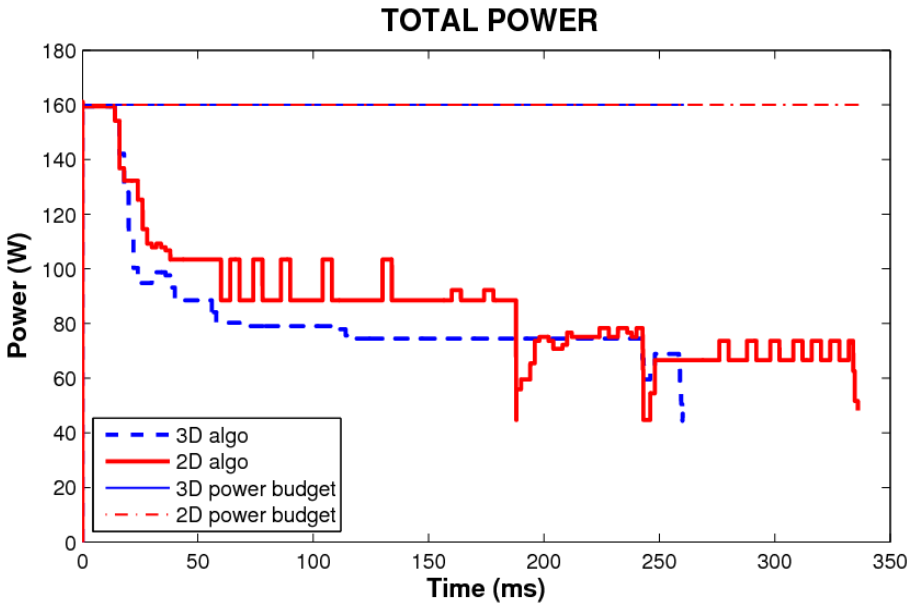
TOTAL POWER
Power (W)
Time (ms)
3D algo
2D algo
3D power budget
2D power budget
0
20
40
60
80
100
120
140
160
180
0
50
100
150
200
250
300
350

RESULTS (II)

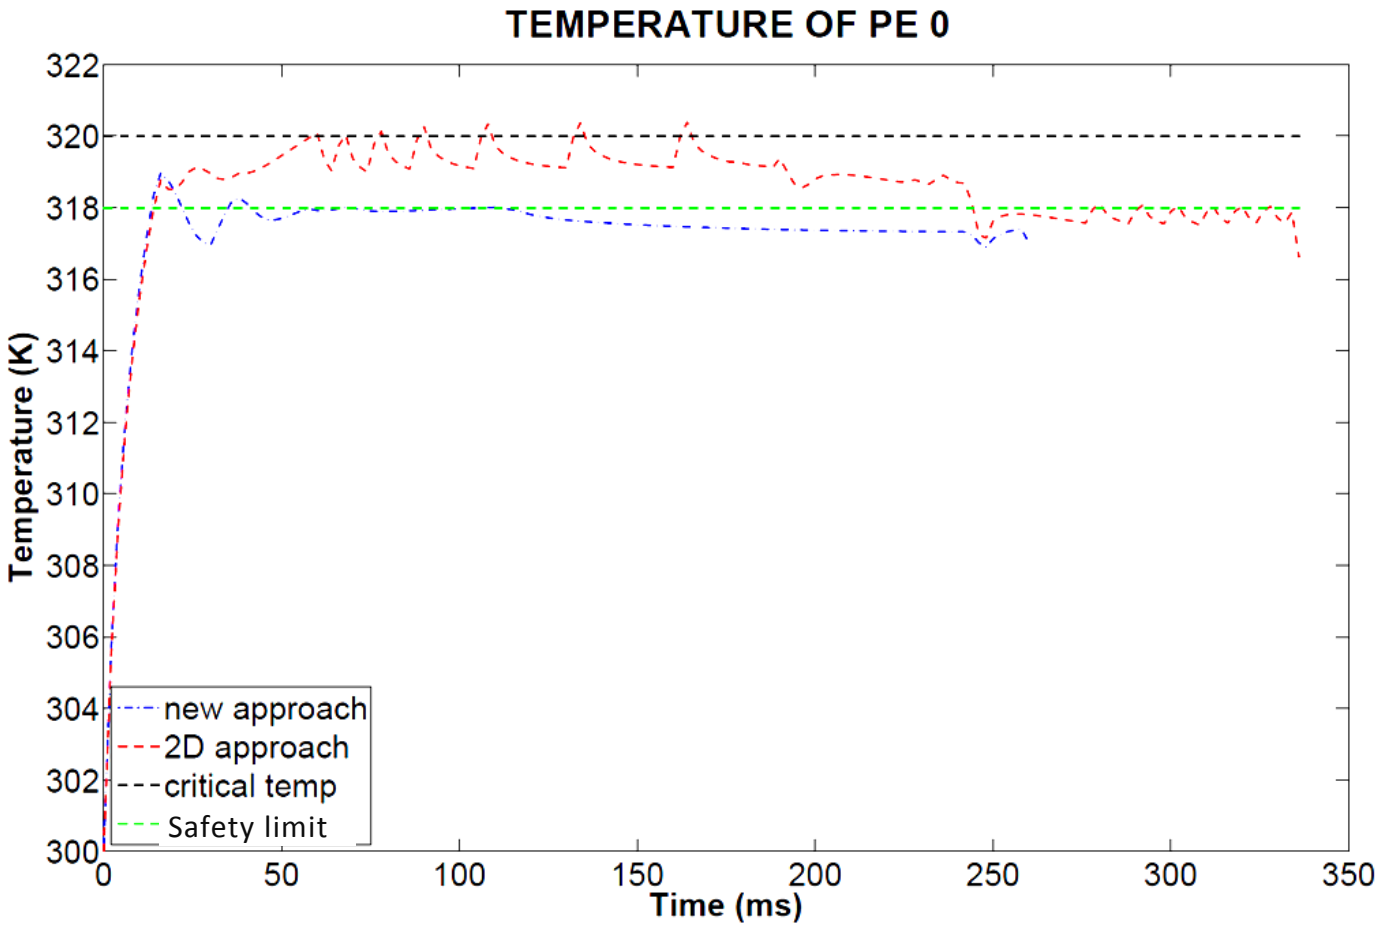
TEMPERATURE OF PE 0
Temperature (K)
Time (ms)
322
320
318
316
314
312
310
308
306
304
302
300
0
50
100
150
200
250
300
350
new approach
2D approach
critical temp
Safety limit

