

micro and nanoelectronics
microsystems
ambient intelligence
image chain
biology and health



2008

Dynamic Voltage and Frequency Scaling Architecture for Units Integration within a GALS NoC

Edith Beigné

F. Clermidy, S. Miermont, P. Vivet



leti

MINATEC®

INSTITUT
CARNOT
CEA LETI

cea

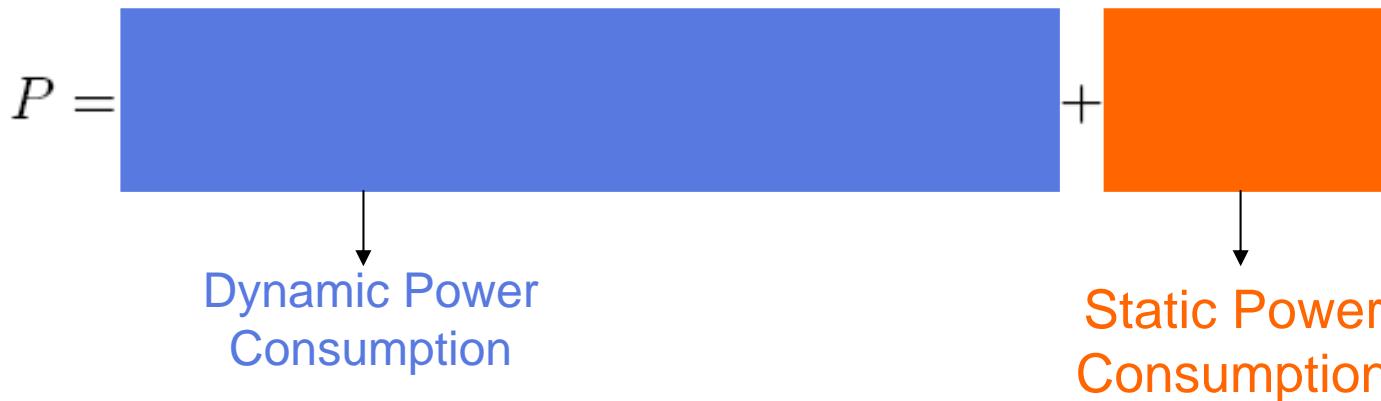
Outline

- Dynamic and Static power consumption issues
- NoC architecture for DVFS support
- NoC Unit architecture
- NoC Unit design
- DVFS execution at system level
- Power gain & physical implementation

Outline

- Dynamic and Static power consumption issues
- NoC architecture for DVFS support
- NoC Unit architecture
- NoC Unit design
- DVFS execution at system level
- Power gain & physical implementation

Dynamic and Static power consumption issues



- Dynamic power consumption reduction :
 - Reduce : switching activity, capacitances, supply voltage, frequency
- Static power consumption reduction :
 - Reduce : supply voltage, dominant leakage currents : I_{STH} , I_{GIDL} , I_{GATE}
 - Multi V_{TH} design
- Low power techniques must exists at all design levels from architecture to physical implementation

- We propose a fully integrated solution to control locally dynamic and static power at Unit level within a GALS NoC

© CEA 2008. Tous droits réservés.
 Toute reproduction totale ou partielle sur quelque support que ce soit ou utilisation du contenu de ce document est interdite sans l'autorisation écrite préalable du CEA
 All rights reserved. Any reproduction in whole or in part on any medium or use of the information contained herein is prohibited without the prior written consent of CEA

Outline

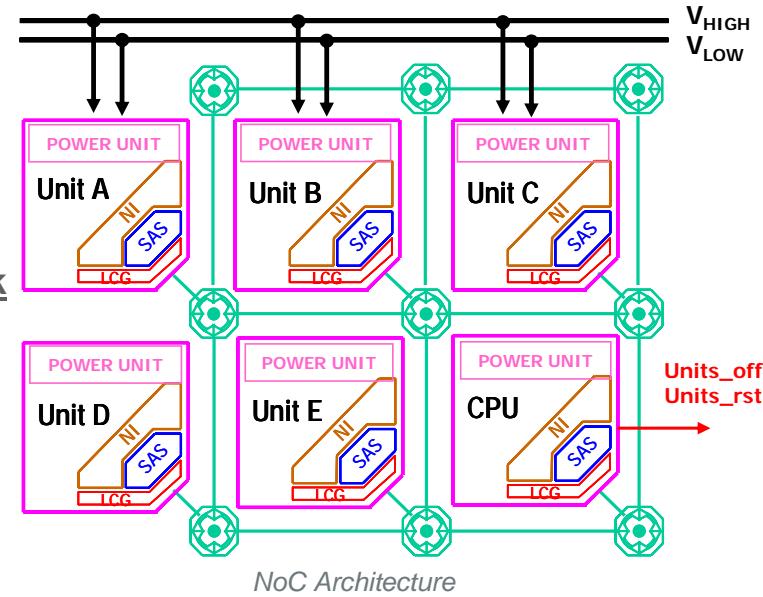
- Dynamic and Static power consumption issues
 - NoC architecture for DVFS support
- NoC Unit architecture
- NoC Unit design
- DVFS execution at system level
- Power gain & physical implementation

