Performance and Fault Tolerance of Preconditioned Iterative Solvers on Low-Power ARM Architectures

José I. Aliaga¹ Sandra Catalán¹ Charalampos Chalios² Dimitrios S. Nikolopoulos² Enrique S. Quintana-Ortí¹

¹Departamento de Ingeniería y Ciencia de los Computadores, Universitat Jaume I, 12.071–Castellón de la Plana, (Spain) {aliaga, catalans, quintana}@uji.es

²School of EEECS, Queen's University of Belfast (United Kingdom)

cchalios01,d.nikolopoulos@qub.ac.uk>

September, 2015



Performance & Technology Evolution

- Although Moore's Law seems exhausted, it still holds today
 - Solutions to reach sub-10nm technology will give more years
- Dennard scaling broke down around 2005-2007
 - Doubling of the frequency transistor is no longer free
 - Multicores were the alternative to improve performances
 - Challenge: reduce power consump. & energy dissipation per *mm*²
- Dark Silicon and heterogeneous designs:
 - Only a fraction of silicon can be active at a moment
 - Use of specialized units reduce the energy consumption

2

Near Threshold Voltage Computing (NTVC)
 Allow to reduce the energy consumption





Performance and Fault Tolerance of ILUPACK on ARM 2

J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí



Performance & Technology Evolution

- Although Moore's Law seems exhausted, it still holds today
 - Solutions to reach sub-10nm technology will give more years
- Dennard scaling broke down around 2005-2007
 - Doubling of the frequency transistor is no longer free
 - Multicores were the alternative to improve performances
 - Challenge: reduce power consump. & energy dissipation per mm²
- Dark Silicon and heterogeneous designs:
 - Only a fraction of silicon can be active at a moment
 - Use of specialized units reduce the energy consumption
- Near Threshold Voltage Computing (NTVC)
 - Allow to reduce the energy consumption
 - Increase the error rates due to logic failures
 ault Tolerance of ILUPACK on ARM 2 J.Aliaga, S.Catalár

Performance and Fault Tolerance of ILUPACK on ARM

J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí



Performance & Technology Evolution

- Although Moore's Law seems exhausted, it still holds today
 - Solutions to reach sub-10nm technology will give more years
- Dennard scaling broke down around 2005-2007
 - Doubling of the frequency transistor is no longer free
 - Multicores were the alternative to improve performances
 - Challenge: reduce power consump. & energy dissipation per mm²
- Dark Silicon and heterogeneous designs:
 - Only a fraction of silicon can be active at a moment
 - Use of specialized units reduce the energy consumption
- Near Threshold Voltage Computing (NTVC)
 - Allow to reduce the energy consumption
 - Increase the error rates due to logic failures

Performance and Fault Tolerance of ILUPACK on ARM

J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Orti





Performance and Fault Tolerance of ILUPACK on ARM 2

J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí



Performance & Technology Evolution

- Although Moore's Law seems exhausted, it still holds today
 - Solutions to reach sub-10nm technology will give more years
- Dennard scaling broke down around 2005-2007
 - Doubling of the frequency transistor is no longer free
 - Multicores were the alternative to improve performances
 - Challenge: reduce power consump. & energy dissipation per mm²
- Dark Silicon and heterogeneous designs:
 - Only a fraction of silicon can be active at a moment
 - Use of specialized units reduce the energy consumption
- Near Threshold Voltage Computing (NTVC)
 - Allow to reduce the energy consumption
 - Increase the error rates due to logic failures



High Performance Benchmarks

- LINPACK benchmark solves a dense system of linear equations
 - Uses to build TOP500 and GREEN500 lists
 - Gives a good correction of peak performance
- Currently, many applications are defined by sparse matrices
 - The LINPACK performances are useless for real applications
- High Performance Conjugate Gradients (HPCG) benchmark
 - Solves a sparse linear system by using a iterative method
 - Computations and data access patterns more commonly found in applications



High Performance Benchmarks

- LINPACK benchmark solves a dense system of linear equations
 - Uses to build TOP500 and GREEN500 lists
 - Gives a good correction of peak performance
- Currently, many applications are defined by sparse matrices
 - The LINPACK performances are useless for real applications
- High Performance Conjugate Gradients (HPCG) benchmark
 - Solves a sparse linear system by using a iterative method
 - Computations and data access patterns more commonly found in applications



High Performance Benchmarks

- LINPACK benchmark solves a dense system of linear equations
 - Uses to build TOP500 and GREEN500 lists
 - Gives a good correction of peak performance
- Currently, many applications are defined by sparse matrices
 - The LINPACK performances are useless for real applications
- High Performance Conjugate Gradients (HPCG) benchmark
 - Solves a sparse linear system by using a iterative method
 - Computations and data access patterns more commonly found in applications

Objective: Analysis of ILUPACK on low-power ARM multicore

- ILUPACK (Incomplete LU PACKage)
 - Computes a Inverse-based Multilevel ILU preconditioner
 - Applies it to accelerate Krylov-subspace methods
 - In this talk, the task-parallel version is considered
 - Its behaviour is similar to HPCG benchmark
 - Data access patterns
 - Arithmetics-to-memory operations ratios

- Better suited to multicore processor than HPCG
- For the low-power multicore ARM processors, we focus on
 - Scalability
 - Energy efficiency
 - Fault resilience



