ParaDIME: Parallel Distributed Infrastructure for Minimization of Energy

Oscar Palomar, Gulay Yalcin, Santhosh Kumar Rethinagiri,
Adrian Cristal, Osman S. Unsal and Gina Alioto
Barcelona Supercomputing Center

DMTM, Viena, 22-01-2014

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 318693

Motivation

- Call on *Minimising Energy Consumption of Computing to the Limit*
- Combine techniques, cooperation and sharing of information between layers (application, runtime, programming model, hardware, device)
  - e.g. annotations for reduced precision and high performance vs low power sections
- New elementary devices operating at the limits of minimum energy consumption

ParaDIME Consortium

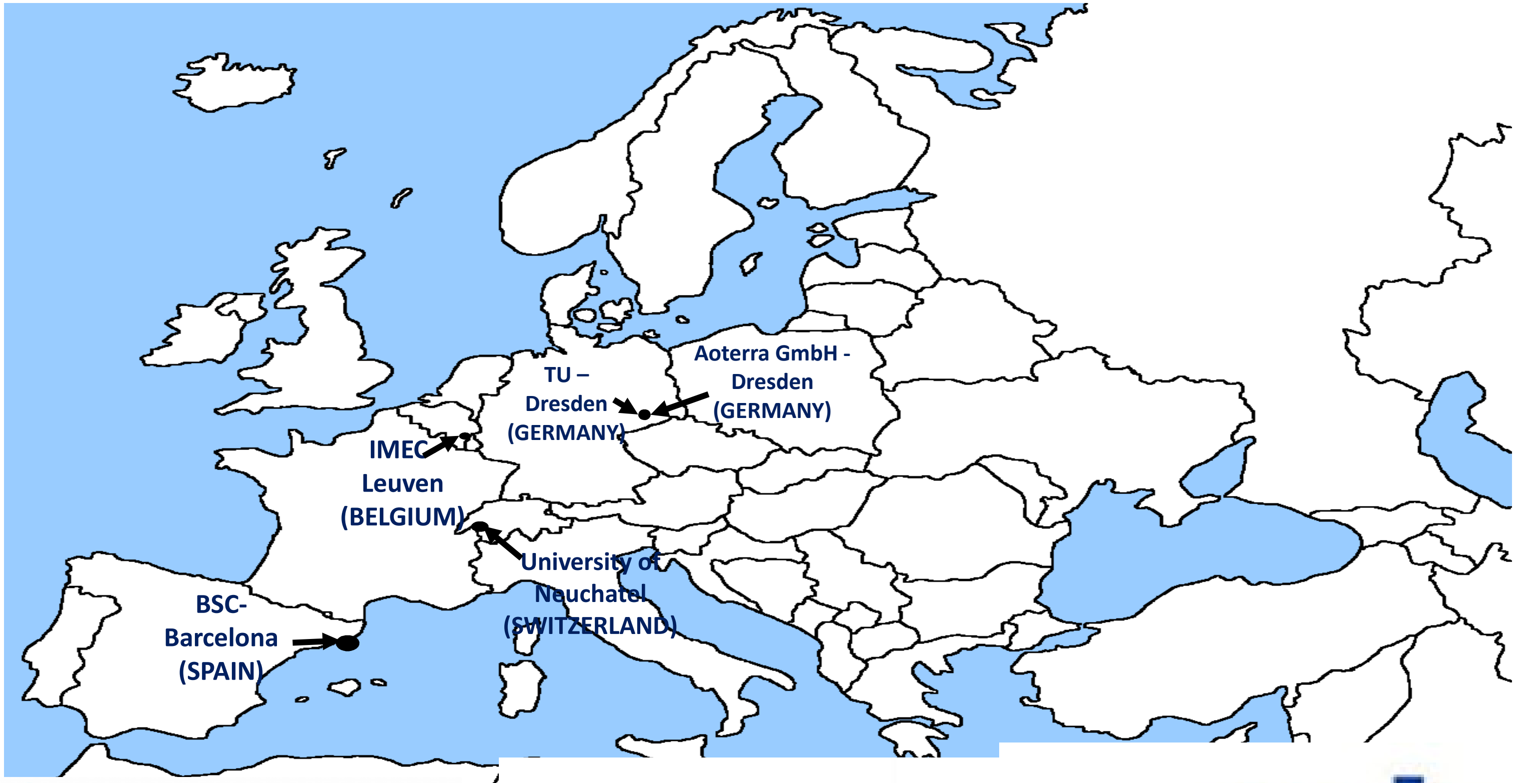

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