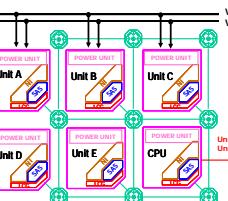
NoC architecture for DVFS support (1)

NoC architecture

- A fully asynchronous Network-on-Chip
- IP units are synchronous islands using programmable Local Clock Generator
- Within the IP unit :
 - Synchronization is done thanks to Pausable Clock
 - A Power Unit manages internal Vcore generated using external V_{high} and V_{low}
 - A Network Interface is in charge of
 - NoC communications
 - Local Power Management
- Main CPU in charge of global power management

➤ Local fine grain power management can be executed during IP computation and communication independently from the others





NoC architecture for DVFS (2)

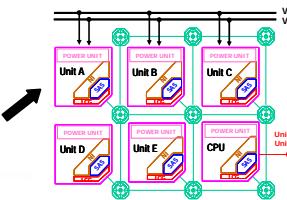
Main Principles

- Each synchronous IP is an independant power and frequency domain
- A local fine grain **Dynamic Voltage Scaling** :
 - Implementation of a local hardware controller to control transitions between V_{high} and V_{low}
 - Ensures smooth DVS transitions for IP safe computation
- A local fine grain **Dynamic Frequency Scaling** :
 - Automatic frequency scaling
 - Use of clock generation re-programming to find the optimal V/F point of operation
- Thanks to pausable clock technique, IP unit **continues** its operation during DVFS phases
 - **GALS architecture and local clock generation is a natural enabler for easy local DVFS**

Outline

- Dynamic and Static power consumption issues
- NoC architecture for DVFS support
- NoC Unit architecture
- NoC Unit design
- DVFS execution at system level
- Power gain & physical implementation