ľi

Objective: Analysis of ILUPACK on low-power ARM multicore

- ILUPACK (Incomplete LU PACKage)
 - Computes a Inverse-based Multilevel ILU preconditioner
 - Applies it to accelerate Krylov-subspace methods
 - In this talk, the task-parallel version is considered
 - Its behaviour is similar to HPCG benchmark
 - Data access patterns
 - Arithmetics-to-memory operations ratios

- Better suited to multicore processor than HPCG
- For the low-power multicore ARM processors, we focus on
 - Scalability
 - Energy efficiency
 - Fault resilience

j'i

Objective: Analysis of ILUPACK on low-power ARM multicore

- ILUPACK (Incomplete LU PACKage)
 - Computes a Inverse-based Multilevel ILU preconditioner
 - Applies it to accelerate Krylov-subspace methods
 - In this talk, the task-parallel version is considered
 - Its behaviour is similar to HPCG benchmark
 - Data access patterns
 - Arithmetics-to-memory operations ratios

- Better suited to multicore processor than HPCG
- For the low-power multicore ARM processors, we focus on
 - Scalability
 - Energy efficiency
 - Fault resilience

j'i

Objective: Analysis of ILUPACK on low-power ARM multicore

- ILUPACK (Incomplete LU PACKage)
 - Computes a Inverse-based Multilevel ILU preconditioner
 - Applies it to accelerate Krylov-subspace methods
 - In this talk, the task-parallel version is considered
 - Its behaviour is similar to HPCG benchmark
 - Data access patterns
 - Arithmetics-to-memory operations ratios

- Better suited to multicore processor than HPCG
- For the low-power multicore ARM processors, we focus on
 - Scalability
 - Energy efficiency
 - Fault resilience

Outline



- 1 Motivation and Introduction
- 2 Task-parallel version of ILUPACK
 - Preconditioner computation
 - Iterative solver
- 3 Setup
- 4 Frequency-Voltage Scaling
 - Computational Performance
 - Power and Energy Efficiency
- 5 Fault tolerance and NTVC
- 6 Conclusions

6

Preconditioner computation Iterative solver

Outline



Motivation and Introduction

- 2 Task-parallel version of ILUPACK
 - Preconditioner computation
 - Iterative solver

3 Setup

- Frequency-Voltage Scaling
 Computational Performance
 - Power and Energy Efficiency
- 5 Fault tolerance and NTVC

6 Conclusions

Preconditioner computation Iterative solver

Inverse-based Multilevel IC factorization

Introducing the inverse-based approach into the IC factorization:

• The (root-free) sparse Cholesky decomposition computes

$$A = LDL^T$$

• In the Incomplete Cholesky (IC) decomp. A is factorized as

$$\mathbf{A} = \tilde{\mathbf{L}}\tilde{\mathbf{D}}\tilde{\mathbf{L}}^{\mathsf{T}} + \mathbf{E}$$

where $E \in \mathbb{R}^{n,n}$ is a "small" perturbation matrix consisting of those entries dropped during factorization

M = *L̃D̃L*^T is applied to the original system *Ax* = *b* in order to accelerate the convergence of the PCG solver

$$(\tilde{D}^{-1/2}\tilde{L}^{-1})A(\tilde{L}^{-T}\tilde{D}^{-1/2})$$



Preconditioner computation Iterative solver

Inverse-based Multilevel IC factorization

Introducing the inverse-based approach into the IC factorization:

- Inverse-based IC factorizations construct $M = \tilde{L}\tilde{D}\tilde{L}^{T}$, $\|\tilde{L}^{-1}\| \leq \kappa$
 - $\kappa = 5$ or $\kappa = 10$ are good choices for many problems
 - Pivoting and **multilevel** methods are required in general



Preconditioner computation Iterative solver

Inverse-based Multilevel IC factorization

Introducing the inverse-based approach into the IC factorization:

- Inverse-based IC factorizations construct $M = \tilde{L}\tilde{D}\tilde{L}^{T}$, $\|\tilde{L}^{-1}\| \leq \kappa$
 - $\kappa = 5$ or $\kappa = 10$ are good choices for many problems
 - Pivoting and multilevel methods are required in general







Preconditioner computation Iterative solver

Inverse-based Multilevel IC factorization

MIC factorization of five-point matrix arising from Laplace PDE discretization



 $\kappa = 5$ and $\tau = 10^{-3}$ leads to a MIC factorization with 5 levels



Preconditioner computation Iterative solver

Parallelization of ILUPACK

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

- Nested dissection (ND) reveals task-level parallelism
- METIS and SCOTCH libraries implement multilevel ND





Preconditioner computation Iterative solver

Parallelization of ILUPACK



The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• The task dependency tree represents the behaviour of the computation



- Each node refers to a block to be factorized, while the edges define the data dependency.
- The nodes without any dependency can be factorized in parallel.
- The non-leaf nodes have to wait until its children have finalized.

Preconditioner computation Iterative solver

Parallelization of ILUPACK



The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• The task dependency tree represents the behaviour of the computation



 Each node refers to a block to be factorized, while the edges define the data dependency.

• The nodes without any dependency can be factorized in parallel.

• The non-leaf nodes have to wait until its children have finalized.

Preconditioner computation Iterative solver

Parallelization of ILUPACK



The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• The task dependency tree represents the behaviour of the computation



- Each node refers to a block to be factorized, while the edges define the data dependency.
- The nodes without any dependency can be factorized in parallel.
- The non-leaf nodes have to wait until its children have finalized.