TECHNISCHE UNIVERSITÄT DRESDEN

Université de Neuchâtel unine

AOTERRA
GREENLOCALEFFICIENT

imec

PARA DIME

The ParaDIME stack

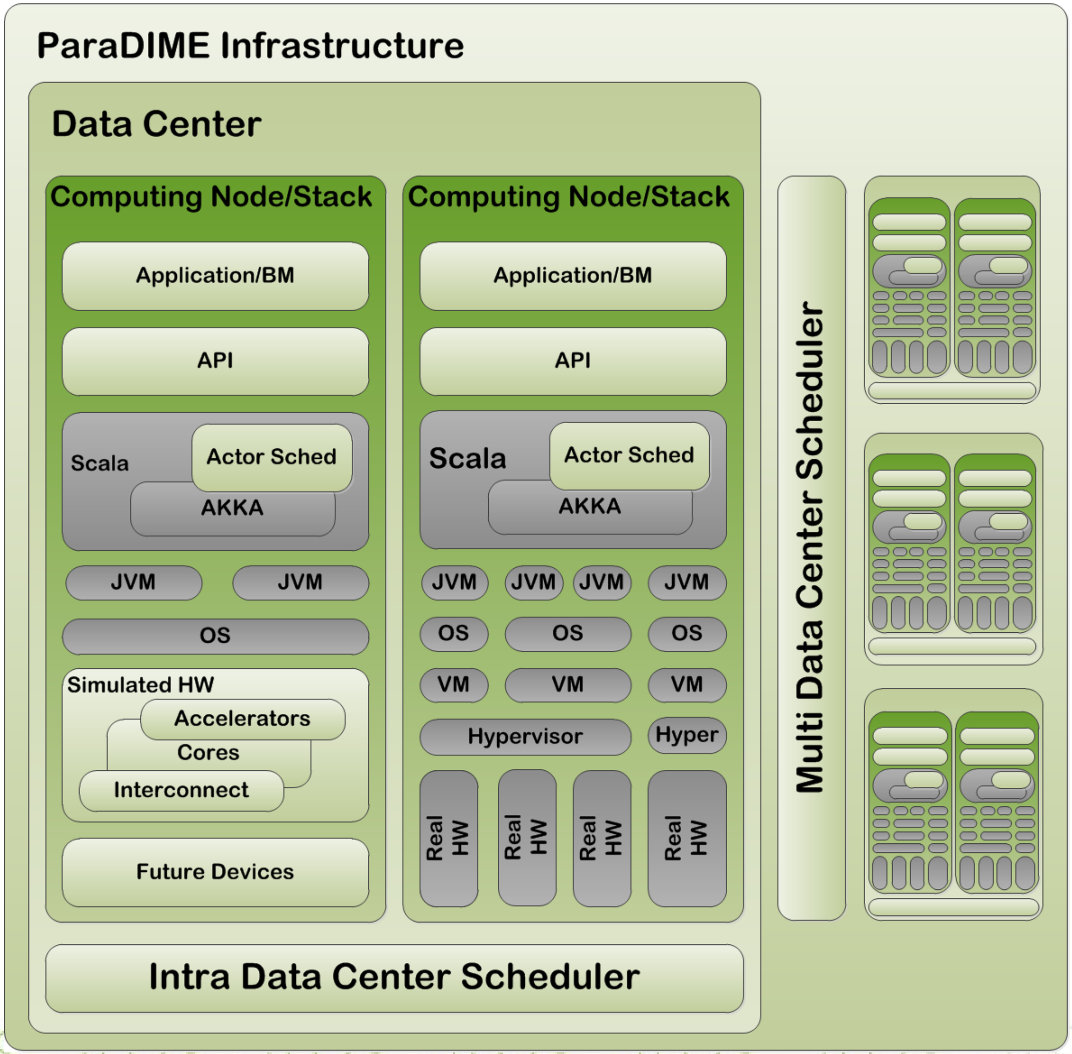

PARA DIME

OPERATION BELOW SAFE Vdd

Introduction

- Reducing the Vdd is a well-known technique for making trade-offs between performance and power
  - P=C*V2*f
  - Limited by safeguard bands to ensure correctness
  - Redundancy can be used to detect and correct errors but full replication incurs significant energy overhead
- Hardware transactional memory as its ability to keep a check-pointed state is very handy for error correction
- Replicated transactions for high error detection capability
  - Selectively replicate when high reliability is required to limit the overhead

