AN INTEGRATED VARIATION-AWARE MAPPING FRAMEWORK FOR FINFET BASED IRREGULAR 2D MPSOCS IN THE DARK SILICON ERA

MS Thesis Final Presentation

PRAMIT RAJKRISHNA

ADVISOR :

Prof. Sudeep Pasricha

COMMITTEE :

Prof. Anura Jayasumana

Prof. Patrick Burns

Colorado State University

Department of Electrical & Computer ENGINEERING

OUTLINE

- INTRODUCTION
- MOTIVATION
- OVERVIEW
- RELATED WORK
- PROBLEM STATEMENT
- METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

INTRODUCTION

- Networks-on-chip (NoC) - an alternative framework to bus based topologies.
- A matrix of interconnected IP cores in a variety of topology configurations.
- Each IP **core** is connected to other cores via a network interface (**NI**) and router (**R**), by bidirectional links.

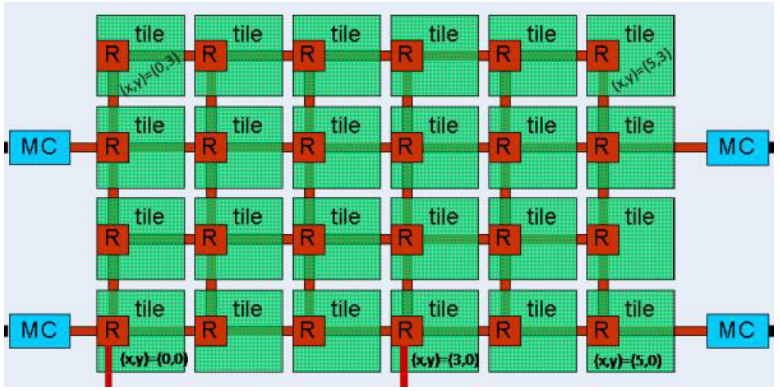

NoC : Intel SCC

NoCs are high performing, scalable, modular and efficient.

Reference : https://www.intel.com

INTRODUCTION

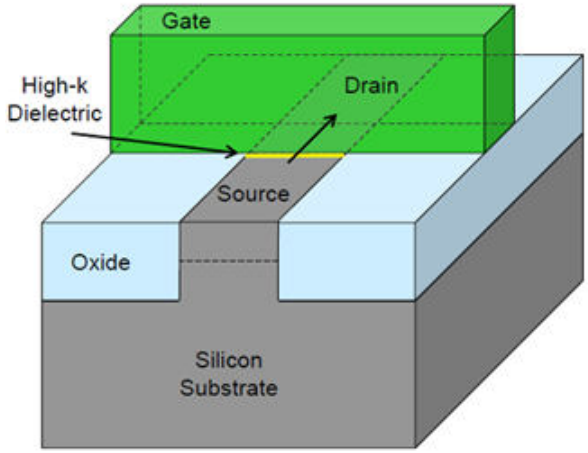

Planar CMOS

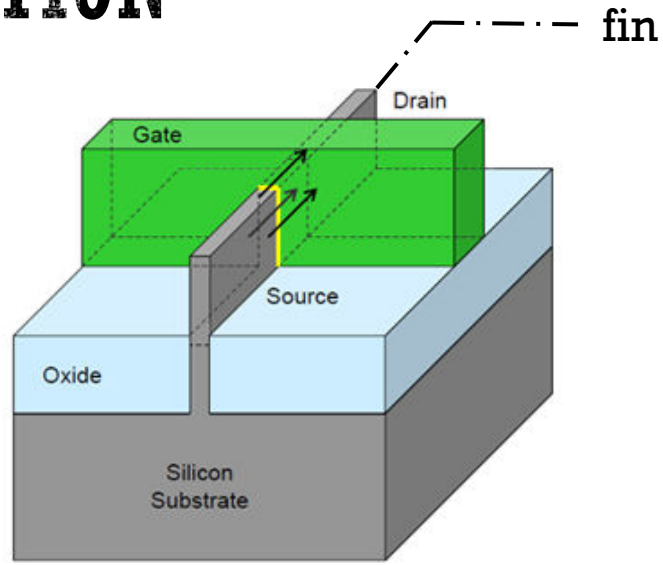

FinFET

- Planar MOSFETs have higher leakage power consumption, beyond 22 nm.
- FinFETs are double gate based devices with 'fins', increasing overall channel width.

Fins enable better control over channel current and minimize leakage.

INTRODUCTION

- Decrease in process geometries - larger scale of integration of components for on-chip networks (NoCs).
- This trend has given rise to :
  - process variations - variations in circuit parameters.
  - design time component failures - affect the regularity of the NoC topology.
  - power constraint - limits on the number of active cores on a die.

These effects have a significant impact during early design time exploration.

INTRODUCTION

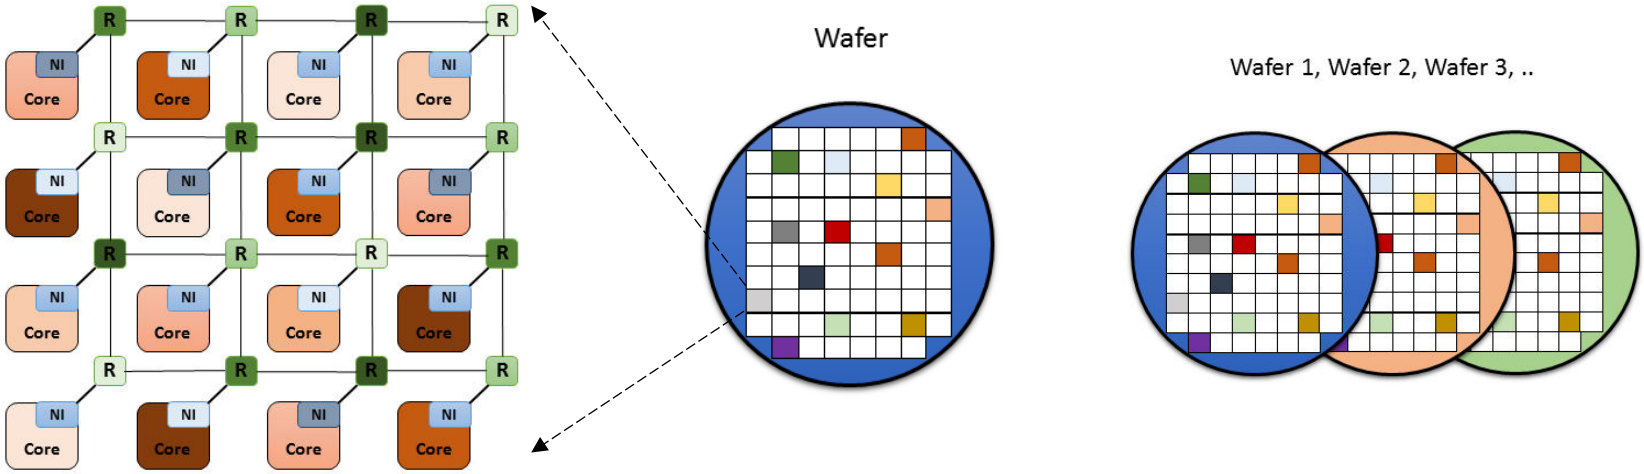

- Process variations manifest as :
  - Core to core (C2C) variations.
  - Die to Die (D2D) variations.
  - Wafer to Wafer (W2W) variations.

We only consider C2C and D2D variations in our work.

INTRODUCTION

- Increasing number of cores on a die lead to the "**power wall**" - power budget for a die.
- Defines the number of cores that can be turned on simultaneously on a die.
- Downscaling threshold voltage - exponentially increases leakage, causing to larger sections of the die to be powered down.

Projections indicate approximately 50% of the chip will need to be turned off at 8 nm.

INTRODUCTION

- Manufacturing defects induce failed links between cores, leading to an irregular topology.
- Routing algorithms for regular topologies - not applicable in this scenario.
- Selection of a topology agnostic routing algorithm is needed to ensure :
  - Connectivity
  - Deadlock freedom
  - Path diversity and selection

INTRODUCTION

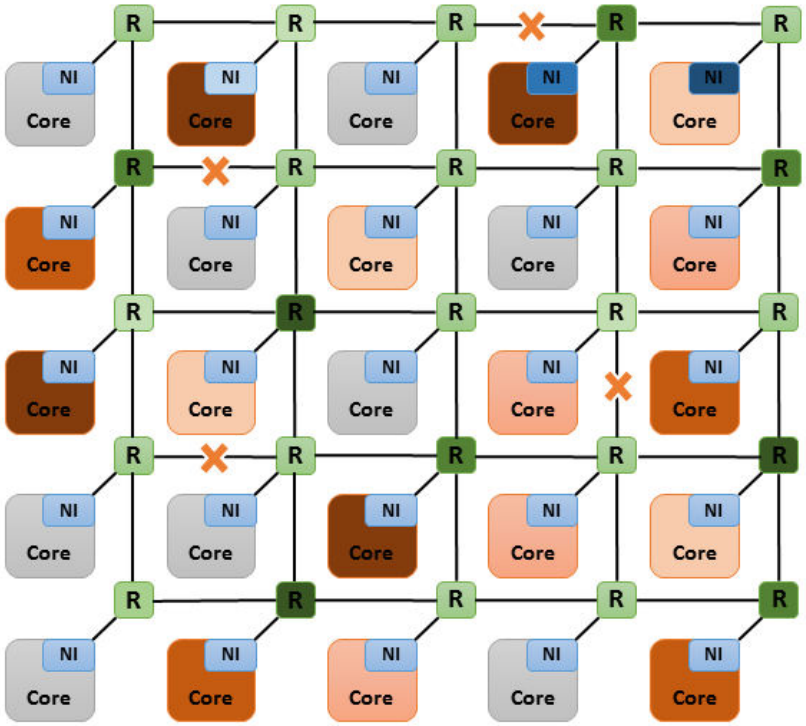

ILLUSTRATION OF A NoC WITH PROCESS VARIATIONS, DARK SILICON CONSTRAINTS AND FAILED LINKS

OUTLINE

- ✓ INTRODUCTION
- MOTIVATION
- OVERVIEW
- RELATED WORK
- PROBLEM STATEMENT
- METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

MOTIVATION

- There is no holistic consideration of the combined effects of
  - dark silicon considerations
  - FinFET based implementations
  - process variation implications

  for early design space exploration on irregular NoC topologies.

- Our work is the first to examine jointly:
  - process variation modeling,
  - on irregular mesh NoC implementations
  - with FinFET devices,
  - with dark silicon constraints.

MOTIVATION

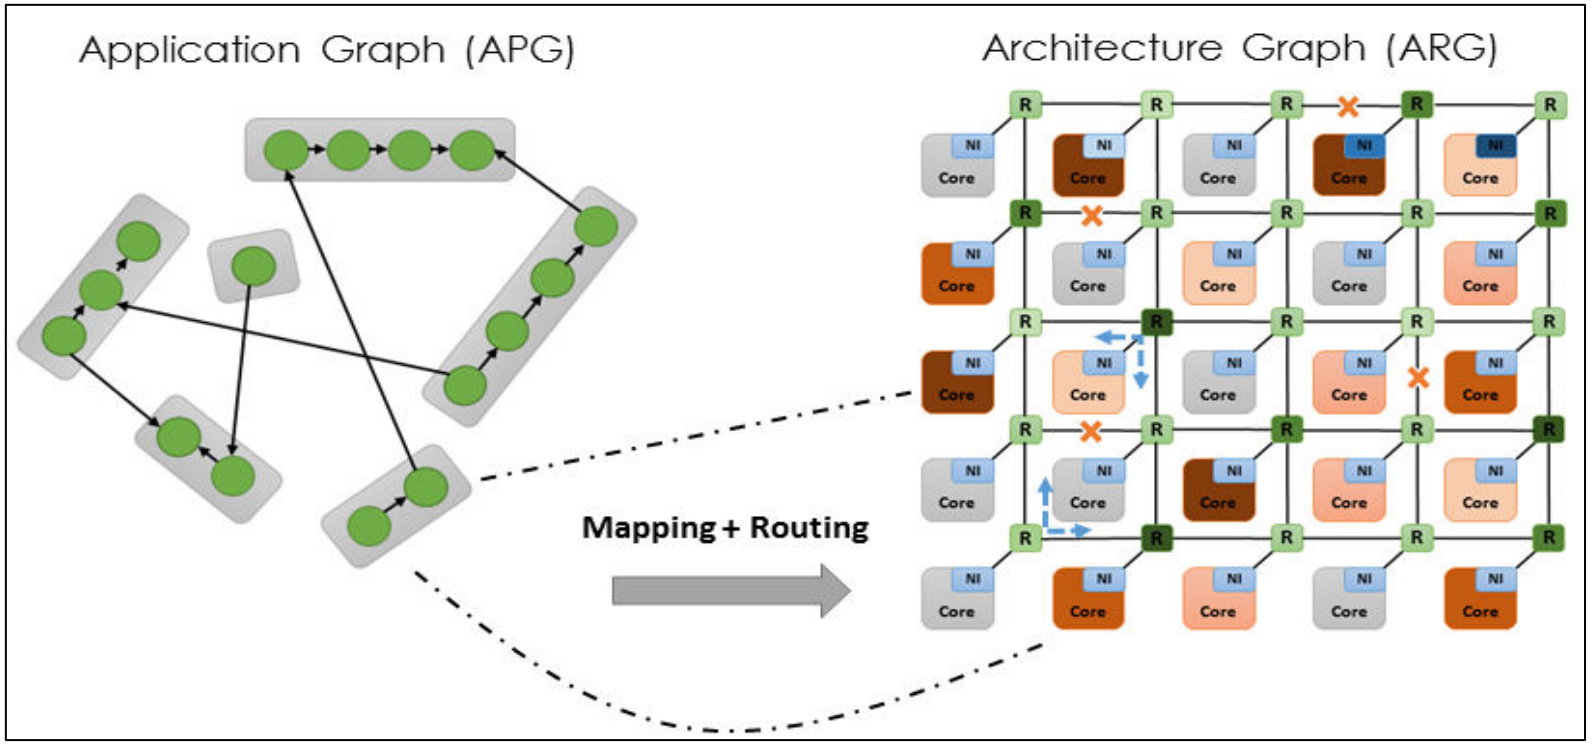

ILLUSTRATION OF HOLISTIC APPLICATION MAPPING AND ROUTING

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- OVERVIEW
- RELATED WORK
- PROBLEM STATEMENT
- METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