Preconditioner computation Iterative solver



inverse-based Multilevel IC factorization



Parallel inverse-based Multilevel IC factorization

Performance and Fault Tolerance of ILUPACK on ARM

11 J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí

Preconditioner computation Iterative solver

Parallel vs Sequential Inverse-based MIC factoriz.

Factorization of five-point matrix arising from Laplace PDE discretization



Performance and Fault Tolerance of ILUPACK on ARM

J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí

バ

Preconditioner computation Iterative solver

Parallel vs Sequential Inverse-based MIC factoriz.

Factorization of five-point matrix arising from Laplace PDE discretization



Performance and Fault Tolerance of ILUPACK on ARM

J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí

Ĵi

Preconditioner computation Iterative solver

Preconditioned Conjugate Gradient

Initialize k,
$$x_0, r_0, z_0, d_0, \beta_0, \tau_0$$

while $((k < k_{max})\&\&(\tau_k > \tau_{min}))$
 $w_k := Ad_k$
 $\gamma_k := d_k^T w_k$
 $\rho_k := \beta_k / \gamma_k$
 $x_{k+1} := r_k - \rho_k w_k$
 $z_{k+1} := r_k - \rho_k w_k$
 $z_{k+1} := \beta_k - \beta_k$
 $\beta_{k+1} := r_{k+1}^T z_{k+1}$
 $\alpha_k := \beta_k$
 $\beta_{k+1} := z_{k+1} + \alpha_k d_k$
 $\tau_{k+1} := ||r_{k+1}||_2$
 $k := k + 1$
end while





Preconditioner computation Iterative solver

Preconditioned Conjugate Gradient

Initialize $k, x_0, r_0, z_0, d_0, \beta_0, \tau_0$ while $((k < k_{max})\&\&(\tau_k > \tau_{min}))$ $W_{k} := Ad_{k}$ SPMV $\gamma_k := \boldsymbol{d}_k^T \boldsymbol{w}_k$ DOT $\rho_{\mathbf{k}} := \beta_{\mathbf{k}} / \gamma_{\mathbf{k}}$ $\mathbf{x}_{k+1} := \mathbf{x}_k + \rho_k \mathbf{d}_k$ $\mathbf{r}_{k+1} := \mathbf{r}_k - \rho_k \mathbf{W}_k$ $Z_{k\perp 1} := M^{-1} r_{k\perp 1}$ $\alpha_{\mathbf{k}} := \beta_{\mathbf{k}}$ $\beta_{k+1} := r_{k+1}^T z_{k+1}$ DOT $\alpha_{k} := \beta_{k+1} / \alpha_{k}$ $d_{k+1} := z_{k+1} + \alpha_k d_k$ $\tau_{k+1} := ||r_{k+1}||_2$ k := k + 1

AXPY AXPY Preconditioning

AXPY-like 2-norm

end while





while $((k < k_{max})\&\&(\tau_k > \tau_{min}))$ $W_{k} := Ad_{k}$ $\gamma_k := \boldsymbol{d}_k^T \boldsymbol{w}_k$ $\rho_{\mathbf{k}} := \beta_{\mathbf{k}} / \gamma_{\mathbf{k}}$ $\mathbf{X}_{k+1} := \mathbf{X}_k + \rho_k \mathbf{d}_k$ $\mathbf{r}_{k+1} := \mathbf{r}_k - \rho_k \mathbf{W}_k$ $z_{k+1} := M^{-1}r_{k+1}$ $\alpha_{\mathbf{k}} := \beta_{\mathbf{k}}$ $\beta_{k+1} := r_{k+1}^T z_{k+1}$ $\alpha_{k} := \beta_{k+1} / \alpha_{k}$ $d_{k+1} := z_{k+1} + \alpha_k d_k$ $\tau_{k+1} := ||r_{k+1}||_2$ k := k + 1end while

 $\begin{array}{l} \text{SPMV} \rightarrow \text{on leaves} \\ \text{DOT} \rightarrow \text{on leaves} + \text{reduction} \end{array}$

 $\begin{array}{l} \text{AXPY} \ \rightarrow \text{ on leaves} \\ \text{AXPY} \ \rightarrow \text{ on leaves} \\ \text{Preconditioning} \ \rightarrow \text{cross. up} + \text{down} \end{array}$

 $\text{DOT } \rightarrow \text{on leaves} + \text{reduction}$

Preconditioner computation Iterative solver

Preconditioner Application and/or Reduction

Preconditioner Forward Resolution (and/or Reduction Step):



Preconditioner Backward Resolution (and/or Copy Result):





Outline



- Motivation and Introduction
- Task-parallel version of ILUPACK
 - Preconditioner computation
 - Iterative solver

3 Setup

- Frequency-Voltage Scaling
 Computational Performance
 - Power and Energy Efficiency
- 5 Fault tolerance and NTVC

6 Conclusions

Hardware Platform



- Exynos5 Odroid-XU development board assembled by Samsung
 - Single big.LITTLE system-on-chip (SoC)
 - ARM Cortex-A15 quad-core cluster
 - ARM Cortex-A7 quad-core cluster
 - The frequency reported by cpufreq driver determines the active cluster
 - 2 Gbytes of DDR3 RAM
- We target on low-power scenarios
 - Only consider the ARM Cortex-A7 cores
 - Each core has 32+32 (data+instruction)-KByte L1 cache
 - Each core shares a 512 KByte-L2 cache with its siblings
 - Frequencies $\in \{250, 300, 350, \dots, 600\}$ MHz

