Dynamic Resource Allocation for Many-Core SoCs

Axel Jantsch

TU Wien, Vienna, Austria

Huawei Sweden 5G IC Design Trend Workshop Stockholm, Sweden

26 October 2017

Nikil Dutt, Mohammad-Hashem Haghbayan, Axel Jantsch, Anil Kanduri, Pasi Liljeberg, Antonio Miele, Amir-Mohammad Rahmani, Santanu Sarma, Hannu Tenhunen

TU Wien, Vienna, Austria University of Turku, Finland University of California, Irvine, USA

The Problem

- Large number of resources
- Many tight constraints
- Varying application demands, both within and between applications;
- Functional Aberrations:
 - Design errors or omissions;
 - Malicious attacks;
 - Aging;
 - Soft errors;
- Non-functional Aberrations:
 - Performance;
 - Power consumption;





Santanu Sarma et al. "On-Chip Self-Awareness Using Cyberphysical-Systems-On-Chip (CPSoC)". . In: Proceedings of the 12th International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS). New Delhi, India, Oct. 2014













Autonomy and Adaptivity

Autonomy is the ability to operate independently, without external control.

Adaptivity is the ability to effect run-time changes and handle unexpected events.

Goals for Dynamic Task Mapping



Performance Driven Throughput Driven Lifetime Reliability Driven

Dynamic Task Mapping



MapPro Objectives:

- Maximize performance for all applications;
- Minimize communication latency in the new application;
- Minimize fragmentation.

Mohammad-Hashem Haghbayan et al. "MapPro: Proactive Runtime Mapping for Dynamic Workloads by Quantifying Ripple Effect of Applications on Networks-on-Chip". In: *Proceedings of the International Symposium on Networks on Chip.* Vancouver, Canada, Sept. 2015



























MapPro: Heuristic to minimize application internal communication delay and to minimize fragmentation.

- 1. First Node selection: Identifies a first node and a region for a new application;
- 2. Allocates specific cores around the first node;
- 3. Maps tasks to cores.



Experiments with 12x12 - 16x16 networks.

AWMD: Average Weighed Manhattan Distance: Measures the communication cost based on traffic volume.

NMRD: Normalized Mapped Region Dispersion is the normalized average of pairwise Manhattan distances of all communication nodes of a mapped application: measures the compactness of a region.

Example 2: Power- and Thermal Constrained Task Mapping



The patterning algorithm disperses mapped cores to maximize the Thermal Safe Power budget.

Anil Kanduri et al. "Dark Silicon Aware Runtime Mapping for Many-core Systems: A Patterning Approach". In: Proceedings of the International Conference on Computer Design (ICCD). New York City, USA, Oct. 2015, pp. 610–617

Example 2: Dark Silicon



Silicon Melting Chip Malfunctioning Unreliability and Ageing

Example 2: Thermal Design Power



Design Time Estimate

Example 2: Fixed Power Budget



Steady State Chip

Working Chip

Conservative Budgeting

Dark Silicon

Example 2: Variable Power Budget



Example 2: Efficient Budgeting

Tightly packed Cores

55°	55°	54°	50°	
58°	58°	57°	56°	
78°	79°	77°	57°	
(12.7)	(12.7)	(12.7)		
79°	80°	78°	58°	
(12.7)	(12.7)	(12.7)		

Neighbors accumulating temperature

Utilized Power Budget = 76.2 W

Spreadout Cores

80° (14.6)	60°	59°	79° (14.6)
60°	<mark>79°</mark> (14.6)	60°	60°
59°	60°	61°	80° (14.6)
<mark>79°</mark> (14.6)	60°	80° (14.6)	60°

Neighbors dissipating temperature

Utilized Power Budget = 87.6 W

- ✓ 15% Better Utilization
- ✓ Activate more cores
- ✓ Reduce temperatures
- ✓ Minimize Dark Silicon

Example 2: Implications of Mapping





Application 2

Application 1





Tightly Packed – Greedy First Node



Power Budget = 66W

Spread Out – Adaptive First Node



Power Budget = 74.6W

¢

Example 2: Power Budget Improvement

Percentage Power Budget Improvement for PAT over SC

Network Size	90% Dark		75% Dark		50% Dark	
	Avg.	Best	Avg.	Best	Avg.	Best
16x16	5.74	13.9	4.15	11.3	2.19	7.68
20x20	6.54	17.17	5.06	8.55	2.63	4.28

Percentage Power Budget Improvement for PAT over TSP-WC

Network Size	90% Dark		75% Dark		50% Dark	
	Avg.	Best	Avg.	Best	Avg.	Best
16x16	32.33	34.92	22.02	24.14	11.73	13.2
20x20	38.70	40.83	22.40	27.4	12.5	13.33