OVERVIEW – FINFET

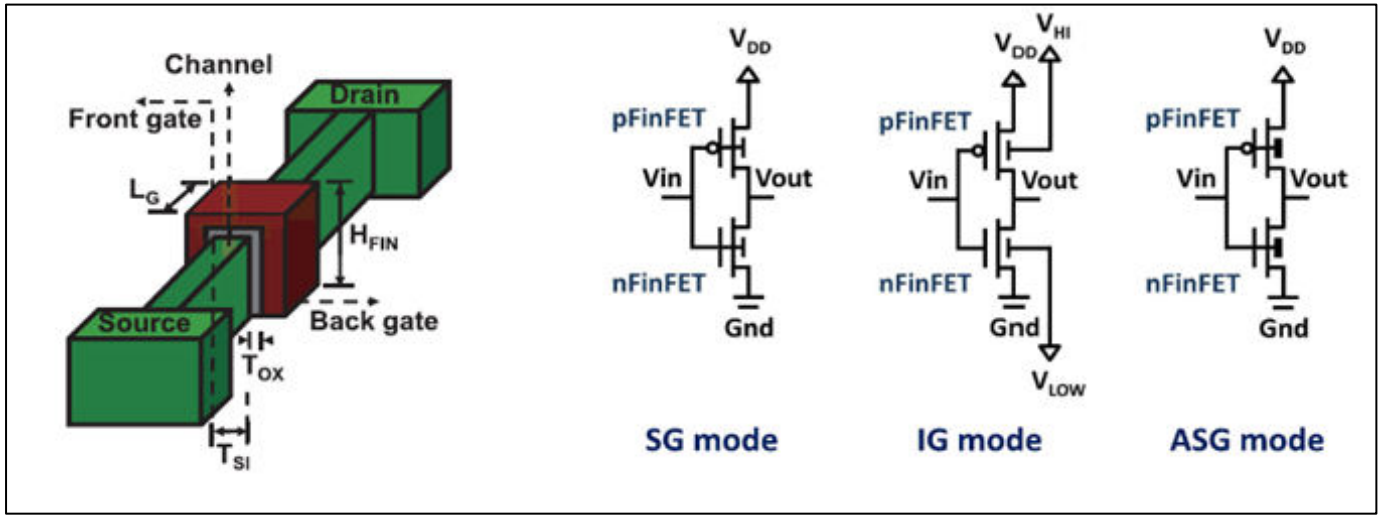

- SG (Shorted Gate) Mode
- IG (Independent Gate) Mode
- ASG (Asymmetric Shorted Gate) Mode

For our work, we choose ASG mode because of the best tradeoff between SG and IG modes.

OVERVIEW – PROCESS VARIATIONS

- Process variations modeling - done by addition of offset values to the original parameter.
- The equation to model variations in parameter (P) is:

$$\Delta P = \Delta P_{WID} + \Delta P_{D2D} = (\Delta P_{sys} + \Delta P_{rand}) + \Delta P_{D2D}$$

- WID (Within-die) variations - jointly modeled by systematic and random components.
- D2D (Die-to-Die) variations - modeled by an offset value for the parameter that differs from die to die.

Link (VARIUS) : http://web.cse.ohio-state.edu/~teodores/arch/tools/varius/varius.html

OVERVIEW - FINCANON

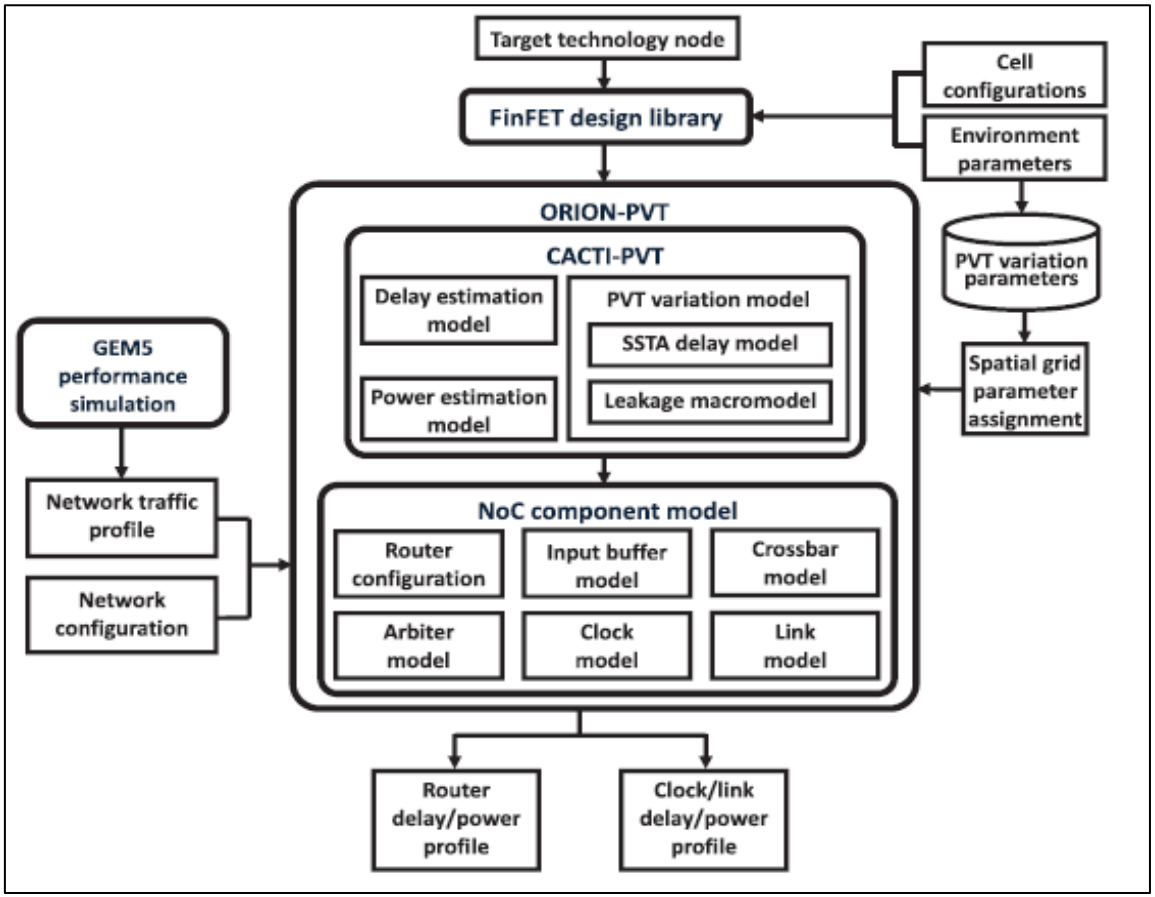

FinCANON incorporates ORION-PVT and CACTI-PVT.

Link (FinCANON) : https://www.princeton.edu/~jha/files/tools.html

OVERVIEW – MCPAT-PVT

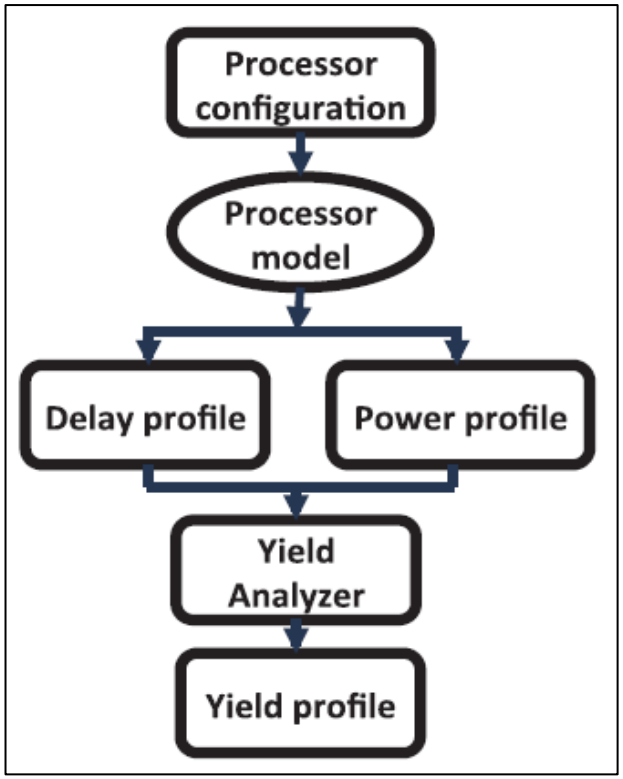

McPAT-PVT is an extension to McPAT with FinFET based PVT variations.

Link (McPAT-PVT) : https://www.princeton.edu/~jha/files/tools.html

OVERVIEW – ROUTER ARCHITECTURE

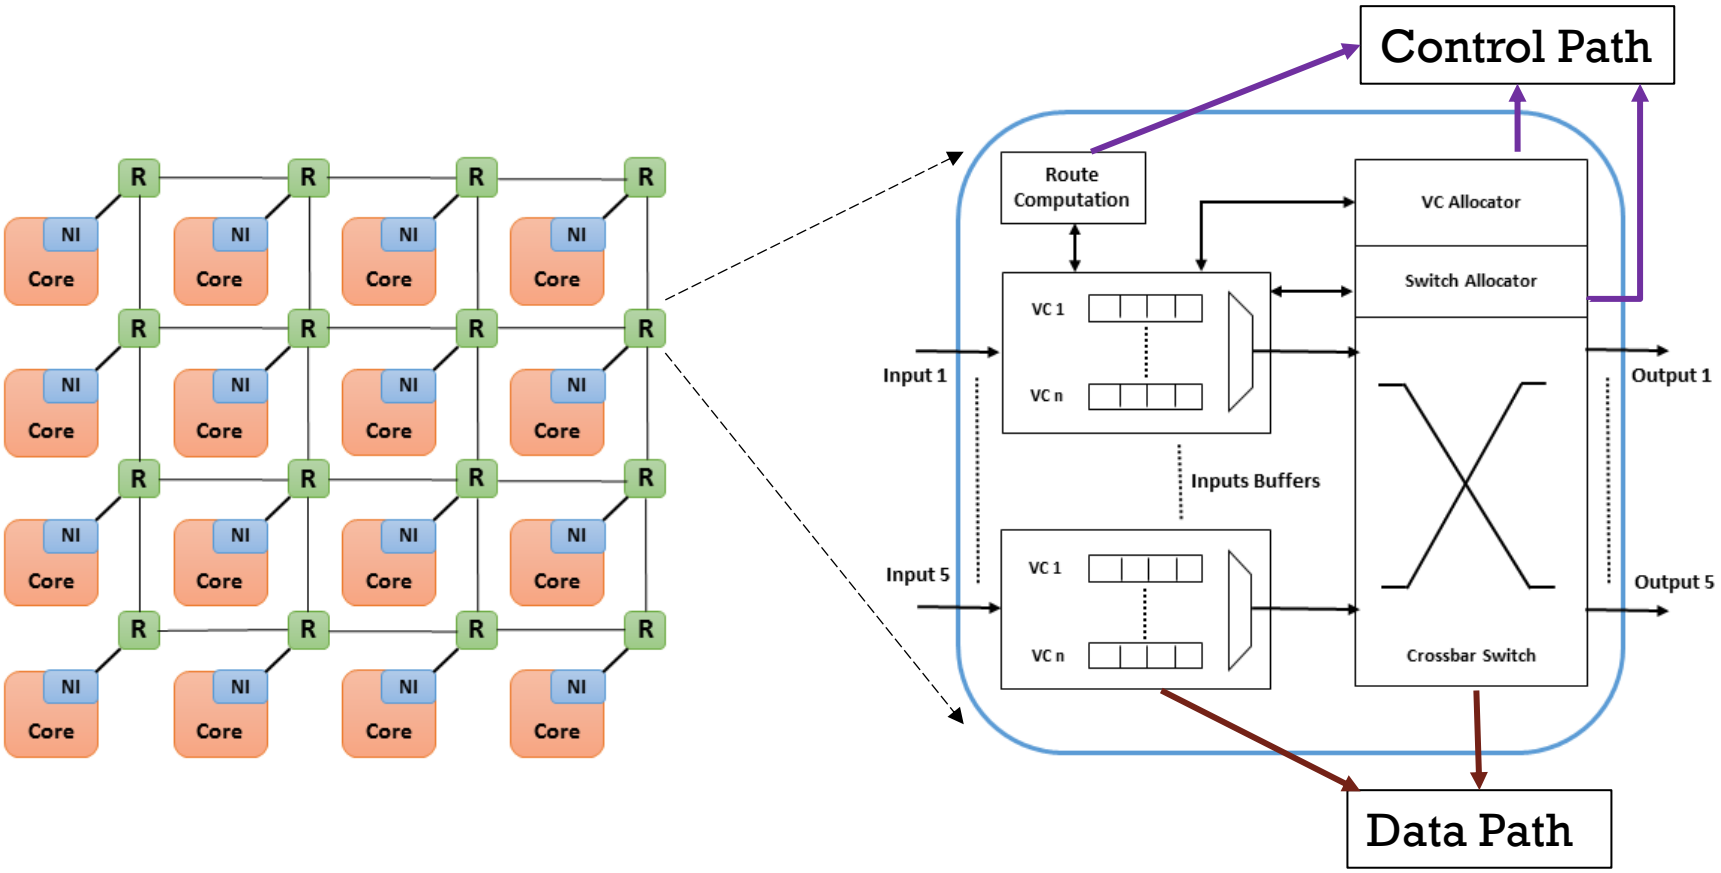

ROUTER ARCHITECTURE

Two types of router functional units : control path and data path.