Operation Below Safe Vdd

- Automatic HW lowering of Vdd
  - When it does not degrade performance (high cache miss rate)
- SW-guided
  - High-performance vs low-power annotation
  - More aggressive
- Errors detection and correction
  - Selective replication using transactional memory
  - Combine with approximate to reduce the overhead of replication

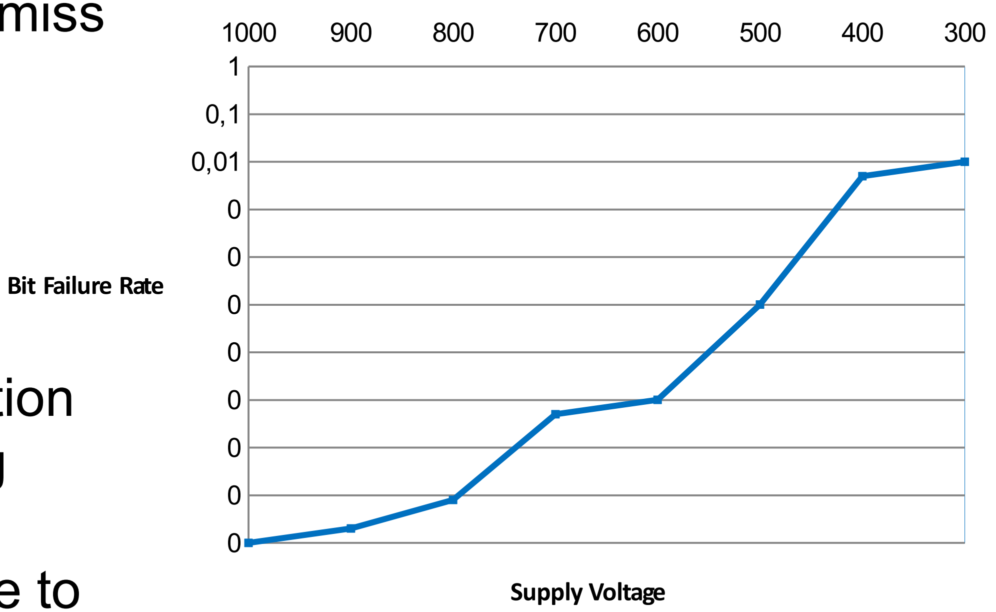

Error Detection Mechanisms

Method	Memory Overhead	Processing Overhead
DMR/TMR	High	High
Assertions/Invariants	Medium	High
Encoded Processing	Medium	High
Symptoms	Low	Low

Error probability

- Need to study error probability as a function of Vdd
- $PTXf = 1-(P_{ALU})^{size}$
- $E_{recovery} = E_{TX} * PTXf * (1/(1-PTXf))$
- FaulTM on gem5-x86
- SPLASH