NoC Unit architecture

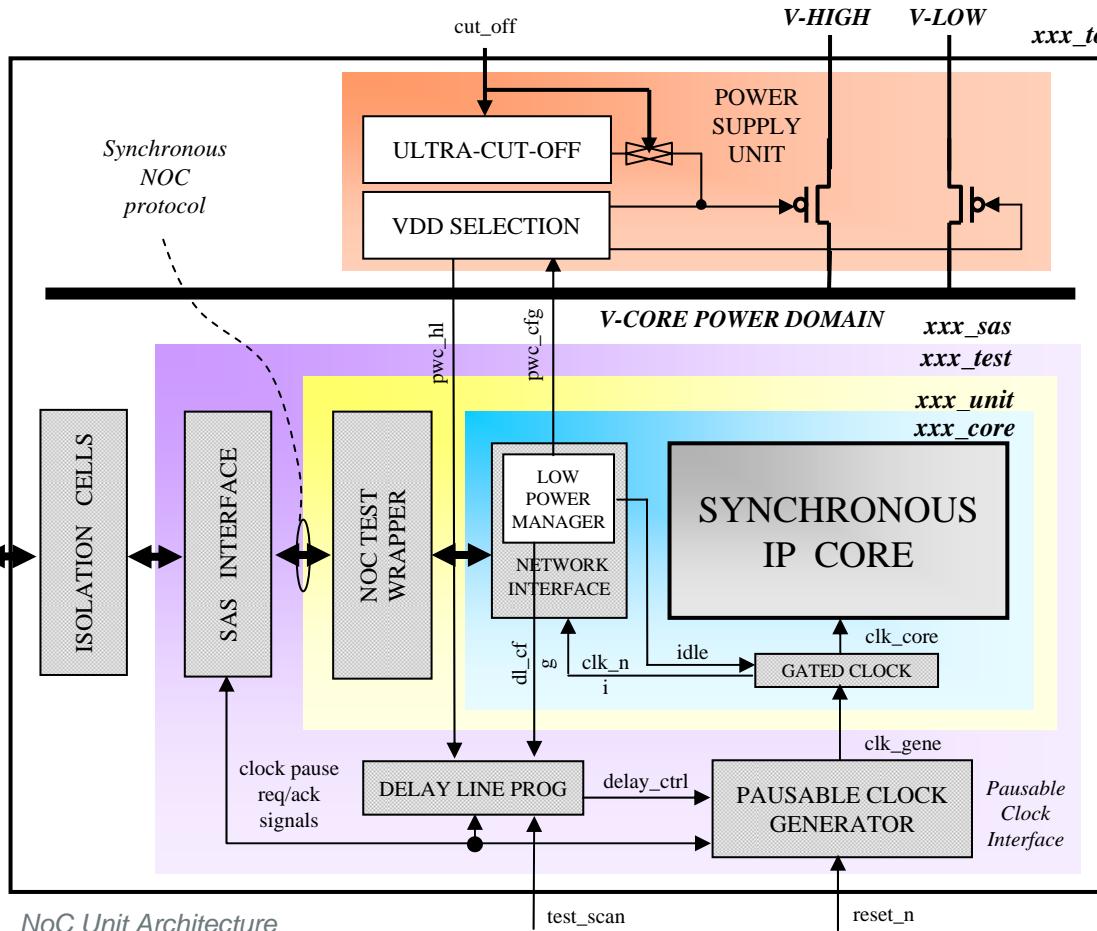


- Each IP core encapsulated with

- Network Interface,
- Test Wrapper,
- Pausable Clock,
- Power Supply Unit

- IP units have 5 supply modes

- **Init** : reset at *Vhigh* (1.2V)
- **High** : *Vhigh* supply
- **Low** : *Vlow* supply (0.8V)
- **Hopping** : switch *Vhigh* / *Vlow* for DVFS
- **Idle** : retention state at *Vlow* (no clock)
- **Off** : stand-by mode

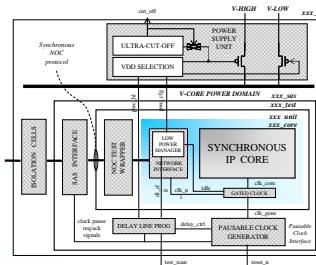


Outline

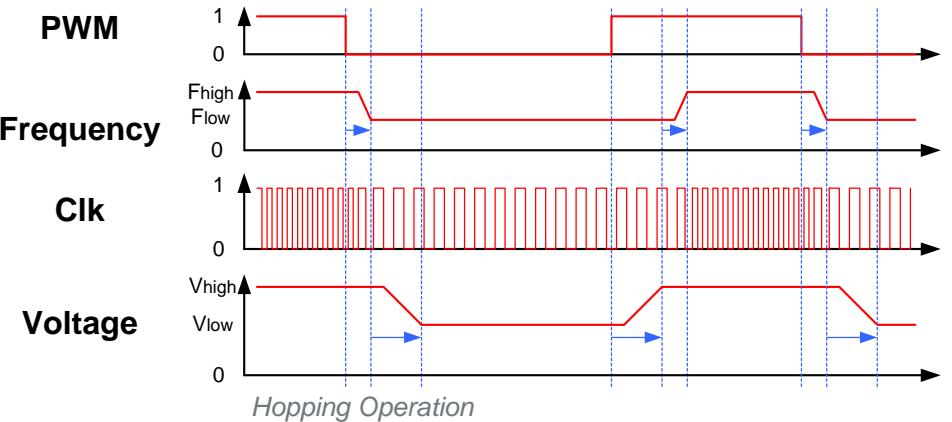
- Dynamic and Static power consumption issues
- NoC architecture for DVFS support
- NoC Unit architecture
- NoC Unit design
- DVFS execution at system level
- Power gain & physical implementation

NoC Unit design (1)

Local Power Manager



- Local Power Manager handles unit power modes
- A set of programmable registers, through the NoC
- Configuration of :
 - Programmable delay line
 - Power Supply Unit
- Pulse Width modulator used to control the Hopping mode

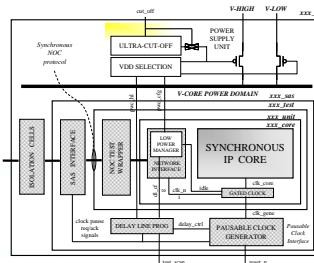


© CEA 2008. Tous droits réservés.