OVERVIEW – POWER MODELING

- The equation to model combined power for a die:

$$P_{Combined} = P_{Total}^{Compute} + P_{Total}^{NoC}$$

- The total power between two tiles t_a and t_b for n bits:

$$P_{comm}(between\ t_a\ and\ t_b) = B * [n * (P_{Switch} + P_{Buffer}) + (n-1) * P_{Link}]$$

- The total power consumed by the NoC fabric, on the die, by M communications:

$$P_{Total}^{NoC} = \sum_{j=1}^{M} P_{comm_j}$$

OVERVIEW - PPY

- Power performance yield (PPY) - the number of dies meeting the power and performance constraints, from the set of test dies.
- The power performance yield (PPY) is represented by the equation:

$$PPY = \left(\frac{X}{Y} * 100\right)$$

No. of dies meeting constraints

Total no. of dies

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- ✓ OVERVIEW
- RELATED WORK
- PROBLEM STATEMENT
- METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

RELATED WORK

Process Variation Modeling, FinFETs and Dark Silicon

- E. Humenay, D. Tarjan, K. Skadron, "**Impact of process variations on multicore performance symmetry**," *IEEE Proceedings Design, Automation &Test in Europe*, pp. 1653-1658, Apr. 2007.
- B. Swahn and S. Hassoun, "**Gate sizing: FinFET versus. 32nm bulk MOSFETs**," *IEEE Proceedings Design, Automation &Test in Europe*, pp. 528–531, Jul. 2006.
- H. Esmaeilzadehy, E. Blemz, R. St. Amantx, K. Sankaralingamz, D. Burger, "**Dark Silicon and the End of Multicore Scaling**", pp. 365 – 376, *International Symposium on Computer Architecture*, Jun. 2011.

The works formed the basis of the process variation, FinFET and dark silicon based platforms for our exploration.

RELATED WORK

Process Variation Modeling, FinFETs and Dark Silicon

- S. R. Sarangi, B. Greskamp, R. Teodorescu, J. Nakano, A. Tiwari, J. Torrellas, "**VARIUS: A Model of Process Variation and Resulting Timing Errors for Microarchitects**", *IEEE Transactions on Semiconductor Manufacturing*, pp. 3-13 ,Feb 2008.
- C. Lee, N. Jha, "**FinCANON: A PVT-Aware Integrated Delay and Power Modeling Framework for FinFET-Based Caches and On-Chip Networks**", *IEEE TVLSI*, Apr. 2014.
- A. Tang, Y. Yang, C. Lee, N. Jha, "**McPAT-PVT: Delay and Power Modeling Framework for FinFET Processor Architectures Under PVT Variations**", *IEEE TVLSI*, Sept. 2014.

The works provided the process variation based modeling tools for our exploration.

RELATED WORK

Work on Routing Techniques for Irregular Topologies

- A. Mejia, J. Flich, J. Duato, S. Reinemo, T. Skeie, "**Segment-based routing: an efficient fault tolerant routing algorithm for meshes and tori**," *IEEE International Parallel and Distributed Processing Symposium*, Apr. 2006.
- A. Mejia, J. Filch, J. Duato "**On the Potentials of Segment-Based Routing for NoCs**," *IEEE International Conference on Parallel Processing*, pp 594-603, Sept. 2008.
- J. Flich, "**A Survey and Evaluation of Topology-Agnostic Deterministic Routing Algorithms**", *IEEE Transactions on Parallel & Distributed Systems*, pp. 405-425, Mar. 2012.

These works helped choose the appropriate routing algorithm for the exploration.

RELATED WORK

Application Mapping for Various Optimization Objectives

- J. Hu, R. Marculescu, "**Energy and Performance-Aware Mapping for Regular NoC Architectures**," *DATE*, Jun. 2003.
- T. Lei, S. Kumar ,"**A Two-step Genetic Algorithm for Mapping Task Graphs to a Network on Chip Architecture**," *EUROMICRO,* 2003
- R. Tornero, J. Ordua, A. Mejia, J. Flich, J. Duato ,"**CART: Communication-Aware Routing Technique for Application-Specific NoCs**," *EUROMICRO*, Sept. 2008
- P.K. Sahu, P. Venkatesh, S. Gollapalli, S. Chattopadhyay, "**Application Mapping onto Mesh Structured Network-on-Chip using Particle Swarm Optimization**", *IEEE Computer Society Annual Symposium on VLSI*, pp 335-336, Jul. 2011

These works are generic mapping and routing techniques, without any process variation considerations.

Our work is the first to integrate process variations, dark silicon and topology agnostic routing/path selection.

RELATED WORK

Process Variation Aware Application Mapping

- N. Kapadia, S. Pasricha, "**Process Variation Aware Synthesis of Application-Specific MPSoCs to Maximize Yield**", *International Conference on Embedded Systems and International Conference on Embedded Systems*, pp 270-275, Jan. 2014.
- N. Kapadia and S. Pasricha, "**VERVE: A Framework for Variation-Aware Energy Efficient Synthesis of NoC-based MPSoCs with Voltage islands**," *ISQED*, Mar. 2013.
- S. Majzoub, R. Saleh, S. Wilton, and R. Ward, "**Energy Optimization for Many-Core Platforms: Communication and PVT Aware Voltage-Island Formation and Voltage Selection Algorithm**", *IEEE TCAD*, May 2010.
- D. Mirzoyan, B. Akesson, and K. Goossens, "**Process-variation aware mapping of real-time streaming applications to MPSoCs for improved yield**," *ISQED*, 2012.

These works do not incorporate FinFET based NoCs.

Our work builds on these works by adding the FinFET perspective.

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- ✓ OVERVIEW
- ✓ RELATED WORK
- PROBLEM STATEMENT
- METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

PROBLEM STATEMENT

The objective of our framework is to :

- **generate an one-to-one core to tile mapping**,
- and **map the communication flows between app. cores to the links.**

on a 2-D irregular mesh NoC, statistically modeled with process variations, for a given set of parallel applications, such that:

- the system energy is minimized and PPY is maximized,
- while the network connectivity and deadlock freedom is guaranteed,
- the link bandwidth and dark silicon based power (computation and communication) constraints are met.

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- ✓ OVERVIEW
- ✓ RELATED WORK
- ✓ PROBLEM STATEMENT
- METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

METHODOLOGY - FLOW

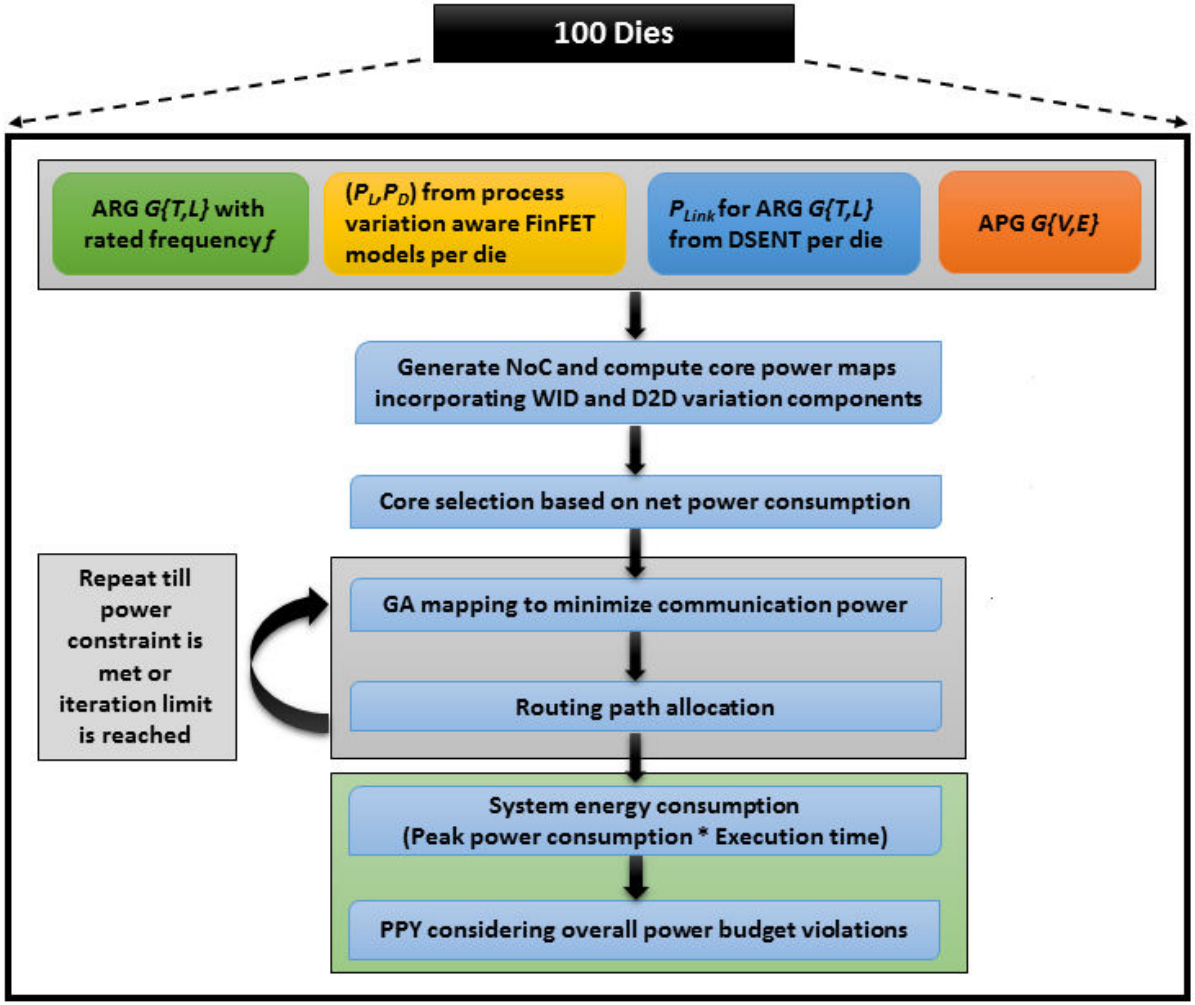

Link (DSENT) : https://sites.google.com/site/mitdsent/

METHODOLOGY

inputs: Dynamic and leakage power values for compute cores on the die

1: Compute the total compute power (P_L +P_D) for compute cores for each tile and store the tile ids, for all the tiles in the die.

2: Perform sorting of the power values and store tile ids in an ascending order with the total compute power.

3: Choose an equal number of tiles from the sorted set, as the number of application cores, to perform mapping.

4: Calculate and store the total compute power from the equation:

$$P_{Selected\ Tiles}^{Compute} = \sum_{i=1}^{C} P_{Li} + P_{Di}$$

output : Sorted list of selected NoC tiles with power values

CORE SELECTION ALGORITHM

METHODOLOGY

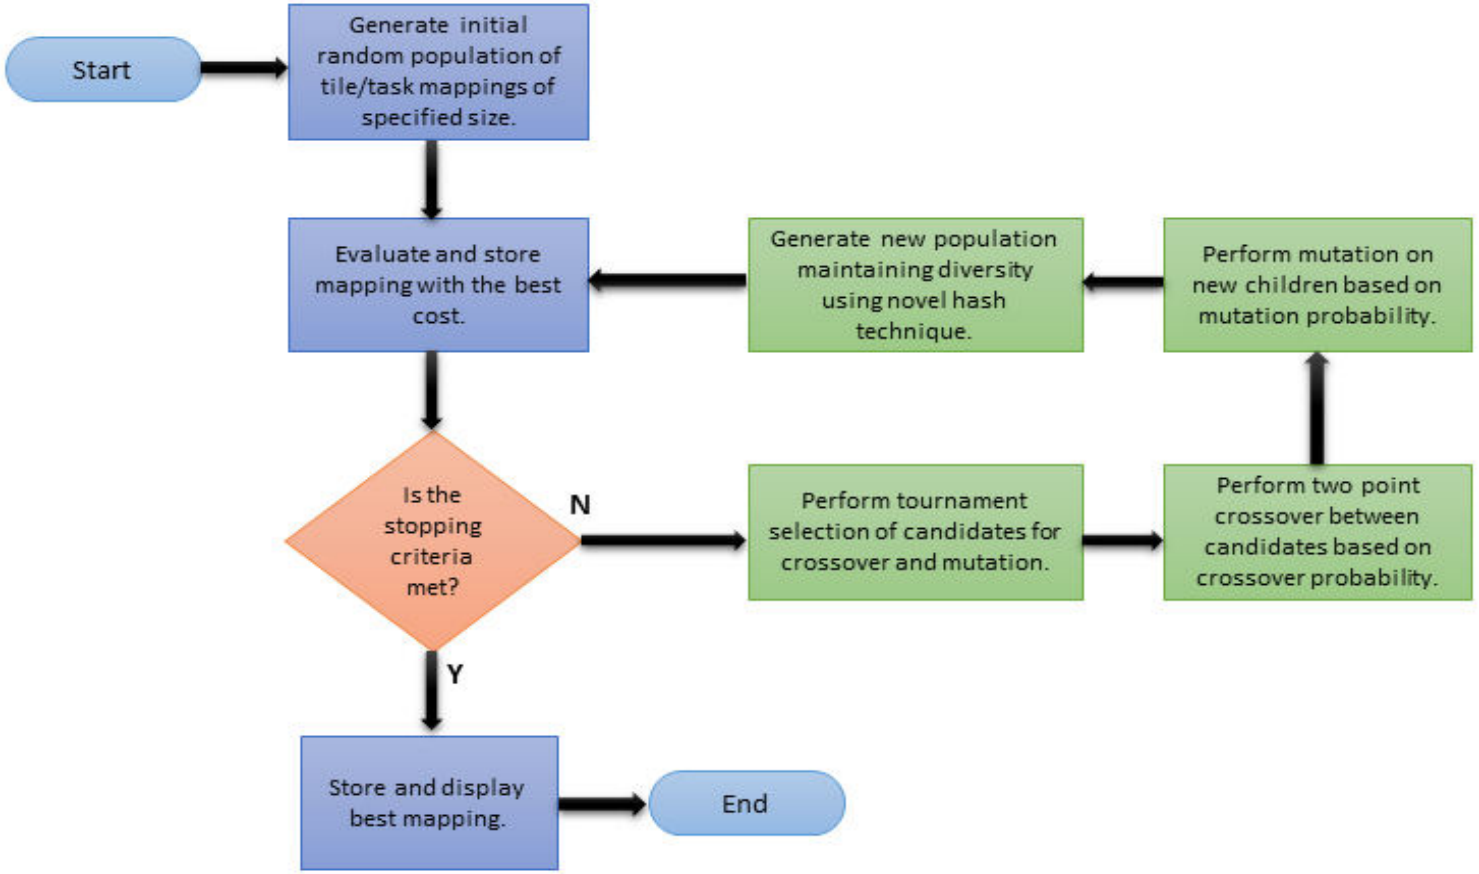

CUSTOM GENETIC ALGORITHM

METHODOLOGY

inputs: ***core tile mapping as genotype x***

1: Choose a prime number c, typically greater than 10 times the population size N.
2: Set the start index i =1 and initial hash of the mapping as r = 0.
3: Calculate the hash value

$$r = \left(r * l + x(i)\right) mod\ c$$

where l is the number of genes per genotype, $x(i)$ is the gene value at location l in the genotype x.
4: Increment the index value i by 1.
5: If i is less or equal than 1, go to step2.
6: Store the value of r as the hash value for the given mapping.

output : ***generated hash value for genotype x***

HASH GENERATION ALGORITHM

Reference : S. Ronald, "Preventing Diversity Loss in a Routing Genetic Algorithm with Hash Tagging," Complex Systems: Mechanism of Adaptation, pp 133-140, 1994.