Hardware Platform



- Exynos5 Odroid-XU development board assembled by Samsung
 - Single big.LITTLE system-on-chip (SoC)
 - ARM Cortex-A15 quad-core cluster
 - ARM Cortex-A7 quad-core cluster
 - The frequency reported by cpufreq driver determines the active cluster
 - 2 Gbytes of DDR3 RAM
- We target on low-power scenarios
 - Only consider the ARM Cortex-A7 cores
 - Each core has 32+32 (data+instruction)-KByte L1 cache
 - Each core shares a 512 KByte-L2 cache with its siblings
 - Frequencies $\in \{250, 300, 350, \ldots, 600\}$ MHz

Experiments



- Features of the parallel implementation:
 - Compiler: gcc 4.8.1 with the appropriate optimization flags
 - IEEE double precision arithmetic
 - Task tree with 8 leaves
 - Only the PCG results are reported
- Benchmark Matrix:
 - Finite difference discretization of a 3D Laplace problem
 - The grid size determines the size of the matrix
 - *n*=1,000,000 and *nnz*=6,940,000

Experiments



- Features of the parallel implementation:
 - Compiler: gcc 4.8.1 with the appropriate optimization flags
 - IEEE double precision arithmetic
 - Task tree with 8 leaves
 - Only the PCG results are reported
- Benchmark Matrix:
 - Finite difference discretization of a 3D Laplace problem
 - The grid size determines the size of the matrix
 - *n*=1,000,000 and *nnz*=6,940,000

Computational Performance Power and Energy Efficiency

Outline



Motivation and Introduction

- 2 Task-parallel version of ILUPACK
 - Preconditioner computation
 - Iterative solver

3 Setup

- Frequency-Voltage Scaling
 - Computational Performance
 - Power and Energy Efficiency
- 5 Fault tolerance and NTVC

6 Conclusions

Computational Performance Power and Energy Efficiency

First experiment



- Analyze the scalability of the solver:
 - Changing the number of cores and the frequency
 - The speed-up is computed with respect to the sequential code
 - The results show fair speed-ups
 - The speed-up are mostly independent of the operating frequency
- Features of the frequency-voltage scaling analysis:
 - Execution for different frequencies on 4 Cortex-A7 cores
 - Several metrics:
 - Execution time
 - Power and Energy, considering for both Total: ⇒ Exynos5 SoC plus the memory DIMMs CPU: ⇒ Cortex-A7 cluster
 - Polynomial regression is applied to interpolate the results

Computational Performance Power and Energy Efficiency

First experiment



Task parallel ILUPACK for different frequencies 250 MHz 3.5 400 MHz 600 MHz 3.0 2.5 Speed-up 1.5 1.0 0.5 2 1 4 Number of Cortex-A7 cores

Performance and Fault Tolerance of ILUPACK on ARM 19 J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí

Computational Performance Power and Energy Efficiency

First experiment



- Analyze the scalability of the solver:
 - Changing the number of cores and the frequency
 - The speed-up is computed with respect to the sequential code
 - The results show fair speed-ups
 - The speed-up are mostly independent of the operating frequency
- Features of the frequency-voltage scaling analysis:
 - Execution for different frequencies on 4 Cortex-A7 cores
 - Several metrics:
 - Execution time
 - Power and Energy, considering for both Total: ⇒ Exynos5 SoC plus the memory DIMMs CPU: ⇒ Cortex-A7 cluster
 - Polynomial regression is applied to interpolate the results

Computational Performance Power and Energy Efficiency

Analysis for Computational Performance

- The time and the frequency are usually inversely proportional
- To be more precise, some features have to be analyzed
 - To know if the code is compute- or memory-bounded
 - The effect of frequency on the memory bandwidth
- In our experiments,
 - ILUPACK combines compute- and memory-bounded operations
 - In Cortex-A7 cluster the frequency governs the memory bandwidth
- A quadratic polynomial as a function of the frequency works

$$T(f) = 5.80\text{E-4} f^2 - 7.33\text{E-1} f + 3.27\text{E+2} s$$
$$I - r^2 = 6.69\text{E-3}$$

Computational Performance Power and Energy Efficiency

Analysis for Computational Performance

- The time and the frequency are usually inversely proportional
- To be more precise, some features have to be analyzed
 - To know if the code is compute- or memory-bounded
 - The effect of frequency on the memory bandwidth
- In our experiments,
 - ILUPACK combines compute- and memory-bounded operations
 - In Cortex-A7 cluster the frequency governs the memory bandwidth
- A quadratic polynomial as a function of the frequency works

 $T(f) = 5.80E-4 f^2 - 7.33E-1 f + 3.27E+2 s$ $1 - r^2 = 6.69E-3$



Computational Performance Power and Energy Efficiency

Analysis for Computational Performance



Task-parallel ILUPACK on 4 Cortex-A7 cores



Performance and Fault Tolerance of ILUPACK on ARM 20 J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí