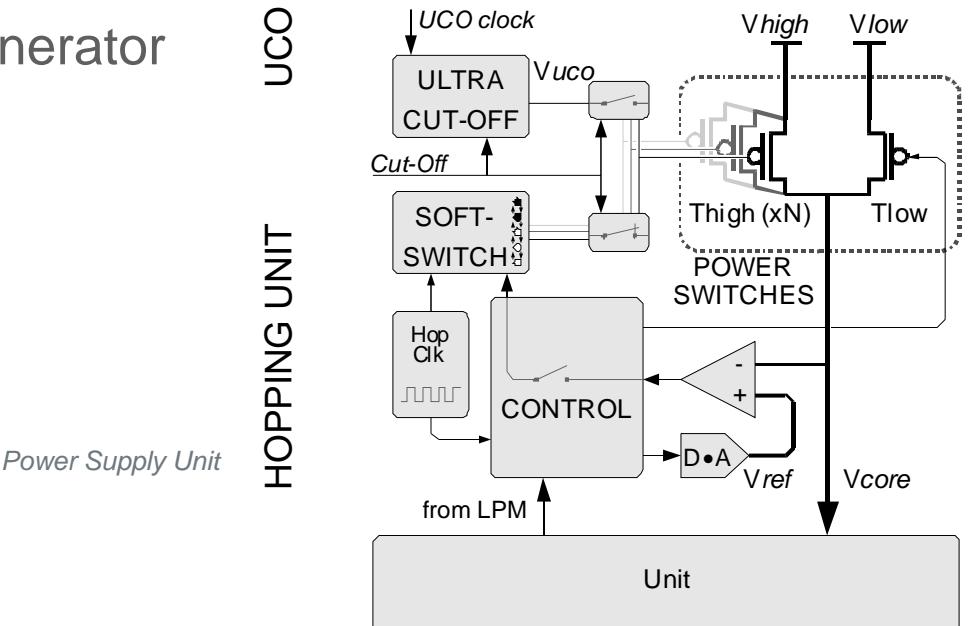
Toute reproduction totale ou partielle sur quelque support que ce soit ou utilisation du contenu de ce document est interdite sans l'autorisation écrite préalable du CEA
All rights reserved. Any reproduction in whole or in part on any medium or use of the information contained herein is prohibited without the prior written consent of CEA

NoC Unit design (2)

Power Supply Unit



- Power Supply Unit manages V_{core}
- Two power switches T_{high} and T_{low} LVT transistors
- A Hopping Unit
- An Ultra Cut-Off Generator



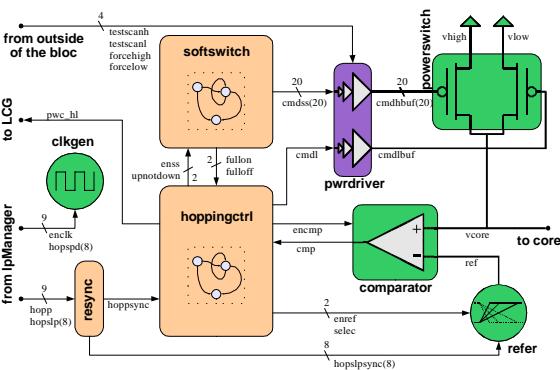
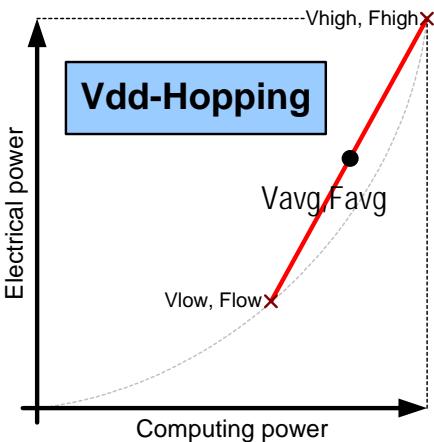
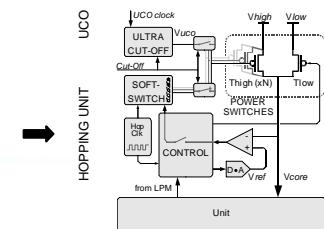
➤ Local Power Supply Unit offers a safe control of internal power supply depending on pre-defined power modes

© CEA 2008. Tous droits réservés.
 Toute reproduction totale ou partielle sur quelque support que ce soit ou utilisation du contenu de ce document est interdite sans l'autorisation écrite préalable du CEA
 All rights reserved. Any reproduction in whole or in part on any medium or use of the information contained herein is prohibited without the prior written consent of CEA

NoC Unit design (3)

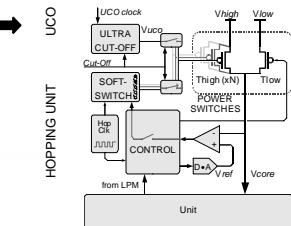
Hopping Unit

- Energy per operation scales with V^2
 - Decrease Voltage (and Frequency) to be energy efficient
- « Triple state » power supply
 - Proposed by Tokyo Univ. (2000)
 - ◆ But was not integrated on chip
 - Use of two PMOS power switches
 - ◆ Vhigh (1.2 V), Vlow (0.7 V), or off (0 V)
- Switch between Vhigh and Vlow :
 - Transitions take less than 100 ns
 - Mean speed / mean power of the IP is programmed by a PWM
- Compatible with synchronous and asynchronous IPs
 - For GALS system : coordination done with local clock generator
- Can easily be integrated in any CMOS circuit
 - No inductor contrary to traditional DC/DC converters
 - No capacitor contrary to charge pump implementation