METHODOLOGY

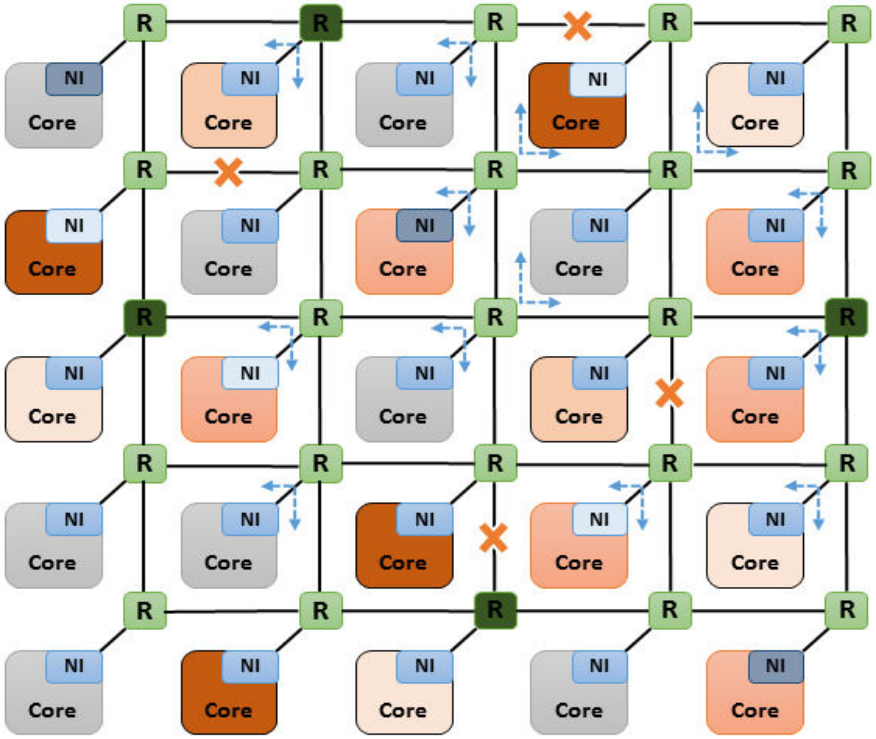

SEGMENT ROUTING TURN RESTRICTIONS

Segments preserve connectivity. Turn restrictions prevent deadlocks and offer path diversity.

METHODOLOGY

Min cost max flow (MCF) Formulation:

- ***Capacity Constraints:*** $\sum_{m=1}^{K} f_m(l_i) \leq c(t_i, t_j)$
- ***Flow conservation:*** $\sum_{n \in T} f_m(l_i) = 0$

  $when\ u \neq t_i, t_j\ \forall\ u, v, f_m(l_i) = -f_m(l_i)$
- ***Demand satisfaction:*** $\sum_{w \in T} f_m(t_i, w) = \sum_{w \in T} f_m(w, t_j) = h_i$
- ***Minimize hop count * communication volume:***

  $\sum_{(l_i) \in E} (h(l_i) * \sum_{m=1}^{K} f_m(l_i))$

MCF formulation identifies the shortest bandwidth restricted path.

Link (lp_solve) : http://sourceforge.net/projects/lpsolve/

METHODOLOGY

- Computation power is evaluated as the sum of the dynamic and leakage powers of the mapped compute cores.
- The communication power is estimated from the proposed power model, for NoC data path components.
- Total power is evaluated for all 100 dies and evaluated for the number of dies that meet the power constraint.
- The system energy (for the best case die)

$$E_{system} = (P_{Total} * T_{Sim})$$

The PPY evaluation step is performed to estimate the die yield.

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- ✓ OVERVIEW
- ✓ RELATED WORK
- ✓ PROBLEM STATEMENT
- ✓ METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

RESULTS

- Simulations performed on eight parallel benchmarks from the SPLASH 2 and PARSEC benchmark suites
  - Less comm. intensive – *(blackscholes, fluidanimate, cholesky and ocean)*
  - More comm. intensive – (*dedup, canneal, fft and lu*)
- We consider a 144 tile NoC topology with 100 active cores and 10% failed links as the base case. The power budgets are set to :
  - *comm_int_low* set:170W (stringent) and 180W (nominal)
  - *comm_int_high* set :230W (stringent) and 240W (nominal)
- Power budgets are scaled to perform sensitivity analysis with variance in system size and fault rate.

RESULTS

- We run HERMES (GA) for a population range of 50 to 100, with a limit on the number of generations to 500. The crossover probability was set to 0.6 and the mutation probability was set to 0.02.
- HERMES is compared with 3 other techniques:
  - Simulated Annealing (SA) - start temperature of 100 and a total limit on the iterations of 300. [J. Hu, R. Marculescu, "Energy and Performance-Aware Mapping for Regular NoC Architectures," DATE, Jun. 2003.]
  - Particle Swarm Optimization (PSO) - swarm size of 50 to 100, with epoch limit of 500. [P.K. Sahu, P. Venkatesh, S. Gollapalli, S. Chattopadhyay, "Application Mapping onto Mesh Structured Network-on-Chip using Particle Swarm Optimization", IEEE Computer Society Annual Symposium on VLSI, pp 335-336, Jul. 2011]
  - Communication Aware Routing Technique (CART) - iteration limit of 100 for 3 random perturbations. [R. Tornero, J. Ordua, A. Mejia, J. Flich, J. Duato ,"CART: Communication-Aware Routing Technique for Application-Specific NoCs," EUROMICRO, Sept. 2008]

RESULTS

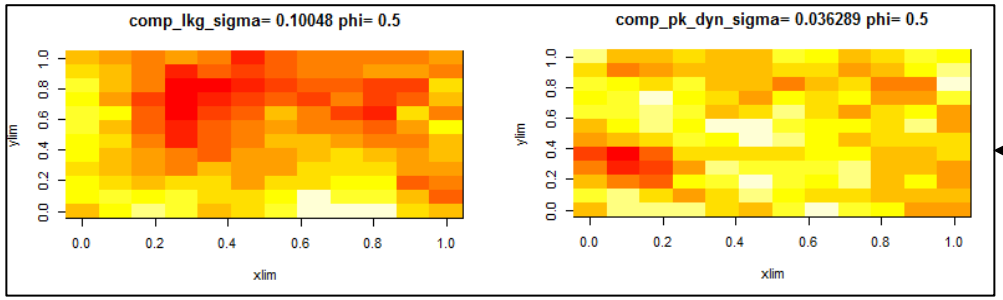

Compute Core

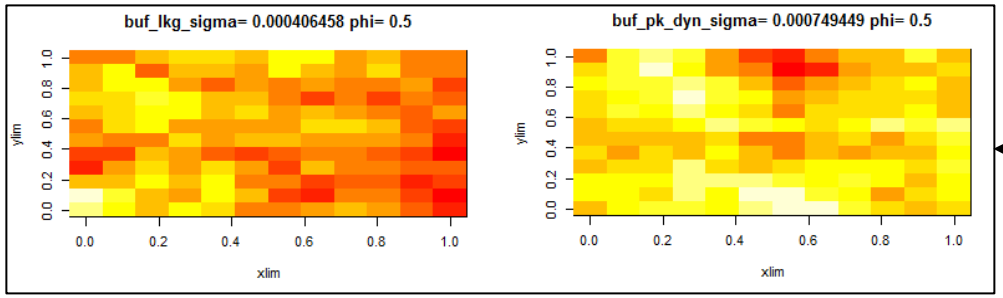

NoC Buffer

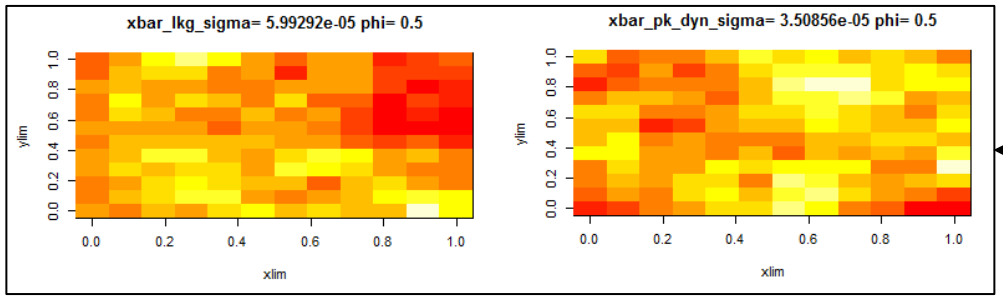

NoC Switch

RESULTS

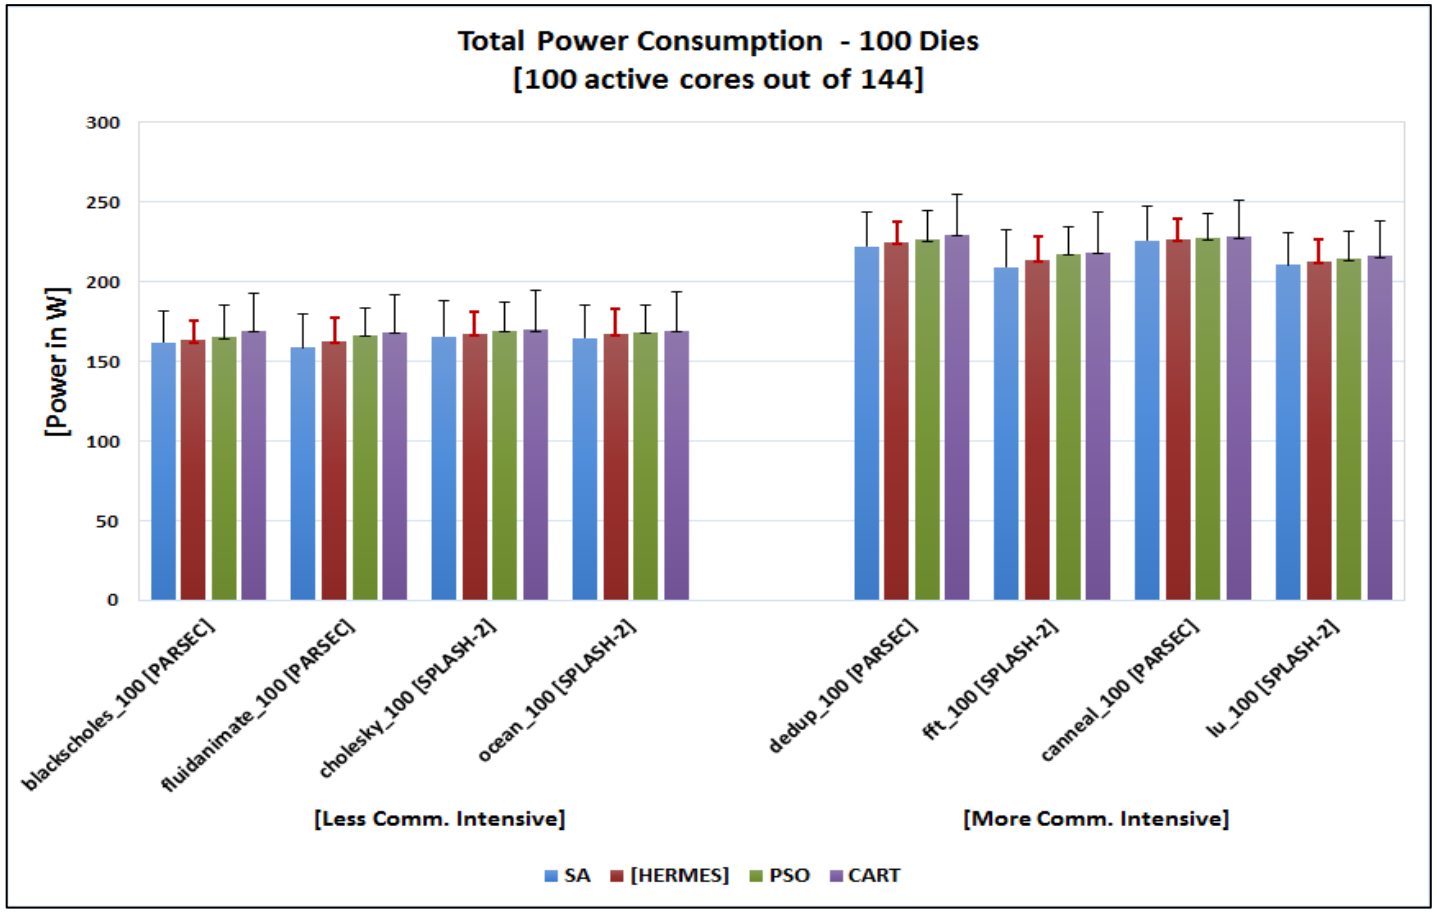

HERMES offers the least variance in power values among the SA, CART and PSO techniques

RESULTS

Execution Time (in seconds)	SA	HERMES	PSO	CART
blackscholes_100 **[PARSEC]**	8665	**6700**	7123	8523
dedup_100 **[PARSEC]**	8723	**6965**	7333	8632
fluidanimate_100 **[PARSEC]**	8821	**6854**	7435	8535
canneal_100 **[PARSEC]**	8735	**6745**	7395	8594
cholesky_100 **[SPLASH-2]**	8698	**6731**	7112	8412
fft_100 **[SPLASH-2]**	8780	**6774**	6992	8292
ocean_100 **[SPLASH-2]**	8432	**6814**	7213	8213
lu_100 **[SPLASH-2]**	8541	**6790**	7011	8311

- HERMES offers better execution times, compared to the SA, PSO and CART techniques.
- HERMES consistently performs better than the other techniques, across all the benchmarks in the SPLASH-2 and PARSEC sets.