Computational Performance Power and Energy Efficiency

Analysis for Computational Performance

- The time and the frequency are usually inversely proportional
- To be more precise, some features have to be analyzed
 - To know if the code is compute- or memory-bounded
 - The effect of frequency on the memory bandwidth
- In our experiments,
 - ILUPACK combines compute- and memory-bounded operations
 - In Cortex-A7 cluster the frequency governs the memory bandwidth
- A quadratic polynomial as a function of the frequency works

$$T(f) = 5.80\text{E-4} f^2 - 7.33\text{E-1} f + 3.27\text{E+2} s$$

1 - r2 = 6.69E-3



Computational Performance Power and Energy Efficiency

Analysis for power and Energy Efficiency

• The relationship between power and frequency is cubic

$$P \propto V^2 * f$$
, $V \propto f \Rightarrow P \propto f^3$, $P \propto V^3$

obtained that

$$P_T(f) = 1.46E-9 f^3 + 2.76E-7 f^2 - 7.60E-5 f + 2.18E-1 W$$

1 - r2 = 5.72E-4

• The energy can be estimated by a quadratic polynomial:

$$E_T(f) = 2.43\text{E-}4 f^2 - 1.70\text{E-}1 f + 7.04\text{E+}1$$

1 - r2 = 7.36E-3

• The optimal configuration is different for total and for cluster: $E_T^{opt} = E_T(350)$, $E_{AT}^{opt} = E_{AT}(250) \Rightarrow f_{opt} = 300$



パ

Computational Performance Power and Energy Efficiency

Analysis for power and Energy Efficiency



50 100 150 200 250 300 350 400 450 500 550 600 Frequency (MHz)

Computational Performance Power and Energy Efficiency

Analysis for power and Energy Efficiency

• The relationship between power and frequency is cubic

$$P \propto V^2 * f$$
, $V \propto f \Rightarrow P \propto f^3$, $P \propto V^3$

obtained that

$$P_T(f) = 1.46E-9 f^3 + 2.76E-7 f^2 - 7.60E-5 f + 2.18E-1 W$$

1 - r2 = 5.72E-4

• The energy can be estimated by a quadratic polynomial:

$$E_T(f) = 2.43\text{E-}4 f^2 - 1.70\text{E-}1 f + 7.04\text{E+}1$$

1 - r2 = 7.36E-3

• The optimal configuration is different for total and for cluster: $E_T^{opt} = E_T(350)$, $E_{A7}^{opt} = E_{A7}(250) \Rightarrow f_{opt} = 300$



Computational Performance Power and Energy Efficiency

Analysis for power and Energy Efficiency



Task-parallel ILUPACK on 4 Cortex-A7 cores



Performance and Fault Tolerance of ILUPACK on ARM 21 J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí

Computational Performance Power and Energy Efficiency

Analysis for power and Energy Efficiency

• The relationship between power and frequency is cubic

$$P \propto V^2 * f$$
, $V \propto f \Rightarrow P \propto f^3$, $P \propto V^3$

obtained that

$$P_T(f) = 1.46E-9 f^3 + 2.76E-7 f^2 - 7.60E-5 f + 2.18E-1 W$$

1 - r2 = 5.72E-4

• The energy can be estimated by a quadratic polynomial:

$$E_T(f) = 2.43\text{E-}4 f^2 - 1.70\text{E-}1 f + 7.04\text{E+}1$$

1 - r2 = 7.36E-3

• The optimal configuration is different for total and for cluster:

$$E_T^{opt} = E_T(350) , \ E_{A7}^{opt} = E_{A7}(250) \Rightarrow f_{opt} = 300$$



Computational Performance Power and Energy Efficiency

Iso-power configurations



- Not always maximum frequency (*f_M*) implies best performances
- The power consumption of the optimal frequency is smaller
- Iso-power determines how many systems working at f_{opt} dissipate the same power than one system working at f_M

 $P_{\rm T}(f_{\rm M})/P_{\rm T}(f_{\rm opt}) = 0.587/0.259 = 2.26 \times P_{\rm A7}(f_{\rm M})/P_{\rm A7}(f_{\rm opt}) = 0.438/0.137 = 3.19 \times$

Given that the fraction of a system is not useful

 $C_{\rm T}(f_{\rm M}, f_{\rm opt}) = \lfloor 2.26 \rfloor = 2$ $C_{\rm A7}(f_{\rm M}, f_{\rm opt}) = \lfloor 3.19 \rfloor = 3$

Computational Performance Power and Energy Efficiency

Iso-power configurations



- Not always maximum frequency (*f_M*) implies best performances
- The power consumption of the optimal frequency is smaller
- Iso-power determines how many systems working at f_{opt} dissipate the same power than one system working at f_M

 $P_{T}(f_{M})/P_{T}(f_{opt}) = 0.587/0.259 = 2.26 \times P_{A7}(f_{M})/P_{A7}(f_{opt}) = 0.438/0.137 = 3.19 \times$

Given that the fraction of a system is not useful

 $C_{\rm T}(f_{\rm M}, f_{\rm opt}) = \lfloor 2.26 \rfloor = 2$ $C_{\rm A7}(f_{\rm M}, f_{\rm opt}) = \lfloor 3.19 \rfloor = 3$

Computational Performance Power and Energy Efficiency

Iso-power configurations



- Not always maximum frequency (*f_M*) implies best performances
- The power consumption of the optimal frequency is smaller
- Iso-power determines how many systems working at f_{opt} dissipate the same power than one system working at f_M

$$P_{\rm T}(f_{\rm M})/P_{\rm T}(f_{\rm opt}) = 0.587/0.259 = 2.26 \times P_{\rm A7}(f_{\rm M})/P_{\rm A7}(f_{\rm opt}) = 0.438/0.137 = 3.19 \times$$

• Given that the fraction of a system is not useful

$$\begin{split} C_{\mathrm{T}}(f_{\mathrm{M}},f_{\mathrm{opt}}) &= \lfloor 2.26 \rfloor = 2\\ C_{\mathrm{A7}}(f_{\mathrm{M}},f_{\mathrm{opt}}) &= \lfloor 3.19 \rfloor = 3 \end{split}$$