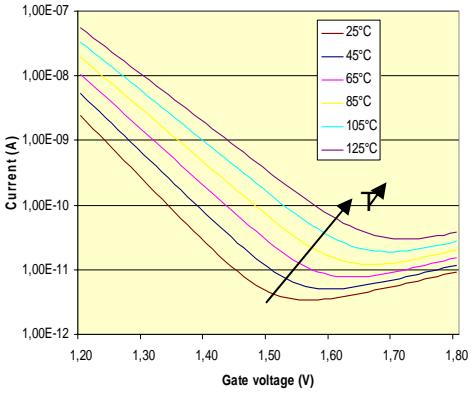


NoC Unit design (4)

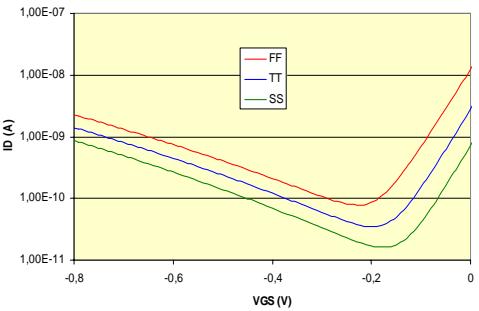
Ultra Cut-Off Generator



- When reverse polarizing the gate, the leakage current goes through a minimum:
 - The subthreshold current I_{STH} diminishes,
 - The I_{GIDL} current increases.
- The optimal polarization point varies with the temperature, the supply voltage and the process corners
- The proposed UCO generator automatically polarizes the gate of the Power switch to its point of minimum leakage
 - two reference transistors respectively representing leakage currents
 - a current sense amplifier to compare those currents
 - a charge pump to generate the polarization voltage V_P
- Compensates for temperature variation, alleviates corners variations.
- The gate oxide reliability is considered by introducing a passive stress reduction mechanism



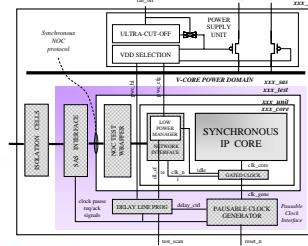
Variation of the total leakage current with temperature.



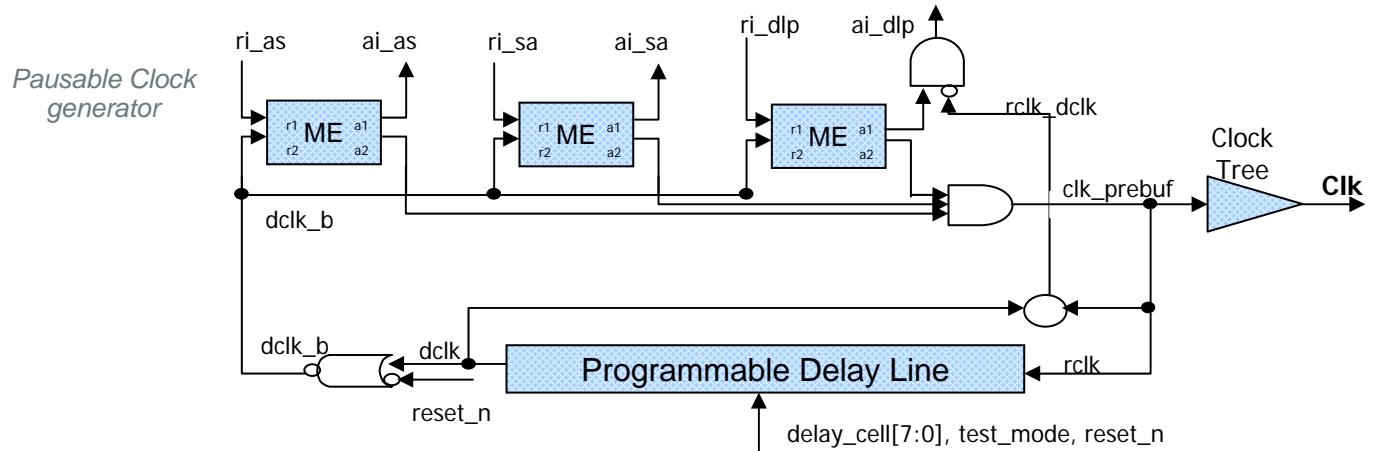
Variation of the total leakage current with corner.