Influence of Transaction Size

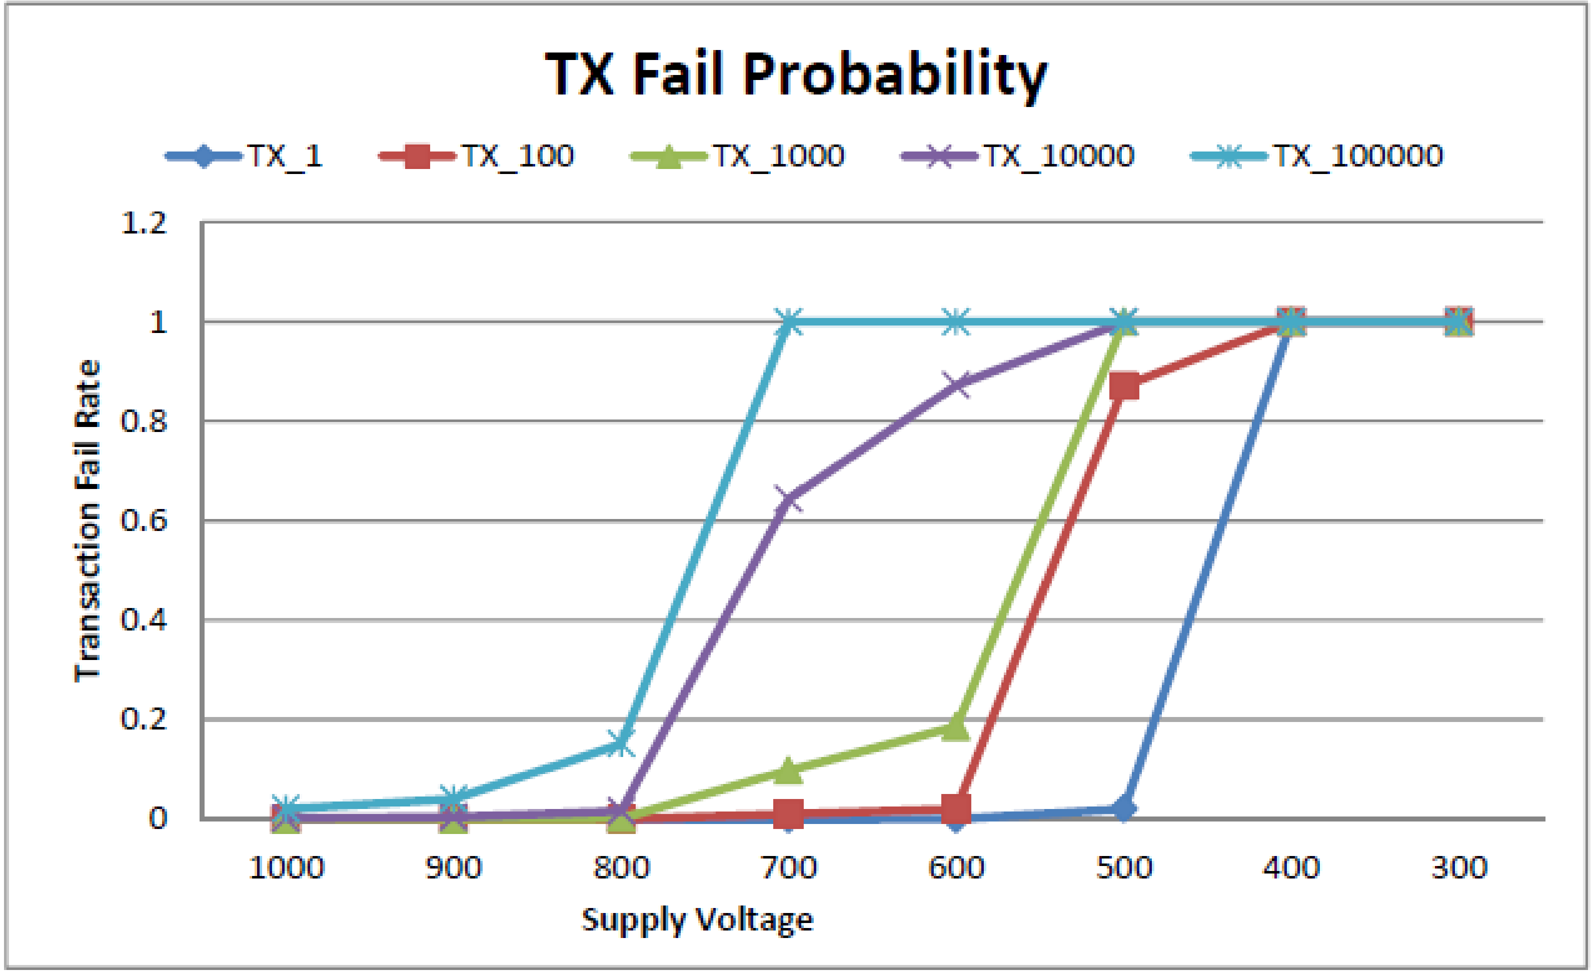

Energy Consumption

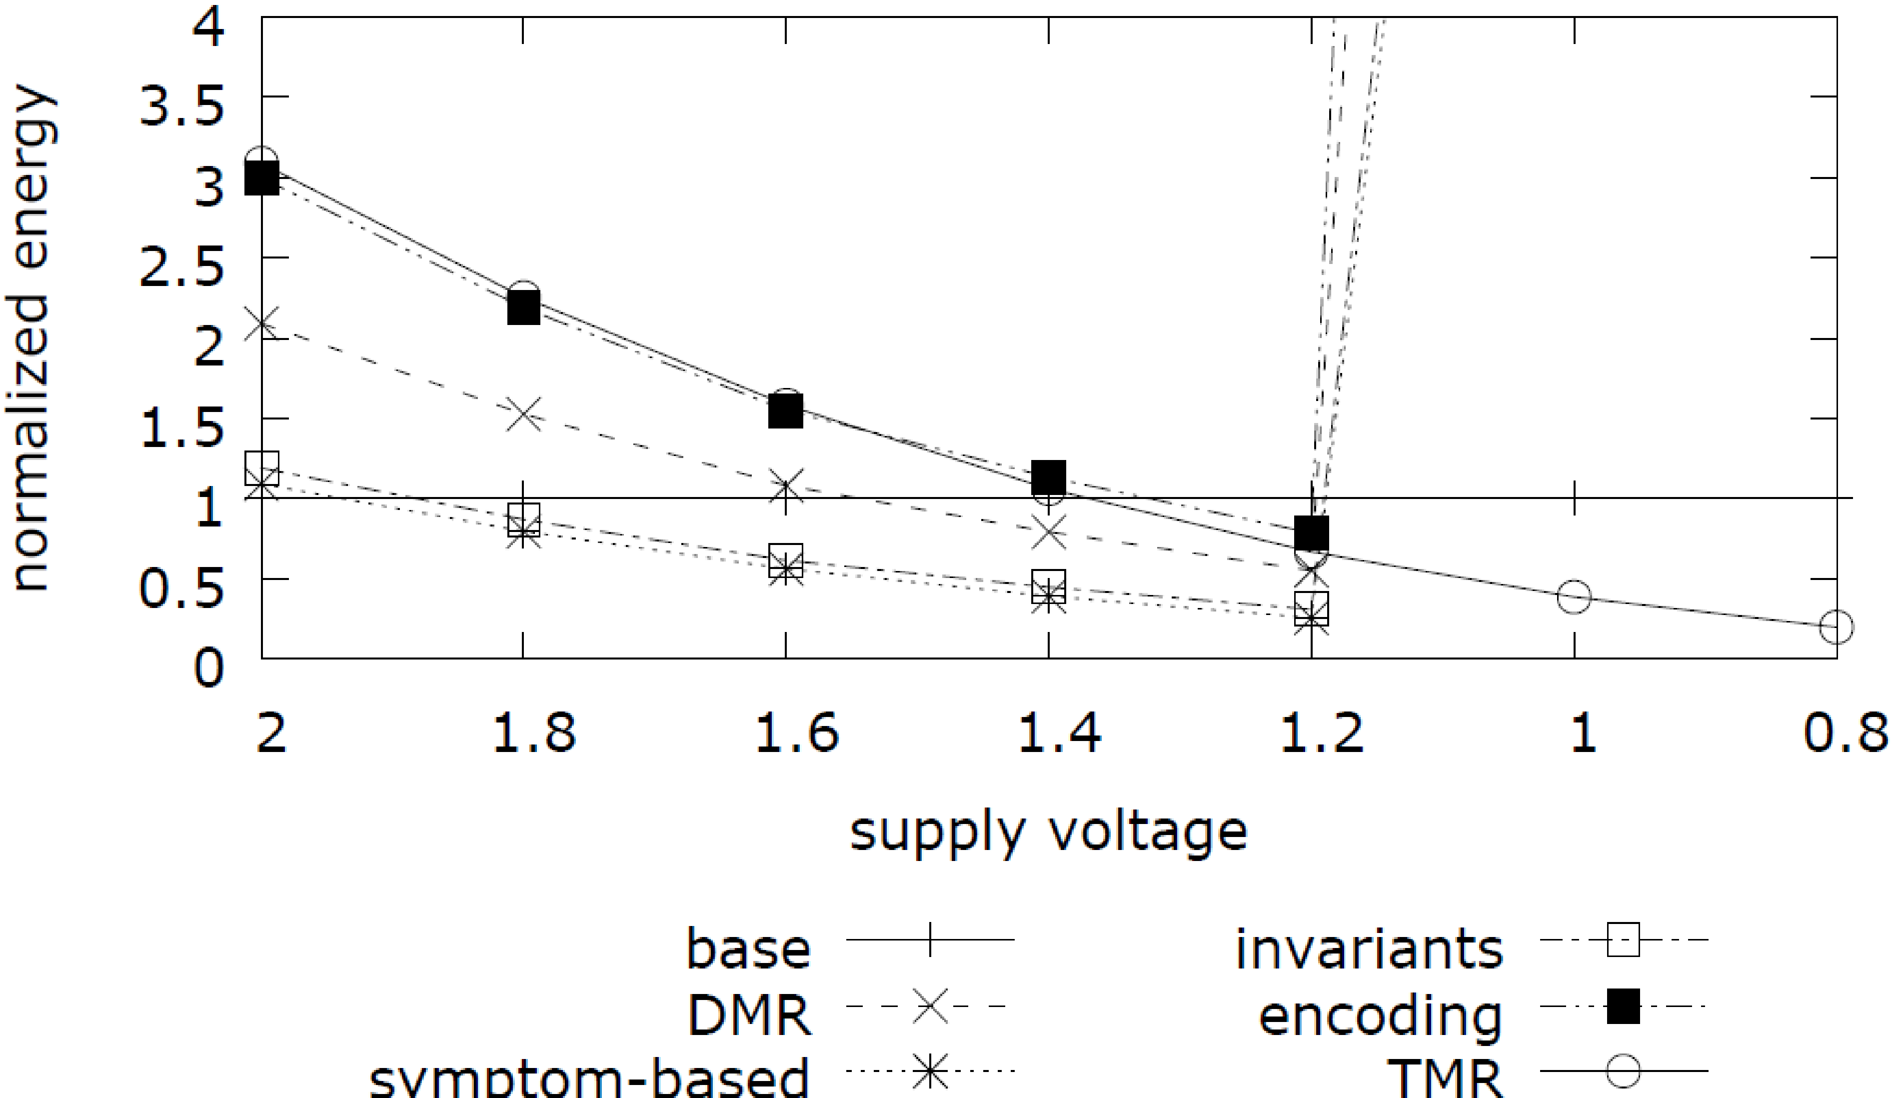

Energy Consumption

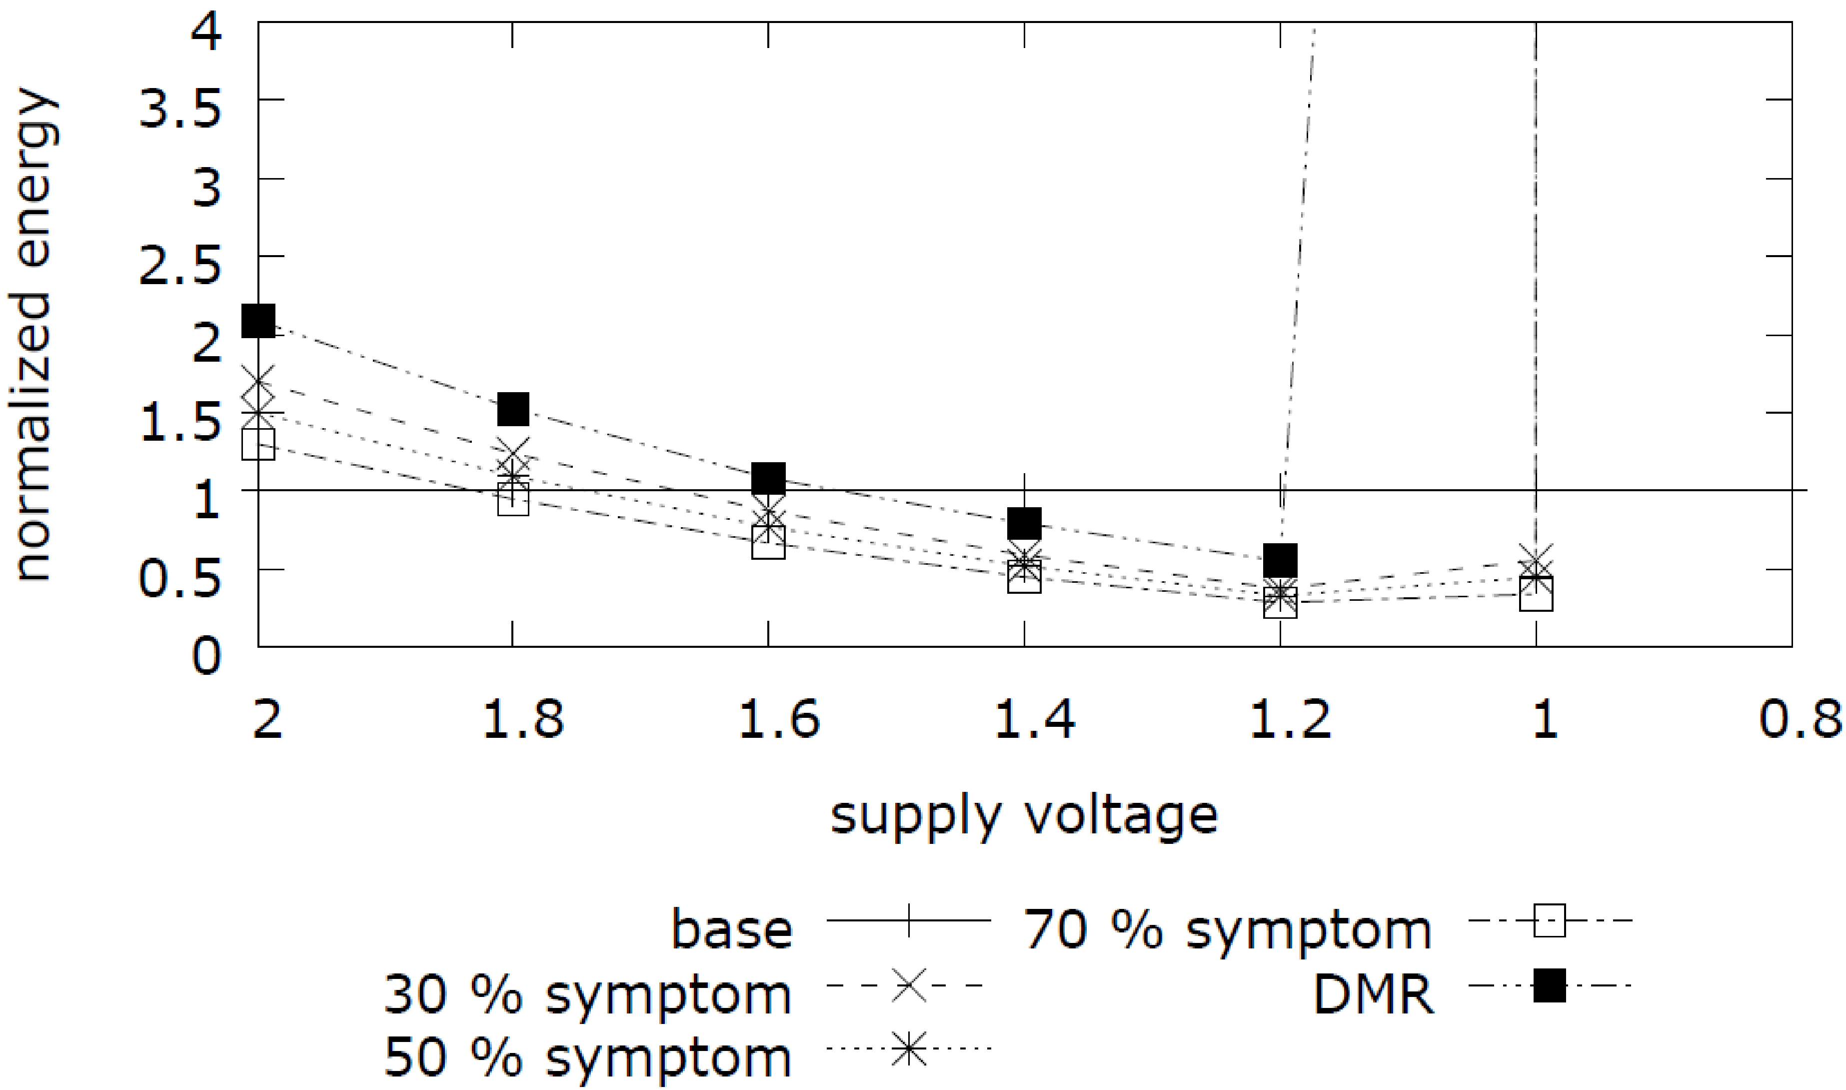

Error probability

- Need to study error probability as a function of Vdd
- Develop error models for future devices used in important processor blocks as a function of Vdd
- Near and Far Future devices
  - FinFet
  - III-V devices
- Probability of not meeting timing constraints at a given Vdd for basic blocks
- Transform into probability of error

ParaDIME: Parallel Distributed Infrastructure for Minimization of Energy

Oscar Palomar, Gulay Yalcin, Santhosh Kumar Rethinagiri,
Adrian Cristal, Osman S. Unsal and Gina Alioto
Barcelona Supercomputing Center
DMTM, Viena, 22-01-2014