Computational Performance Power and Energy Efficiency

Iso-power configurations

- Assuming perfect scalability,
 - The execution time will be reduced

$$\bar{T}(f_{\text{opt}}, C_{\text{T}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/2 = 159.3/2 = 79.6s$$
$$\bar{T}(f_{\text{opt}}, C_{\text{A7}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/3 = 159.3/3 = 53.1s$$

• Where the energy consumption will be

 $\bar{E}_{T}(f_{opt}, C_{T}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 2) \cdot P_{T}(f_{M}) = 79.6 \cdot 0.587 = 46.7W$ $\bar{E}_{A7}(f_{opt}, C_{A7}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 3) \cdot P_{A7}(f_{M}) = 53.1 \cdot 0.438 = 23.2W$

• This represents a gain in energy efficiency

 $E_{\rm T}(f_{\rm M})/\bar{E}_{\rm T}(f_{\rm opt}, C_{\rm T}(f_{\rm M}, f_{\rm opt})) = 55.8/46.7 = 1.19 \times E_{\rm A7}(f_{\rm M})/\bar{E}_{\rm A7}(f_{\rm opt}, C_{\rm A7}(f_{\rm M}, f_{\rm opt})) = 41.3/23.2 = 1.77 \times$

Too optimistic assumption is compensated with rounding down



Computational Performance Power and Energy Efficiency

Iso-power configurations

- Assuming perfect scalability,
 - The execution time will be reduced

$$\bar{T}(f_{\text{opt}}, C_{\text{T}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/2 = 159.3/2 = 79.6s$$
$$\bar{T}(f_{\text{opt}}, C_{\text{A7}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/3 = 159.3/3 = 53.1s$$

Where the energy consumption will be

 $\bar{E}_{T}(f_{opt}, C_{T}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 2) \cdot P_{T}(f_{M}) = 79.6 \cdot 0.587 = 46.7W$ $\bar{E}_{A7}(f_{opt}, C_{A7}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 3) \cdot P_{A7}(f_{M}) = 53.1 \cdot 0.438 = 23.2W$

• This represents a gain in energy efficiency

$$E_{\rm T}(f_{\rm M})/\bar{E}_{\rm T}(f_{\rm opt}, C_{\rm T}(f_{\rm M}, f_{\rm opt})) = 55.8/46.7 = 1.19 \times E_{\rm A7}(f_{\rm M})/\bar{E}_{\rm A7}(f_{\rm opt}, C_{\rm A7}(f_{\rm M}, f_{\rm opt})) = 41.3/23.2 = 1.77 \times$$

• Too optimistic assumption is compensated with rounding down



Computational Performance Power and Energy Efficiency

Iso-power configurations

- Assuming perfect scalability,
 - The execution time will be reduced

$$\bar{T}(f_{\text{opt}}, C_{\text{T}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/2 = 159.3/2 = 79.6s$$
$$\bar{T}(f_{\text{opt}}, C_{\text{A7}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/3 = 159.3/3 = 53.1s$$

Where the energy consumption will be

 $\bar{E}_{T}(f_{opt}, C_{T}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 2) \cdot P_{T}(f_{M}) = 79.6 \cdot 0.587 = 46.7W$ $\bar{E}_{A7}(f_{opt}, C_{A7}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 3) \cdot P_{A7}(f_{M}) = 53.1 \cdot 0.438 = 23.2W$

This represents a gain in energy efficiency

$$E_{\rm T}(f_{\rm M})/\bar{E}_{\rm T}(f_{\rm opt}, C_{\rm T}(f_{\rm M}, f_{\rm opt})) = 55.8/46.7 = 1.19 \times E_{\rm A7}(f_{\rm M})/\bar{E}_{\rm A7}(f_{\rm opt}, C_{\rm A7}(f_{\rm M}, f_{\rm opt})) = 41.3/23.2 = 1.77 \times$$

Too optimistic assumption is compensated with rounding down



Computational Performance Power and Energy Efficiency

Iso-power configurations

- Assuming perfect scalability,
 - The execution time will be reduced

$$\bar{T}(f_{\text{opt}}, C_{\text{T}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/2 = 159.3/2 = 79.6s$$

$$\bar{T}(f_{\text{opt}}, C_{\text{A7}}(f_{\text{M}}, f_{\text{opt}})) = T(f_{\text{opt}})/3 = 159.3/3 = 53.1s$$

Where the energy consumption will be

 $\bar{E}_{T}(f_{opt}, C_{T}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 2) \cdot P_{T}(f_{M}) = 79.6 \cdot 0.587 = 46.7W$ $\bar{E}_{A7}(f_{opt}, C_{A7}(f_{M}, f_{opt})) = \bar{T}(f_{opt}, 3) \cdot P_{A7}(f_{M}) = 53.1 \cdot 0.438 = 23.2W$

This represents a gain in energy efficiency

$$E_{\rm T}(f_{\rm M})/\bar{E}_{\rm T}(f_{\rm opt}, C_{\rm T}(f_{\rm M}, f_{\rm opt})) = 55.8/46.7 = 1.19 \times E_{\rm A7}(f_{\rm M})/\bar{E}_{\rm A7}(f_{\rm opt}, C_{\rm A7}(f_{\rm M}, f_{\rm opt})) = 41.3/23.2 = 1.77 \times$$