NoC Unit design (5)

Pausable Clock Interface (1)



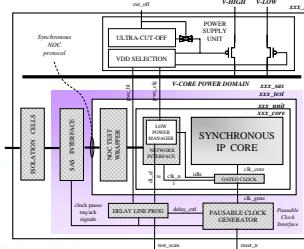
- Pause temporary the clock when a transfer (NoC) or a supply switch is required (LPM)
- Based on :
 - Two GALS ports : Synchronous-to Asynchronous and Asynchronous-to-Synchronous
 - A programmable delay line
 - A pausable clock generator
- Pausable Clock Generator arbitrates pause requests



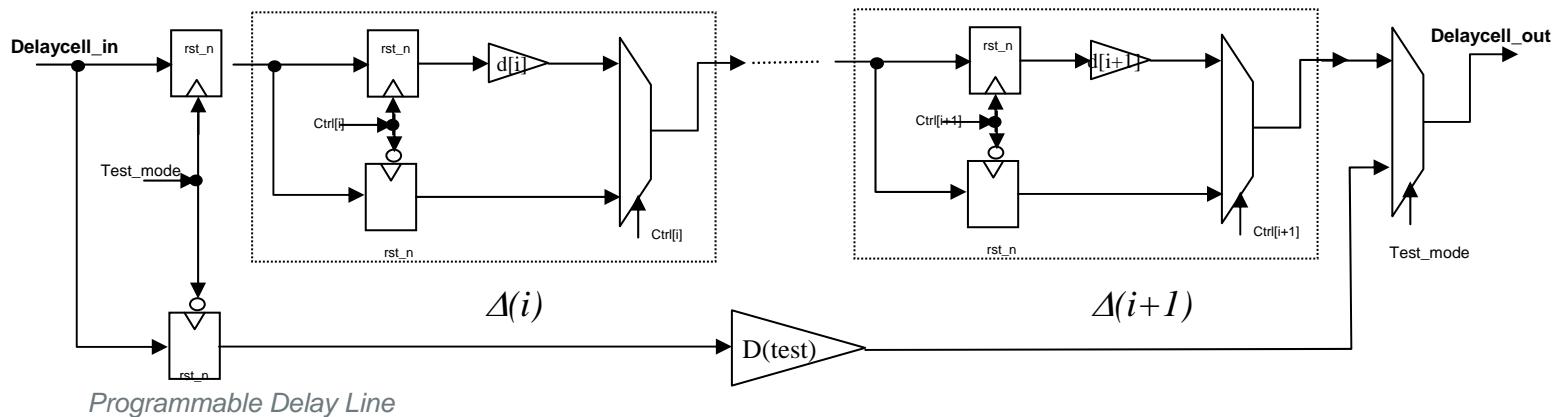
© CEA 2008. Tous droits réservés.
Toute reproduction totale ou partielle sur quelque support que ce soit ou utilisation du contenu de ce document est interdite sans l'autorisation écrite préalable du CEA
All rights reserved. Any reproduction in whole or in part on any medium or use of the information contained herein is prohibited without the prior written consent of CEA

NoC Unit design (3)

Pausable Clock Interface (2)



- Programmable delay line :
 - Precise, small and low power
 - Using Standard cells
 - On the same unit power domain