RESULTS

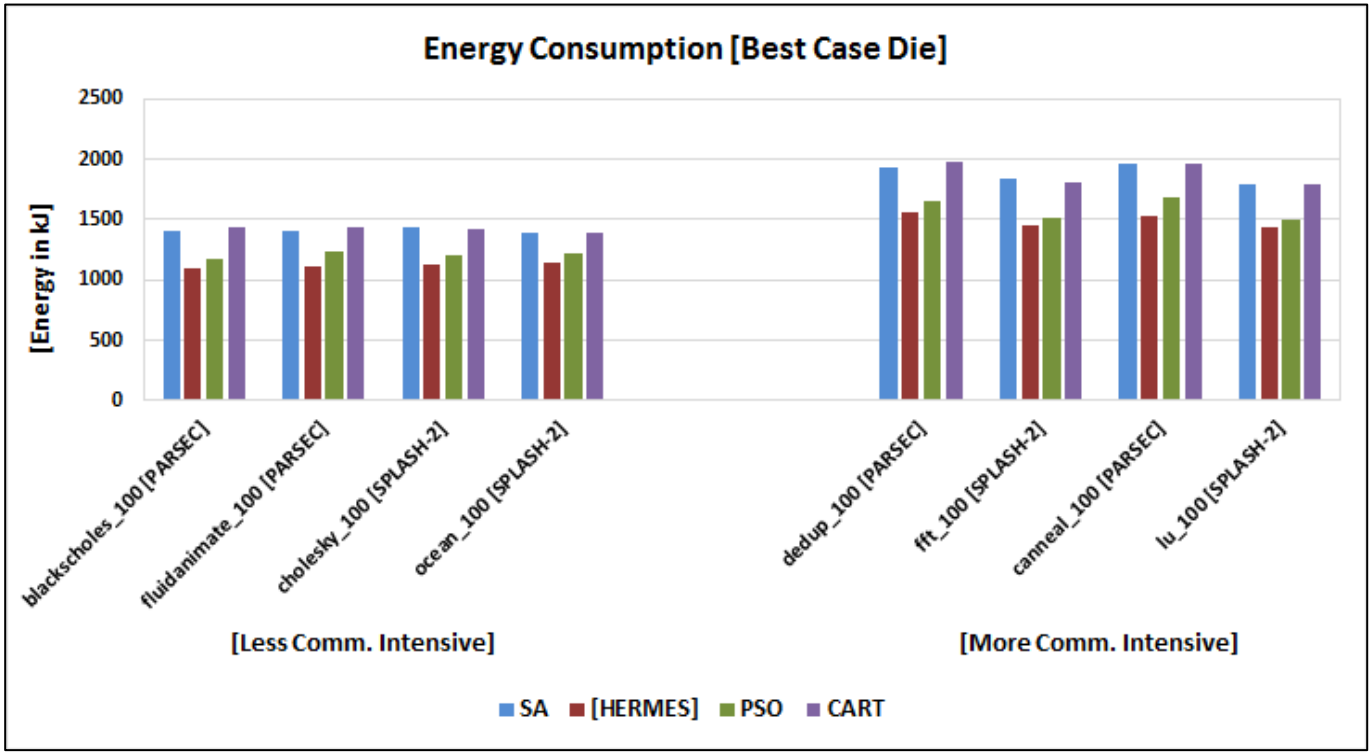

Improvements	SA	PSO	CART
comm_int_low	1.28*x*	1.11*x*	1.32*x*
comm_int_high	1.29*x*	1.10*x*	1.28*x*

RESULTS

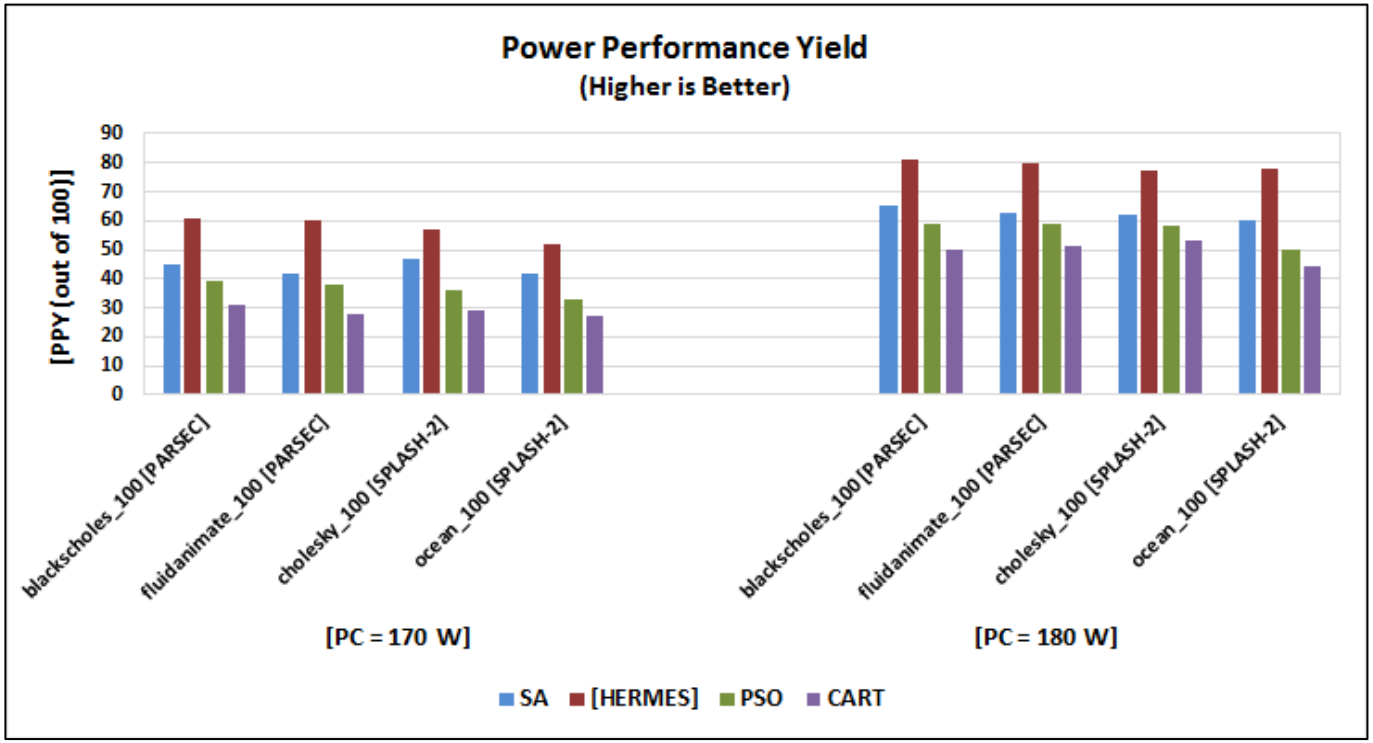

Improvements	SA	PSO	CART
comm_int_low (170W)	30%	36%	53%
comm_int_low (180W)	23%	35.9%	43.5%

RESULTS

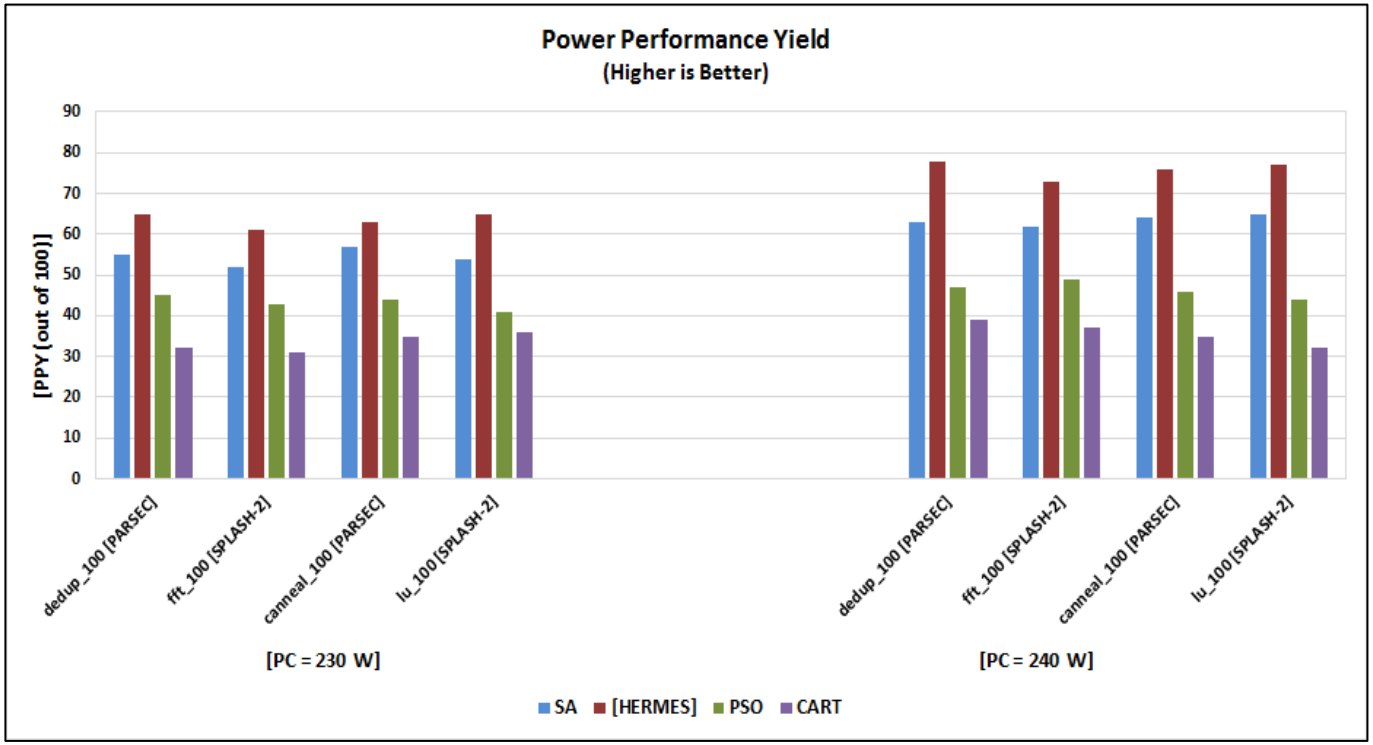

Improvements	SA	PSO	CART
comm_int_high (230W)	16.92%	36.92%	50.77%
comm_int_high (240W)	19.23%	39.74%	58.44%

RESULTS

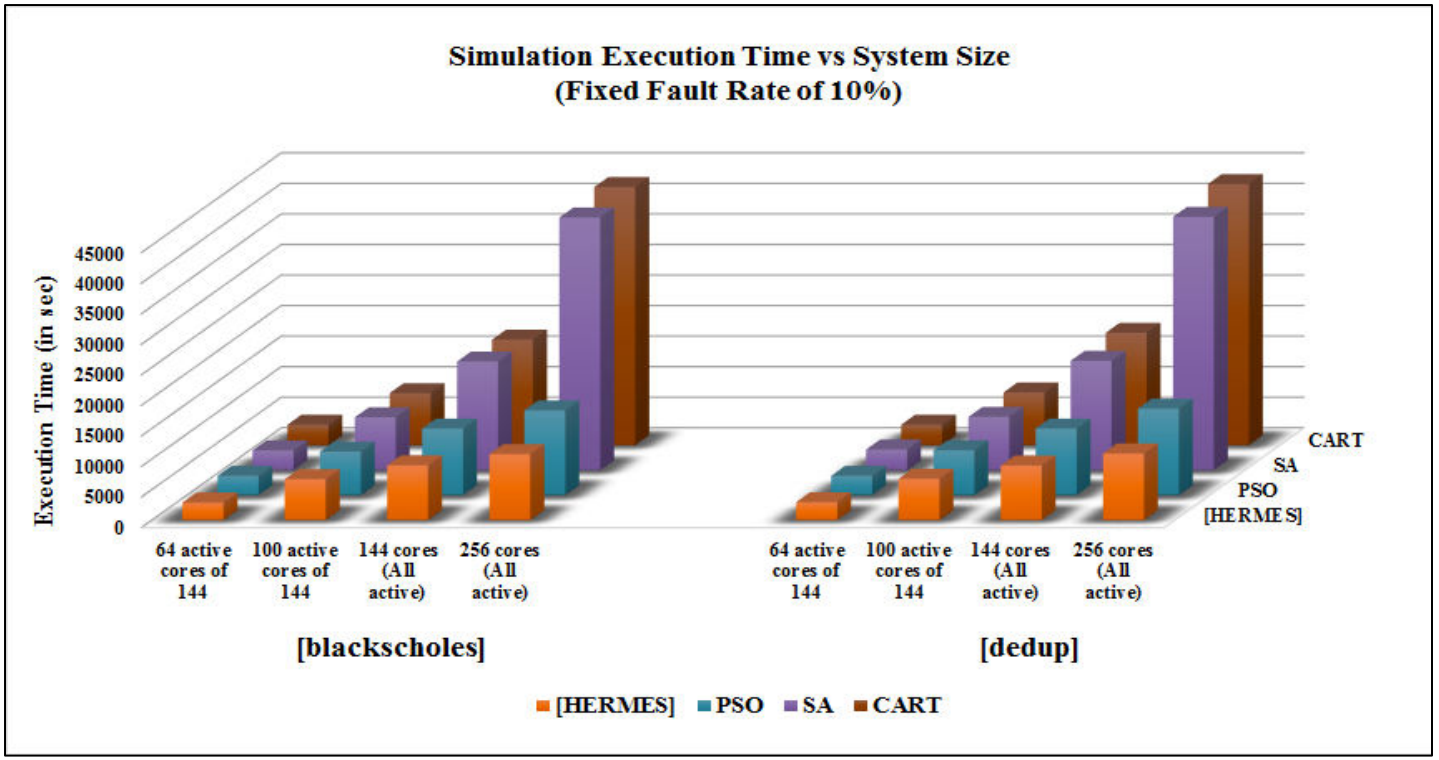

HERMES offers excellent execution time scaling with increasing system sizes.