Too optimistic assumption is compensated with rounding down



Computational Performance Power and Energy Efficiency

Iso-power configurations

11

Freq.	$P(f_{\rm M})/P(f_i)$		$C_i = C_i$	$C(f_{\mathbf{M}}, f_i) =$	$\overline{T}(f_i,$	$C_i) =$	Ē(fi	$, C_i) =$	Gair	n/loss=
(MHz)			LP(f _N	$(1)/P(f_i)$	$T(f_i)$	$C(f_i)$	$\overline{T}(f_i, C$	$F_i) \cdot P(f_M)$	$E(f_{\rm M})$	$/\bar{E}(f_i, C_i)$
f _i	Total	A7	Total	A7	Total	A7	Total	A7	Total	A7
50	2.73	3.71	2	3	145.9	97.2	85.6	42.6	0.65	0.96
100	2.74	3.94	2	3	129.7	86.5	76.1	37.8	0.73	1.09
150	2.70	4.01	2	4	115.0	57.5	67.5	25.1	0.82	1.63
200	2.60	3.91	2	3	101.8	67.8	59.7	29.7	0.93	1.38
250	2.45	3.61	2	3	90.0	60.0	52.8	26.2	1.05	1.57
300	2.26	3.19	2	3	79.6	53.1	46.7	23.2	1.19	1.77
350	2.04	2.72	2	2	70.7	70.7	41.5	30.9	1.34	1.33
400	1.80	2.25	1	2	126.6	63.3	74.3	27.7	0.75	1.48
450	1.57	1.84	1	1	114.6	114.6	67.2	50.1	0.82	0.82
500	1.36	1.50	1	1	105.5	105.5	61.9	46.2	0.90	0.89
550	1.16	1.22	1	1	99.3	99.3	58.2	43.4	0.95	0.94
600	1.00	1.00	1	1	96.0	96.0	55.8	41.3	1.00	1.00

Outline



- 1 Motivation and Introduction
- 2 Task-parallel version of ILUPACK
 - Preconditioner computation
 - Iterative solver
- 3 Setup
- Frequency-Voltage Scaling
 - Computational Performance
 - Power and Energy Efficiency
- 5 Fault tolerance and NTVC
 - Conclusions



- The key is to reduce voltage while keeping the frequency
 - Reduction of power \rightarrow additional core concurrency
 - Hardware less reliable \rightarrow additional mechanism
- We only consider corruptions on floating point arithmetic operations, degrading the convergence rate of the PCG
- The power are related to voltage and frequency

 $P \propto V^2 f$

and current technology sets V proporcionally to f, for safety

 $V \propto f$

therefore any change on V will also change f (not considered)



- The key is to reduce voltage while keeping the frequency
 - Reduction of power \rightarrow additional core concurrency
 - $\bullet \ \ \text{Hardware less reliable} \to \text{additional mechanism}$
- We only consider corruptions on floating point arithmetic operations, degrading the convergence rate of the PCG
- The power are related to voltage and frequency

 $P\propto V^2 f$

and current technology sets V proporcionally to f, for safety

 $V \propto f$

therefore any change on V will also change f (not considered)



- The key is to reduce voltage while keeping the frequency
 - Reduction of power \rightarrow additional core concurrency
 - $\bullet \ \ \text{Hardware less reliable} \to \text{additional mechanism}$
- We only consider corruptions on floating point arithmetic operations, degrading the convergence rate of the PCG
- The power are related to voltage and frequency

 $P \propto V^2 f$

and current technology sets V proporcionally to f, for safety

 $V \propto f$

therefore any change on V will also change f (not considered)

Given the relations between time, power and energy

$$T \propto f \;,\; P \propto V^2 f \;,\; E = T * P$$

If no errors occurs,

$$\hat{V} = V/\sigma$$
, $\hat{f} = f \Rightarrow \hat{T} = T$, $\hat{P} = P/\sigma^2$, $\hat{E} = E/\sigma^2$

• If some errors occurs, there are other configurations:

- Improve power/energy at the expense of time to solution
- Improve time to solution/energy using additional hardware



Given the relations between time, power and energy

$$T\propto f\,,\;P\propto V^2f\,,\;E=T*P$$

• If some errors occurs, there are other configurations:

Improve power/energy at the expense of time to solution

$$\hat{P} = P/s , \ \hat{f} = f \Rightarrow \hat{V} = V/\sqrt{s}$$

 $\hat{e} \to \hat{d} \Rightarrow \hat{T} = T * \hat{d} , \ \hat{E} = E * (\hat{d}/s)$

• Improve time to solution/energy using additional hardware



Ĵ

• Given the relations between time, power and energy

$$T \propto f$$
, $P \propto V^2 f$, $E = T * P$

• If some errors occurs, there are other configurations:

- Improve power/energy at the expense of time to solution
- Improve time to solution/energy using additional hardware

$$\begin{array}{l} \hat{P} = P/s \ , \ \hat{T} = T & \xrightarrow{iso-power} \tilde{P} = \hat{P} * s = P \ , \ \bar{T} = \hat{T}/s = T/s \\ \tilde{e} \to \tilde{c} \ , \ \tilde{e} \to \tilde{d} & \xrightarrow{degradation} \tilde{T} = (\bar{T} + \tilde{c}) * \tilde{d} = T * \tilde{d}/s + \tilde{c} * \tilde{d} \\ \xrightarrow{degradation} \tilde{E} = \tilde{T} * \tilde{P} = E * (\tilde{d}/s) + \tilde{c} * \tilde{d} * P \end{array}$$