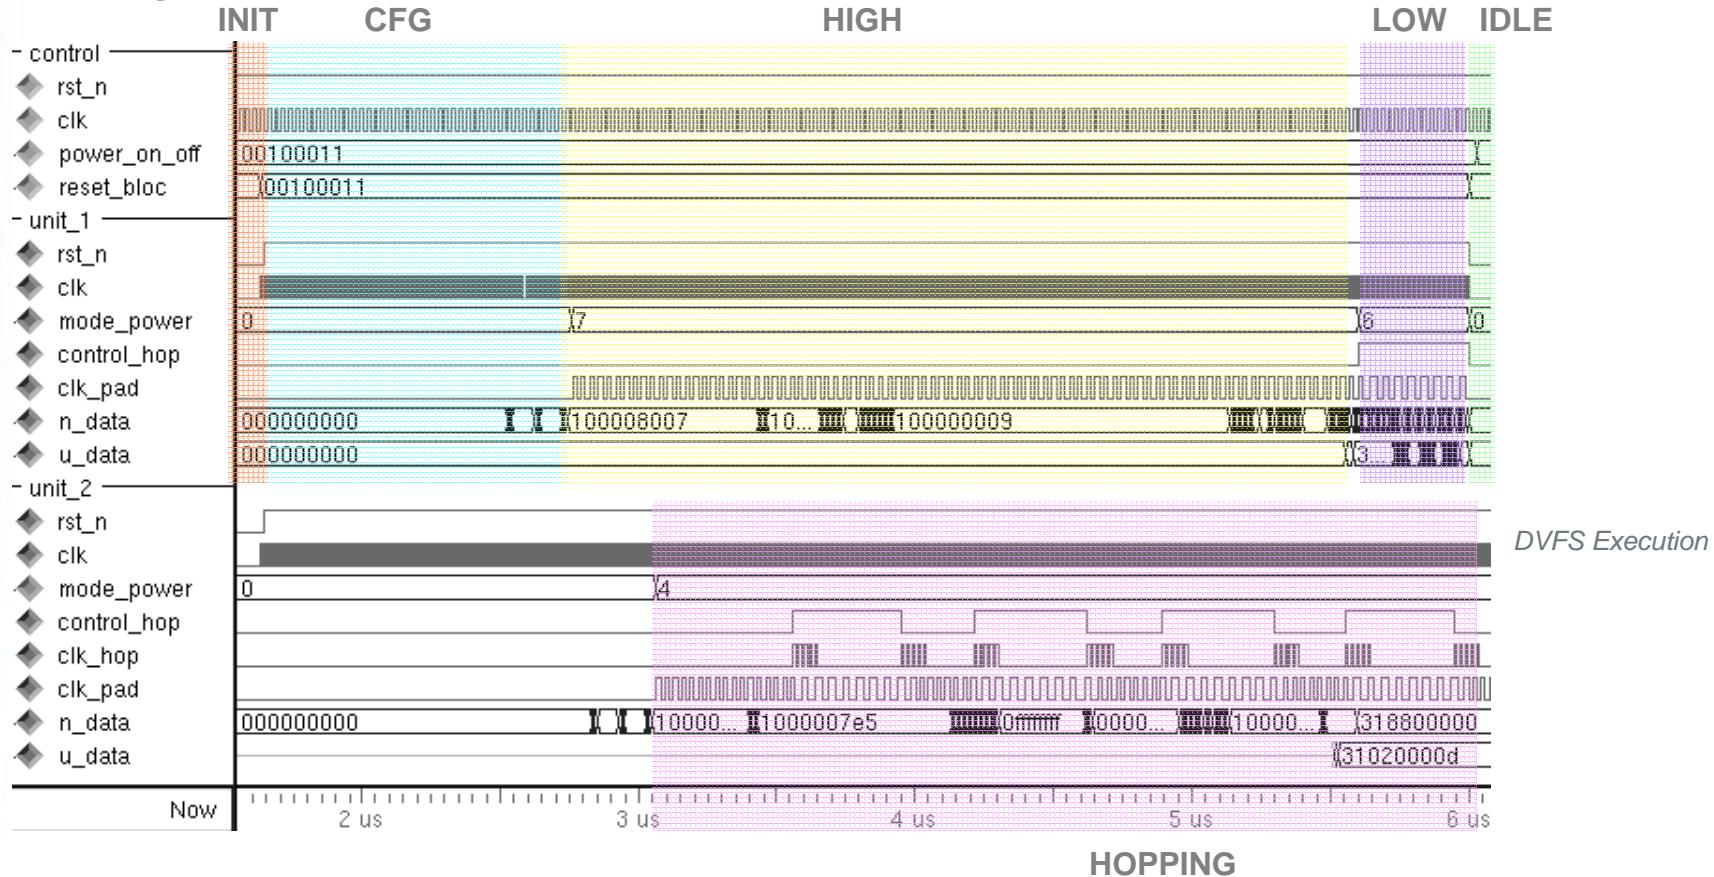
➤ Pausable Clock Interface allows an efficient synchronization and a safe dynamic voltage and frequency scaling with minimal latency cost

Outline

- Dynamic and Static power consumption issues
- NoC architecture for DVFS support
- NoC Unit architecture
- NoC Unit design
- DVFS execution at system level
- Power gain & physical implementation

DVFS execution at system level

- DVFS strategy programmed through NoC, Network Interface and global CPU



➤ Global power management strategy is aiming at reducing the global energy while respecting global latency constraints

© CEA 2008. Tous droits réservés.
Toute reproduction totale ou partielle de ce document est interdite sans l'autorisation écrite préalable du CEA
All rights reserved. Any reproduction in whole or in part on any medium or use of the information contained herein is prohibited without the prior written consent of CEA

Outline

- Dynamic and Static power consumption issues
- NoC architecture for DVFS support
- NoC Unit architecture
- NoC Unit design
- DVFS execution at system level
- Power gain & physical implementation

Power gain & physical implementation (1)