RESULTS (III)

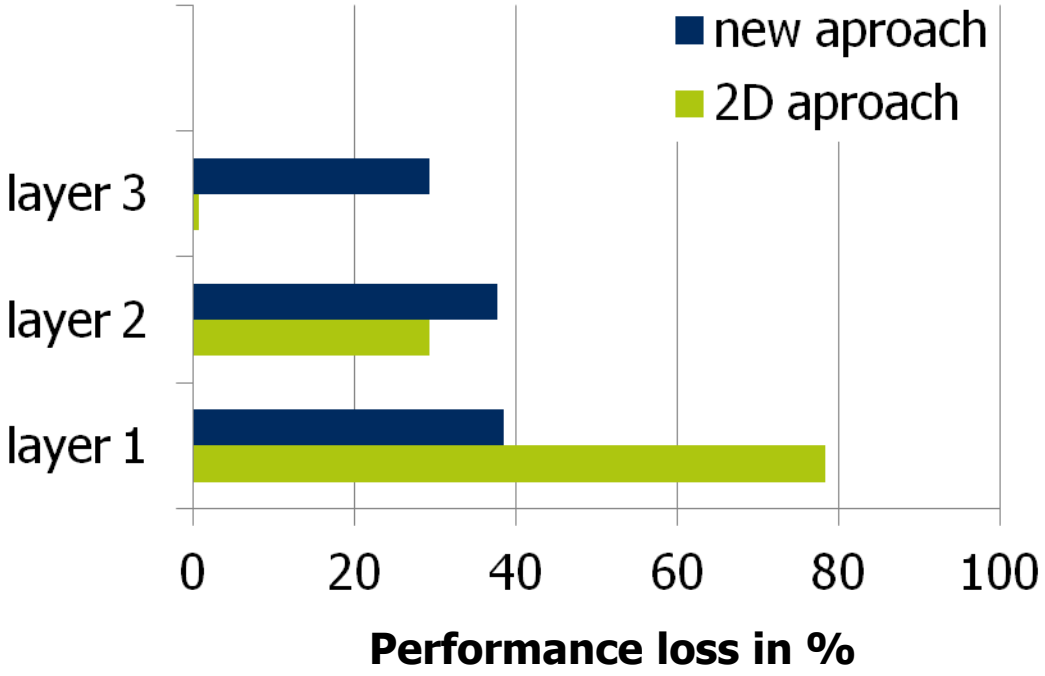

- Similar losses on all layers
- Similar execution time
- Total execution time reduces

19.55% reduction in total execution time

CONCLUSIONS

- A new power management scheme was proposed
- High level thermal model was achieved
- Weighted equations were used for choosing V-F levels
- Power and temperature were maintained below constraints
- Less temperature fluctuations
- Up to 19.55% improvement in total execution time
    - High Sum of Frequencies
    - Turning OFF of PEs was avoided

Temperature-Constrained Power Management Scheme for 3D MPSoC

Arnica Aggarwal, Sumeet S. Kumar, Amir Zjajo, Rene van Leuken

Delft University of Technology
The Netherlands.

14th May 2012