Given the relations between time, power and energy

$$T \propto f \;,\; P \propto V^2 f \;,\; E = T * P$$

If no errors occurs,

$$\hat{V} = V/\sigma$$
 , $\hat{f} = f$ \Rightarrow $\hat{T} = T$, $\hat{P} = P/\sigma^2$, $\hat{E} = E/\sigma^2$

• If some errors occurs, there are other configurations:

Improve power/energy at the expense of time to solution

$$\hat{e}
ightarrow \hat{d} \;, \hat{P} = P/s \;, \; \hat{f} = f \; \Rightarrow \; \hat{T} = T * \hat{d} \;, \; \hat{V} = V/\sqrt{s} \;, \; \hat{E} = E * (\hat{d}/s)$$

• Improve time to solution/energy using additional hardware

$$\begin{split} \tilde{P} &= P \;, \; \tilde{e} \to \tilde{c}, \tilde{d} \; \Rightarrow \; \tilde{T} = (\bar{T} + \tilde{c}) * \tilde{d} = T * \tilde{d}/s + \tilde{c} * \tilde{d} \\ &\Rightarrow \tilde{E} = \tilde{T} * \tilde{P} = E * (\tilde{d}/s) + \tilde{c} * \tilde{d} * P \end{split}$$



Given the relations between time, power and energy

$$T \propto f \;,\; P \propto V^2 f \;,\; E = T * P$$

If no errors occurs,

$$\hat{V} = V/\sigma$$
, $\hat{f} = f$ \Rightarrow $\hat{T} = T$, $\hat{P} = P/\sigma^2$, $\hat{E} = E/\sigma^2$

• If some errors occurs, there are other configurations:

Improve power/energy at the expense of time to solution

$$\hat{e}
ightarrow \hat{d} \;, \hat{P} = P/s \;, \; \hat{f} = f \; \Rightarrow \; \hat{T} = T * \hat{d} \;, \; \hat{V} = V/\sqrt{s} \;, \; \hat{E} = E * (\hat{d}/s)$$

• Improve time to solution/energy using additional hardware

$$\begin{split} \tilde{P} &= P \;,\; \tilde{e} \to \tilde{c}, \tilde{d} \;\Rightarrow\; \tilde{T} = (\bar{T} + \tilde{c}) * \tilde{d} = T * \tilde{d}/s + \tilde{c} * \tilde{d} \\ &\Rightarrow \tilde{E} = \tilde{T} * \tilde{P} = E * (\tilde{d}/s) + \tilde{c} * \tilde{d} * P_c \end{split}$$



Conclusions



• The results determine an optimal frequency for energy efficiency $E_{\tau}^{opt} = E_T(350)$, $E_{AT}^{opt} = E_{AT}(250) \Rightarrow f_{opt} = 300$

• The iso-power configurations improves the performances

- Increase the number of units at f_{opt} to consume as $P(f_M)$
- Reduce the execution time and improve the energy efficiency
- We have analyzed how NTVC can be exploited
 - Without errors, the reductions are directly moved to energy

$$\hat{V} = V/\sigma \ \Rightarrow \ \hat{E} = E/\sigma^2$$

• With errors. the energy are related with the degradation

Conclusions



• The results determine an optimal frequency for energy efficiency

$$E_T^{opt} = E_T(350) , \ E_{A7}^{opt} = E_{A7}(250) \Rightarrow f_{opt} = 300$$

- The iso-power configurations improves the performances
 - Increase the number of units at fopt to consume as P(f_M)
 - Reduce the execution time and improve the energy efficiency

$$E_{T}(f_{M})/\bar{E}_{T}(f_{opt}, C_{T}(f_{M}, f_{opt})) = 1.19 \times E_{A7}(f_{M})/\bar{E}_{A7}(f_{opt}, C_{A7}(f_{M}, f_{opt})) = 1.77 \times$$

- We have analyzed how NTVC can be exploited
 - Without errors, the reductions are directly moved to energy

$$\hat{V} = V/\sigma \ \Rightarrow \ \hat{E} = E/\sigma^2$$

• With errors. the energy are related with the degradation

Conclusions



- The results determine an optimal frequency for energy efficiency $E_T^{opt} = E_T(350)$, $E_{A7}^{opt} = E_{A7}(250) \Rightarrow f_{opt} = 300$
- The iso-power configurations improves the performances
 - Increase the number of units at fopt to consume as P(f_M)
 - Reduce the execution time and improve the energy efficiency $E_T(f_M)/\bar{E}_T(f_{opt},...) = 1.19 \times$, $E_{A7}(f_M)/\bar{E}_{A7}(f_{opt},...) = 1.77 \times$
- We have analyzed how NTVC can be exploited
 - Without errors, the reductions are directly moved to energy

$$\hat{V} = V/\sigma \ \Rightarrow \ \hat{E} = E/\sigma^2$$

• With errors. the energy are related with the degradation

$$\hat{P} = P/s \Rightarrow \hat{E} = E * (\hat{d}/s)$$

 $\tilde{e} \to \tilde{c}, \tilde{d} \Rightarrow \tilde{E} = E * \tilde{d}/s + \tilde{c} * \tilde{d} * P_c$

Thanks for your attention ! Questions ?

Performance and Fault Tolerance of ILUPACK on ARM 27 J.Aliaga, S.Catalán, C.Chalios, D.Nikolopoulos, E.Quintana-Ortí