RESULTS

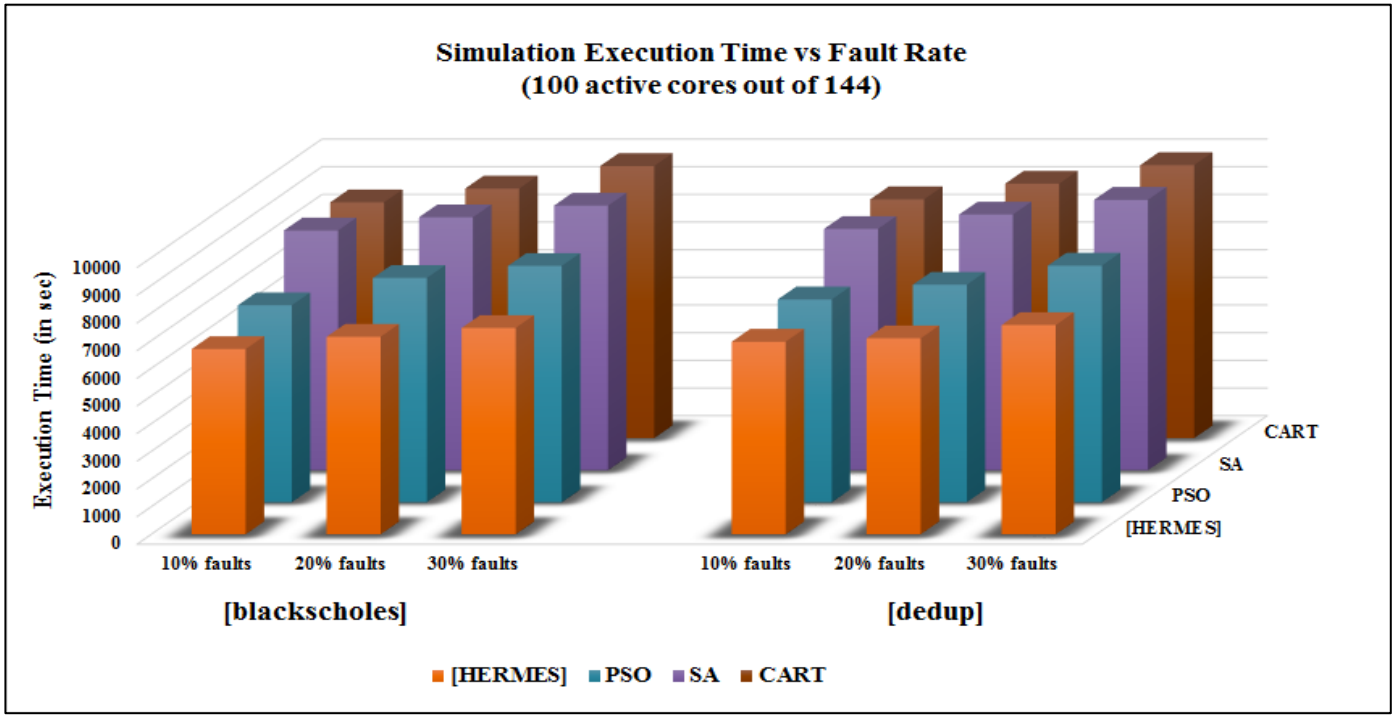

The MCF implementation ensures consistent execution time scaling with increasing fault rates for all 4 algorithms.

RESULTS

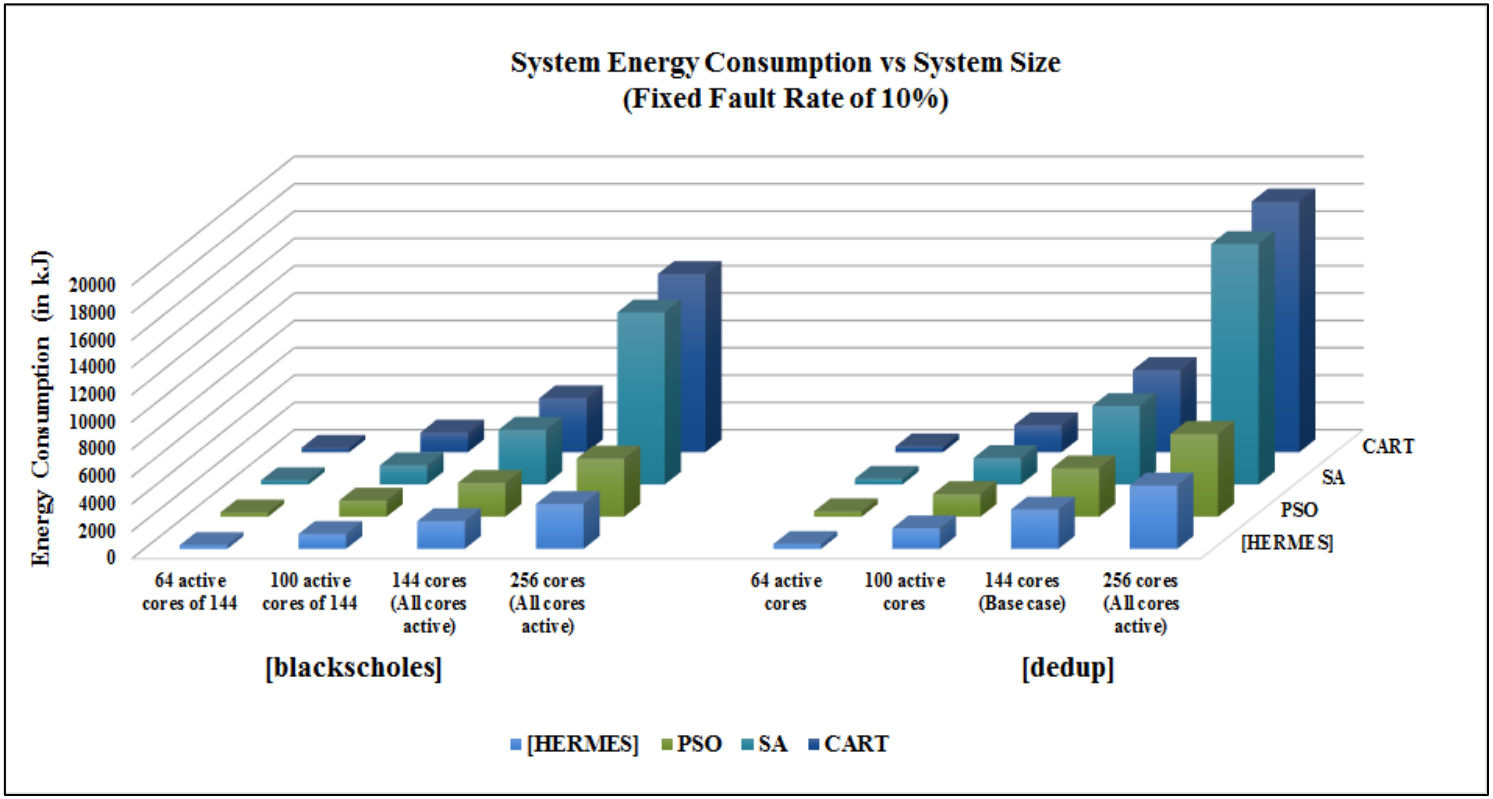

HERMES offers excellent energy figures with increasing system sizes, compared to other schemes.

RESULTS

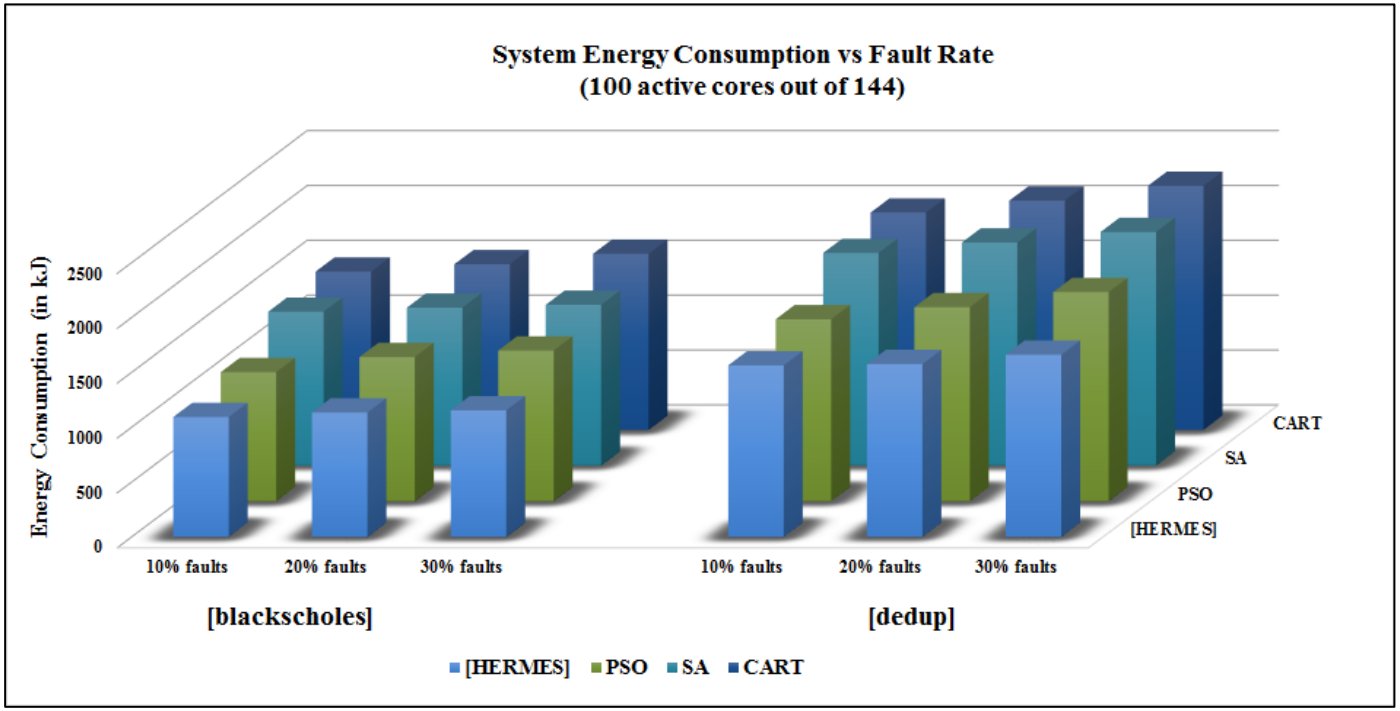

HERMES offers consistent scaling of energy figures with increasing fault rates.

RESULTS

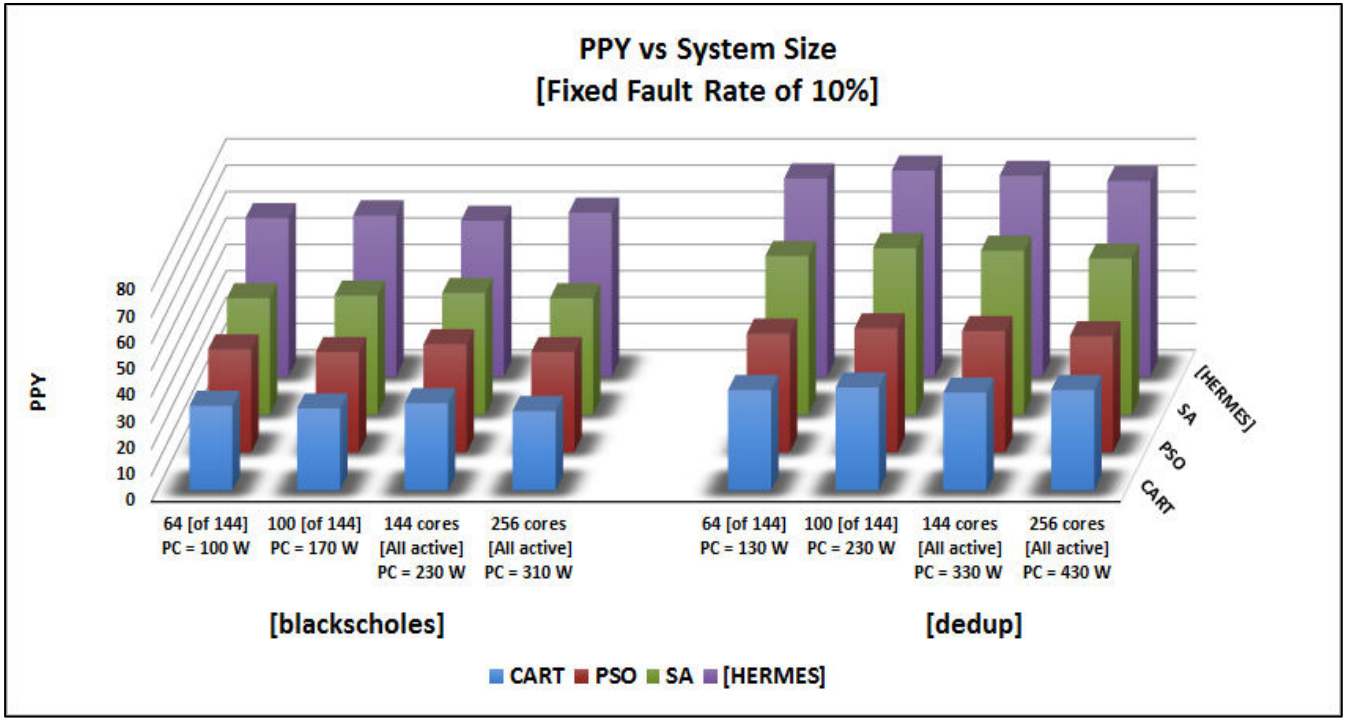

HERMES offers consistent PPY figures with increasing system sizes.