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 318693

IMEC's Contribution

- Emerging devices

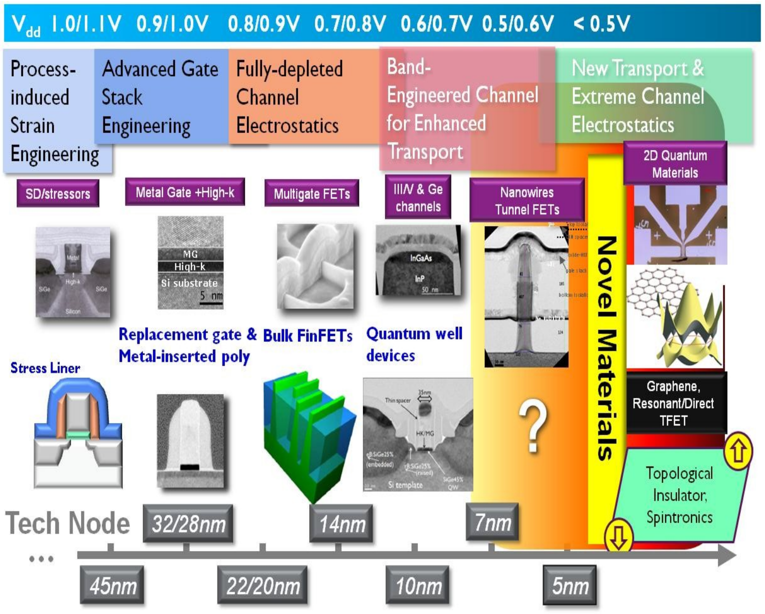

- Voltage limit

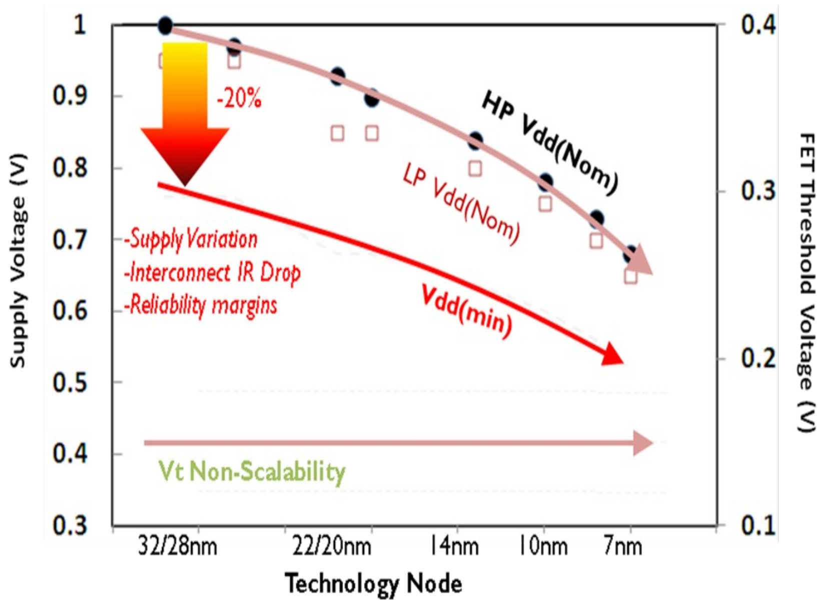

BSC's Contribution – Hardware Level

- Energy-Efficient Message Passing
  - Message passing microarchitecture
  - Message passing accelerator
  - Fast task switch
  - Task passing
- Operation Below Safe Vdd
  - Automatic HW lowering of Vdd
  - SW-guided (low-power annotation)
  - Error detection and correction

Approximate Computing
  - Reduced precision FP (annotation)
  - Reduced error detection / correction
- Heterogeneous Computing
  - Architectural level
  - Device level

TU-Dresden & AoTerra’s Contribution – (Runtime & Data Center Level)

- Fine grain scheduling
  - Schedule treads and processes
- Course grain scheduling
  - Peak load
- Source aware scheduling
  - Carbon footprint (Supply sources: renewable over non renewable)
- Energy-proportional computing
  - Switch-off machines when load is low
- Energy-efficient data storage and replication

University of Neuchatel's Contribution (Programming model)

- Efficient message passing
  - Actor Model (Scala & Akka)
- Error detection/recovery
  - Transactional memory for failure recovery
- Annotations/Metadata
  - Energy Profiles
  - Low-level annotations (e.g. safe sections for lowering Vdd, approximate data types)
- Applications
  - Future energy aware benchmarks
  - Real-world application