Example 2: Throughput Gain

Percentage Throughput gain for PAT over SC

Network Size	90% Dark		75% Dark		50% Dark	
	Avg.	Best	Avg.	Best	Avg.	Best
16x16	7.27	15.64	4.59	13.92	2.42	8.58
20x20	8.5	20.99	5.88	10.21	2.89	4.54

✓ Surplus Budget >> Added latency

✓ Minimal congestion

Per Application Latency

Example 3: Lifetime-Reliability-Driven Task Mapping

- To main limitations of future many-cores:
 - Not enough power to turn on all cores (dark silicon)
 - Increased susceptibility of IC to aging and wear-out
- Goal: Introduce lifetime reliability awareness in the runtime resource management layer
 - Guarantee specified level of reliability
 - Satisfy the power budget
 - Optimize performance

M. H. Haghbayan et al. "A lifetime-aware runtime mapping approach for many-core systems in the dark silicon era". In: Design, Automation Test in Europe Conference Exhibition (DATE). Mar. 2016, pp. 854–857

Example 3: Lifetime-Reliability-Driven Task Mapping

Proposed approach based on two feedback controllers



Example 3: Lifetime-Reliability-Driven Task Mapping





MapPro: lifetime=5.52 years Reliability aware mapping: lifetime=12 years

The plots show the reliability of cores at the end of the system's lifetime. The end of the system's life is reached when the reliability of one core drops below 30%.

Challenges in Complex Many-Core SoCs

- A number and variety of objectives
 - Partially contradicting
 - At different time scales
- Objectives change over time
- The system state has to be known
- Application objectives have to be known

System, Know Thyself



Unknown artist. - Lessing Photo Archive: http://www.lessing-photo.com/p3/110103/11010329.jpg

Cross Layer Sensing and Actuation



Nikil Dutt, Amir M. Rahmani, and Axel Jantsch. "Empowering Autonomy through Self-awareness in MPSoCs". In: Proceedings of the IEEE NEWCAS Conference. Strasbourg, France, June 2017

Axel Jantsch, Nikil Dutt, and Amir M. Rahmani. "Self-Awareness in Systems on Chip – A Survey". In: IEEE Design Test 34.6 (Dec. 2017), pp. 1–19

Hierarchical Goal Mangement



Amir M. Rahmani, Axel Jantsch, and Nikil Dutt. "HDGM: Hierarchical Dynamic Goal Management for Many-Core Resource Allocation". In: IEEE Embedded Systems letters (2017)

Questions ?



References I

Nikil Dutt, Axel Jantsch, and Santanu Sarma. "Towards Smart Embedded Systems: A Self-Aware System-on-Chip Perspective". In: ACM Transactions on Embedded Computing Systems, Special Issue on Innovative Design Methods for Smart Embedded Systems 15.2 (Feb. 2016). invited, pp. 22–27.

Nikil Dutt, Amir M. Rahmani, and Axel Jantsch. "Empowering Autonomy through Self-awareness in MPSoCs". In: *Proceedings of the IEEE NEWCAS Conference*. Strasbourg, France, June 2017.

Mohammad-Hashem Haghbayan et al. "MapPro: Proactive Runtime Mapping for Dynamic Workloads by Quantifying Ripple Effect of Applications on Networks-on-Chip". In: *Proceedings of the International Symposium on Networks on Chip.* Vancouver, Canada, Sept. 2015.

M. H. Haghbayan et al. "A lifetime-aware runtime mapping approach for many-core systems in the dark silicon era". In: *Design, Automation Test in Europe Conference Exhibition (DATE)*. Mar. 2016, pp. 854–857.

Axel Jantsch, Nikil Dutt, and Amir M. Rahmani. "Self-Awareness in Systems on Chip – A Survey". In: *IEEE Design Test* 34.6 (Dec. 2017), pp. 1–19.

References II

.

Axel Jantsch and Kalle Tammemäe. "A Framework of Awareness for Artificial Subjects". In: *Proceedings of the 2014 International Conference on Hardware/Software Codesign and System Synthesis.* CODES '14. New Delhi, India: ACM, 2014, 20:1–20:3.

Anil Kanduri et al. "Dark Silicon Aware Runtime Mapping for Many-core Systems: A Patterning Approach". In: *Proceedings of the International Conference on Computer Design (ICCD)*. New York City, USA, Oct. 2015, pp. 610–617.

Amir M. Rahmani, Axel Jantsch, and Nikil Dutt. "HDGM: Hierarchical Dynamic Goal Management for Many-Core Resource Allocation". In: *IEEE Embedded Systems letters* (2017).

Santanu Sarma et al. "On-Chip Self-Awareness Using Cyberphysical-Systems-On-Chip (CPSoC)". In: Proceedings of the 12th International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS). New Delhi, India, Oct. 2014.