RESULTS

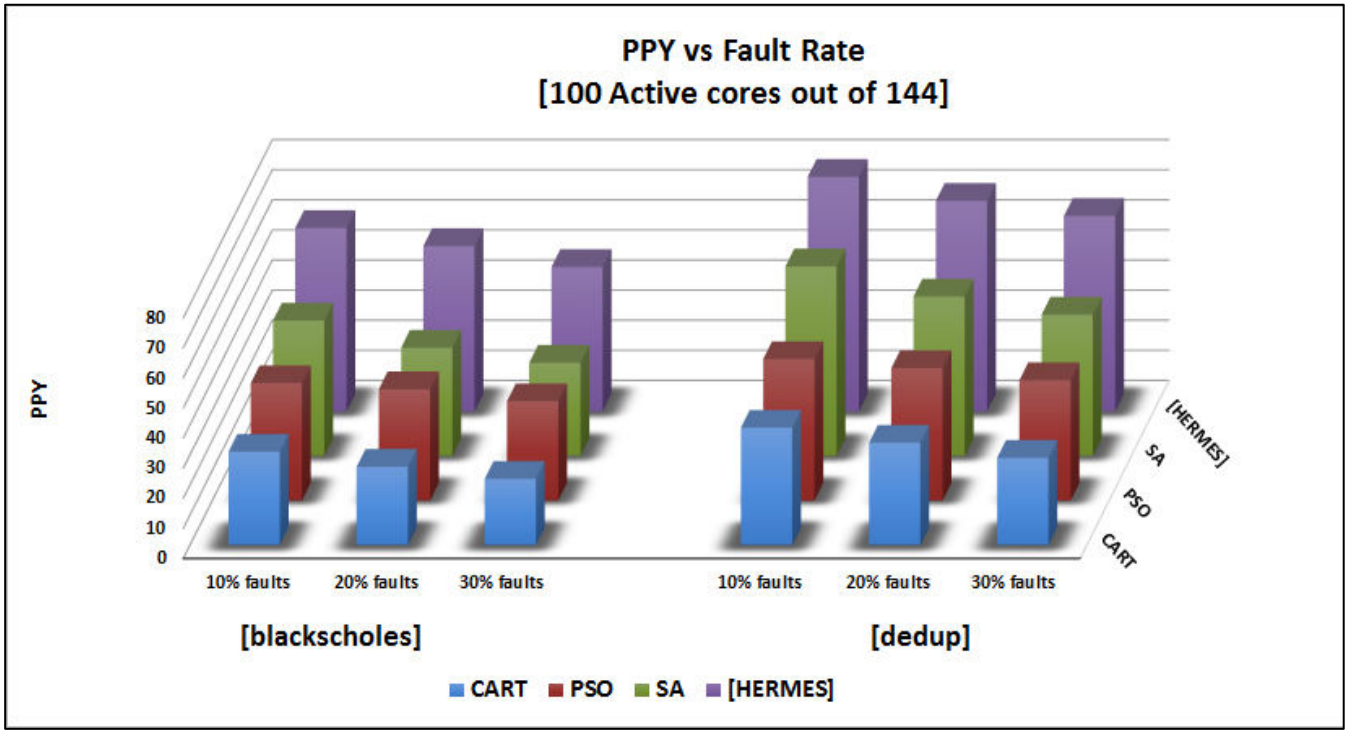

HERMES offers steady scaling of PPY figures with increasing fault rates.

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- ✓ OVERVIEW
- ✓ RELATED WORK
- ✓ PROBLEM STATEMENT
- ✓ METHODOLOGY
- ✓ RESULTS
- ▪ CONCLUSION AND FUTURE WORK

CONCLUSION

- Increasing impact of process variations in FinFET based designs, has placed increasing importance on modeling process variations in application mapping and routing frameworks.
- Conventional design flows without process variation awareness generate suboptimal solutions.
- For the first time, we have incorporated process variation aware FinFET based design methodology into a custom mapping and routing framework for irregular mesh based NoCs.
- Our framework performs up to 1.32*x* better in energy, 1.29*x* in execution time and 58.44% better in PPY statistics respectively, than other proposed schemes.

FUTURE WORK

There are several key extension areas to our work:

- Exploration of our framework for sub 10 nm FinFET designs.
- Exploration of our framework on NTC based methodologies.
- Effect of random variations at very low process geometries.
- Addition of process variation model for active links.
- Inclusion of multi-application mapping on a heterogeneous NoC fabric.

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- ✓ OVERVIEW
- ✓ RELATED WORK
- ✓ PROBLEM STATEMENT
- ✓ METHODOLOGY
- ✓ RESULTS
- ✓ CONCLUSION AND FUTURE WORK

BACKUP SLIDES

INTRODUCTION

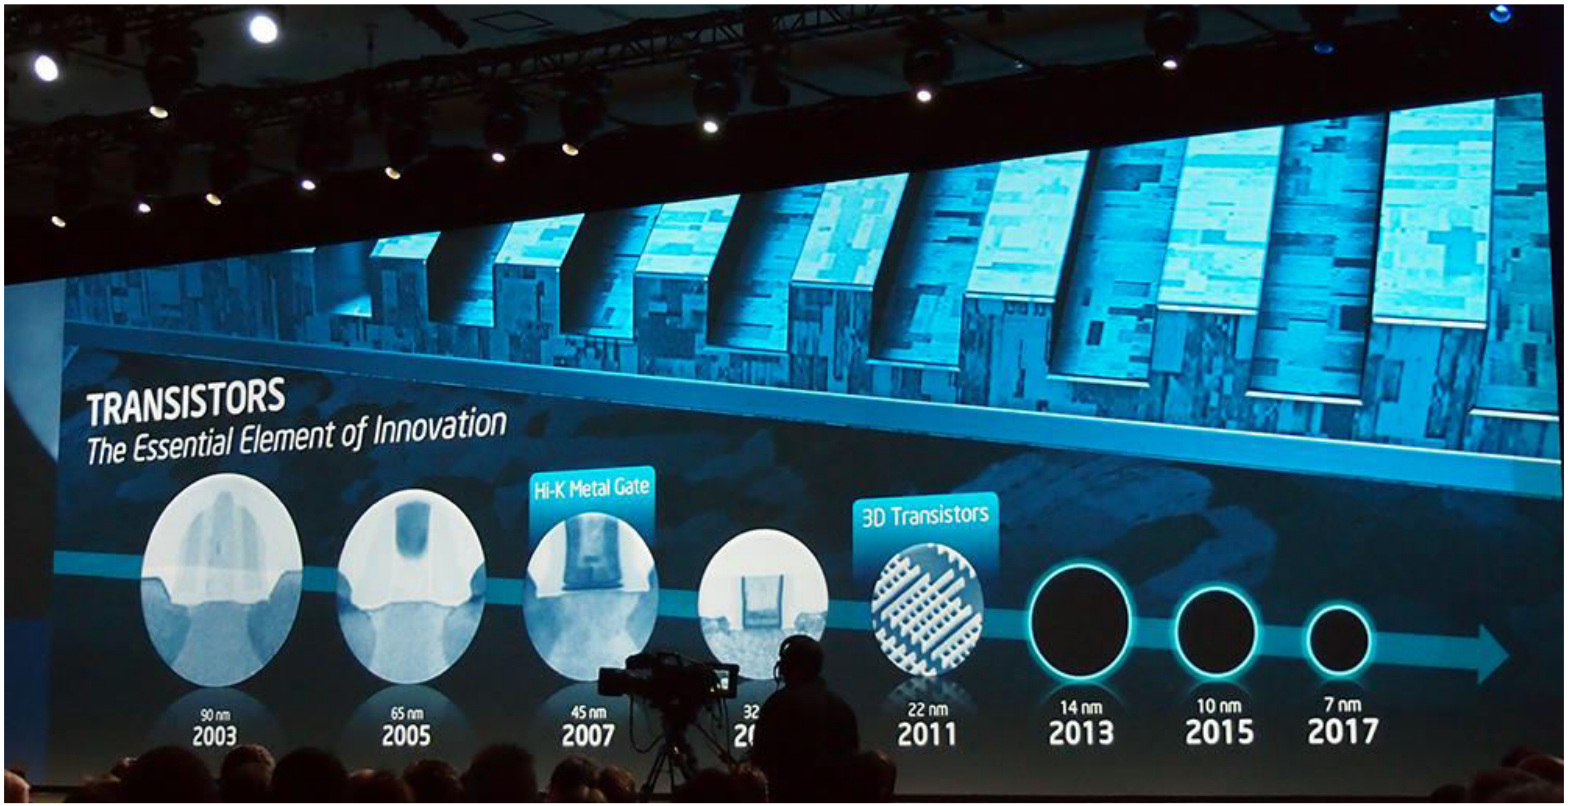

Intel Technology Roadmap

Indicates a clear focus on 3D transistors (FinFETs)

INTRODUCTION

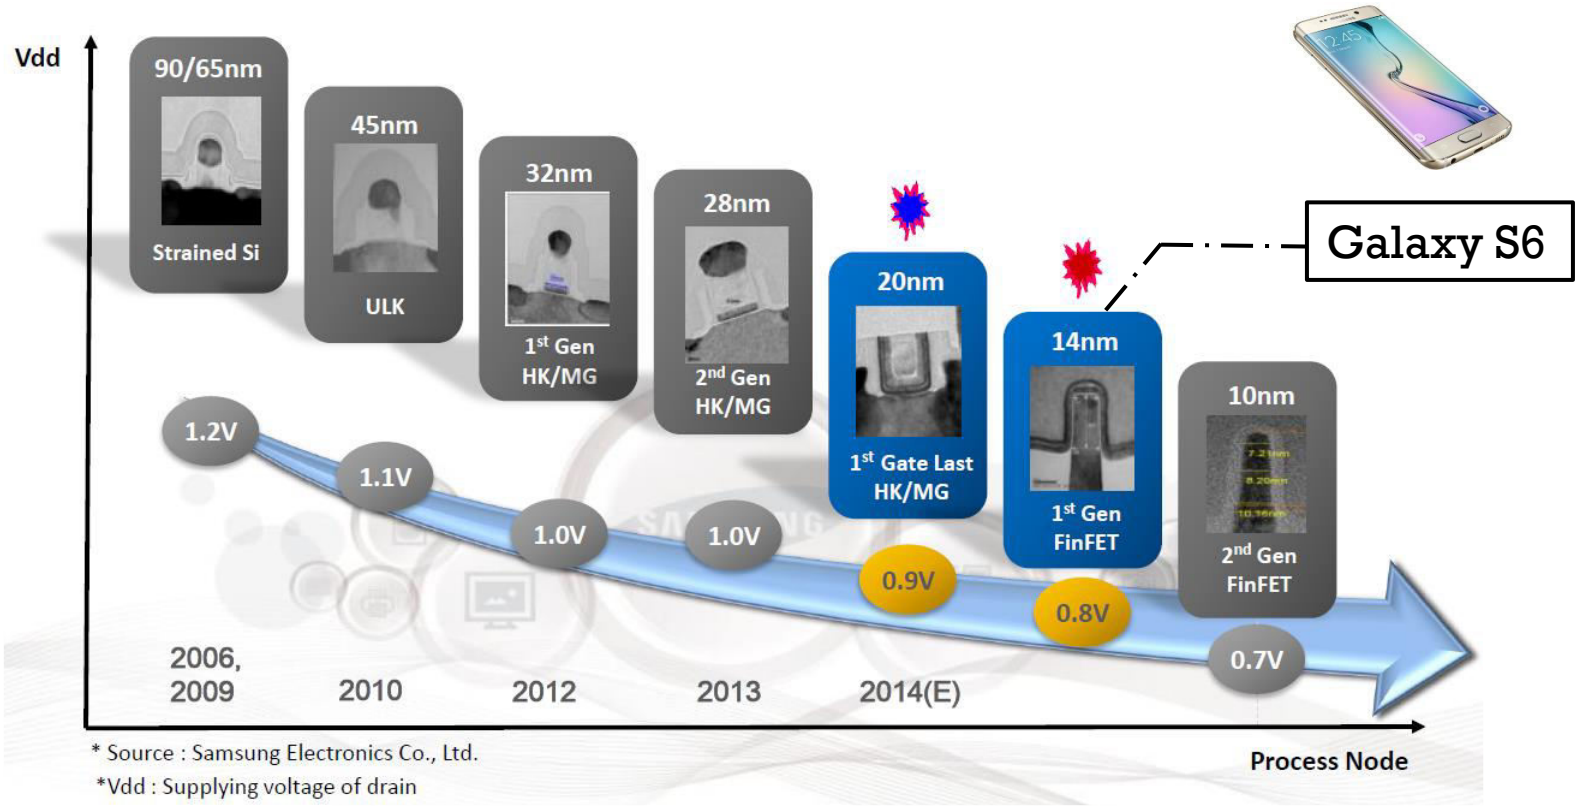

Samsung Technology Roadmap

The Galaxy S6 uses first generation FinFET technology

INTRODUCTION

Utilization Wall:
Dark Silicon's Effect on Multicore Scaling

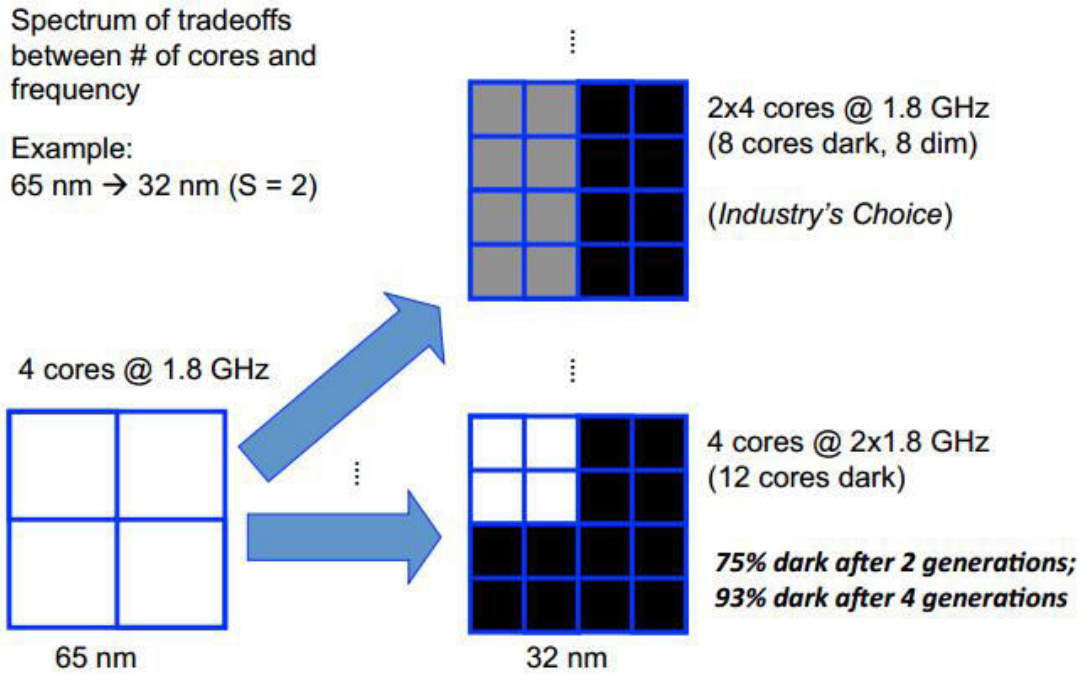

Expectations meets Reality.

INTRODUCTION

- Dark silicon constraints are enforced by switching off specific cores (grey) on the die.

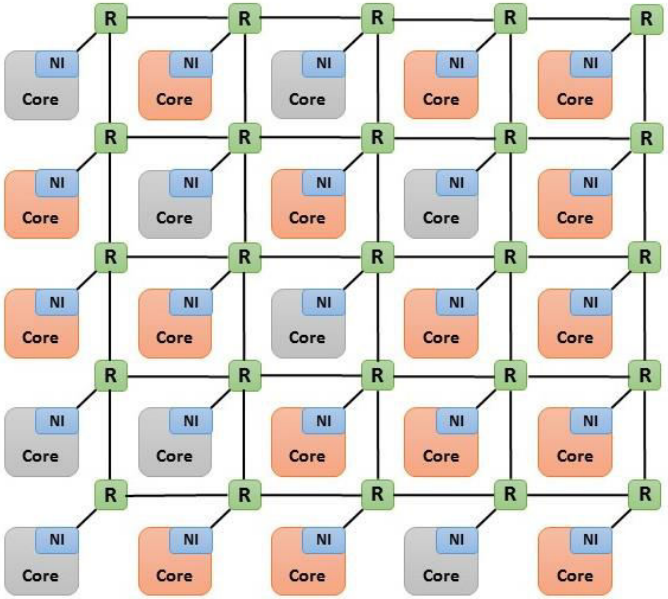

- Routers connected to cores are considered to be on to preserve connectivity in the routing framework.

The power constraint defines the number of cores to be disabled.

OUTLINE

- ✓ INTRODUCTION
- ✓ MOTIVATION
- OVERVIEW
- RELATED WORK
- PROBLEM STATEMENT
- METHODOLOGY
- RESULTS
- CONCLUSION AND FUTURE WORK

OVERVIEW - FINFETS

- SG Mode – The front and back gates, having same work functions, are shorted together.
- IG Mode – The back gates are subjected to a reverse bias low voltage to lower the leakage current.
- ASG Mode – The front and back gates are shorted like SG mode, but have different work functions.

FinFETs offer shorter delay, higher on vs off state current ratio and much lower leakage.

OVERVIEW – PROCESS VARIATIONS

- Systematic variations - lithographic constraints and aberrations.
- Random variations (line edge roughness) - random statistically fluctuating effects.
- The random variation component is expected to have greater impact at very low feature sizes.

For our work, we consider an equal impact of systematic and random variations.

OVERVIEW

- The power modeling of the NoC framework is performed by combined consideration of the compute and NoC components.
- The equation to model power in the NoC fabric is given by:

$$P_{Combined} = P_{Total}^{Compute} + P_{Total}^{NoC}$$

- The total power for the compute tiles is the combination of the leakage and dynamic power components.

$$P_{Total}^{Compute} = \sum_{i=1}^{C} P_{Li} + P_{Di}$$

OVERVIEW

- The total power for the a NoC router and connecting link is given by:

$$P_{bit} = P_{Switch} + P_{Buffer} + P_{Wire} + P_{Link}$$

$$P_{bit} = P_{Switch} + P_{Buffer} + P_{Link}$$

- The total power between two tiles t_a and t_b for a single bit, is given by:

$$P_{bit}(between\ t_a\ and\ t_b) = n * (P_{Switch} + P_{Buffer}) + (n-1) * P_{Link}$$

METHODOLOGY - INPUTS

- A 2D die with an irregular FinFET based 2D mesh NoC *G(T,L)* with *T* tiles and *L* links, with a percentage of failed links.
- An application core graph *G(V,E)* with *V* homogenous cores and *E* communication volumes between cores.
- A power map for the die components, constituting the overall distribution modeling WID (within-die) variations.
- A set of rated core frequencies *f*,, for the compute cores *V*.
- A set of link bandwidth values for the given set of active *L* links.

METHODOLOGY

inputs: ***set of core tile mappings with communication power values***

1: Generate initial random population P of individuals of size N, set crossover and mutation probabilities P_c and P_m and an initial prime number p for hash generation.
2: Evaluate the cost of the individuals for every communicating core-tile pair and store the initial best mapping.
3: for number of generations < max_generations
4: while size of new population P' < N
5: Select two members from P using tournament selection.
6: if ($P_{Xover} > P_c$), perform two point crossover for the members.
7: if ($P_{Mutation} > P_m$), perform random swap on the new members.
8: Generate hash values using p for new members and add to P', if there are no duplicates and update the best mapping.
9: end while
10: end for
11: Evaluate and return the best mapping.

output : ***core tile mapping with minimum communication power***

CUSTOM GENETIC ALGORITHM

METHODOLOGY

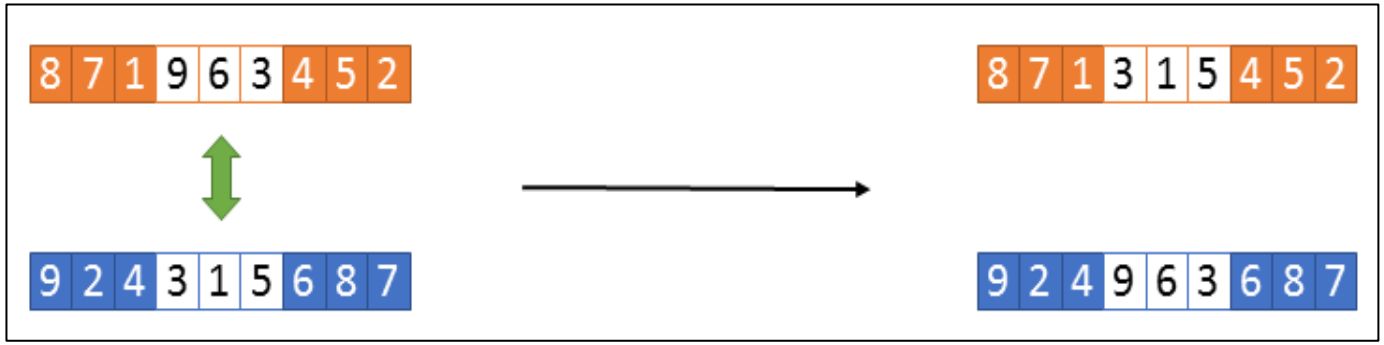

- The crossover operation for GA is performed as a integral step in the creation of the child genes.
- The conventional crossover selects crossover points and switches the values, between the points, for the parent genes.

Conventional crossover produces illegal mappings for ordered genes.

METHODOLOGY

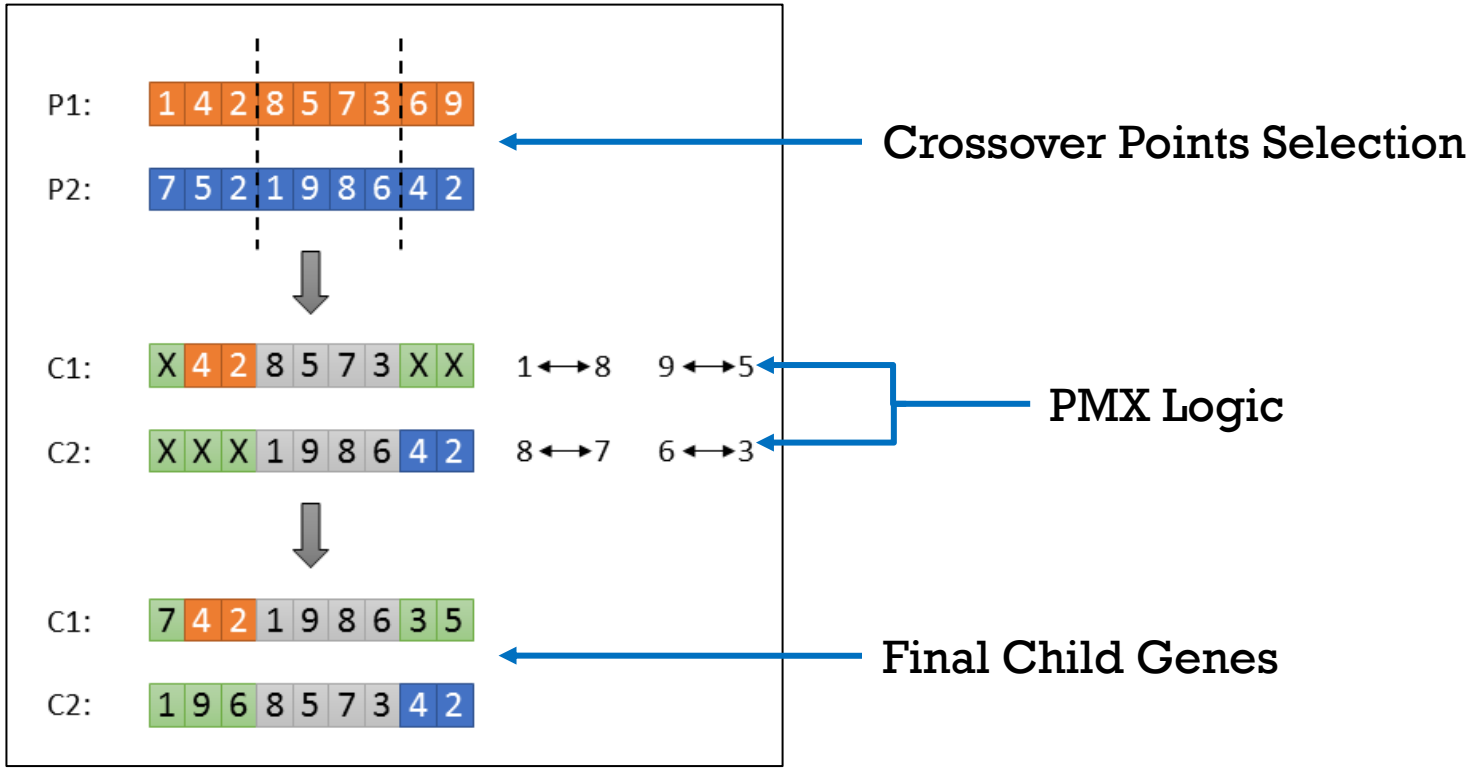

PMX CROSSOVER

PMX Crossover preserves the order of the child genes.

RESULTS – SIMULATION DETAILS

- We use the single core, superscalar, out-of-order ALPHA processor from McPAT-PVT, modeled using a 22 nm ASG FinFET model, with a core frequency of 1.2 GHz, for the compute cores.
- The CPU has 4 ALUs and 1 FPU per core.
- The L1 instruction and data cache capacities are 64 bits with the L2 cache capacity set to 1.75 MB. The cache is configured with a block size of 16 bytes, associativity as 8, access model as UCA and MOESI coherence protocol.
- The link power figures are generated using DSENT for the 22 nm technology node for a frequency of 1.2 GHz, with 64 bit data width and a link length of 1 mm.
- The router buffers are assumed to have a capacity of up to 48 flits, with each flit size being 64 bits.

RESULTS – SIMULATION DETAILS

- The link bandwidth constraints are empirically set at
  - 47 MBps for *cholesky_100* and *ocean_100*
  - 12.5 MBps for *blackscholes_100* and *fluidanimate_100*
  - 160 MBps for *dedup_100*
  - 134 MBps for *lu_100*
  - 155 MBps for *canneal_100*
  - 136 MBps for *fft_100*