Power gain

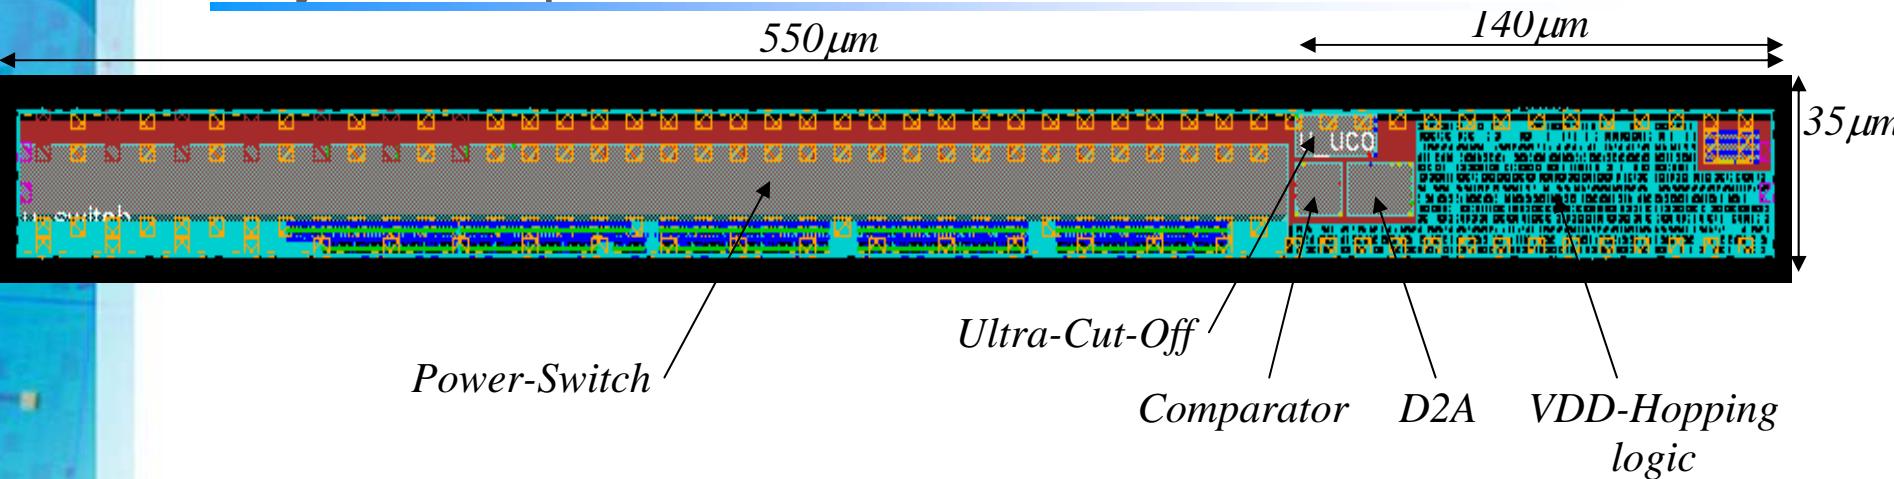
- Programmable delay line matches with unit logic on the same power domain :
 - Compensates any mismatch thanks to re-programmation
 - Power reduction :
 - $V_{high}=1.2V$ and $V_{low}=0.8V$
 - 35 % dynamic power reduction between *High* and *Low* modes
 - *Hopping* mode is used to save power without any latency cost
 - Leakage power thanks to UCO is reduced by 2 decade
 - Power Supply Unit efficiency :
 - Hopping Unit :
 - ◆ Only resistive losses in the power transistors
 - ◆ About 1 mW dynamic power
 - => more than 95 % power efficiency
 - 90 % total efficiency (external DC-DC taken into account)
- An adaptive and reliable Power Supply Unit giving high power reduction factor and high power efficiency

© CEA 2008. Tous droits réservés.

Toute reproduction totale ou partielle sur quelque support que ce soit ou utilisation du contenu de ce document est interdite sans l'autorisation écrite préalable du CEA
All rights reserved. Any reproduction in whole or in part on any medium or use of the information contained herein is prohibited without the prior written consent of CEA

Power gain & physical implementation (1)

Physical Implementation



■ Power Switch :

- One single Power-Switch for the complete power domain
- Sized to get a speed loss < 5%
- Area : about < 5% of the power domain

■ Hopping Unit

- Area : 140μm*35μm
- Hopping Transition : < 100 ns

Conclusion

- A fully integrated DVFS architecture within a GALS NoC
- We are able to handle leakage problems due to technology scaling
 - insertion of power switches
- Dynamic power reduction is possible through voltage scaling:
 - Hopping
 - Management of multi power domains in a complex SoC
- Our knowledge of GALS systems lead us to automatic frequency scaling
 - Pausable clock interfaces
 - Asynchronous implementation

micro and nanoelectronics

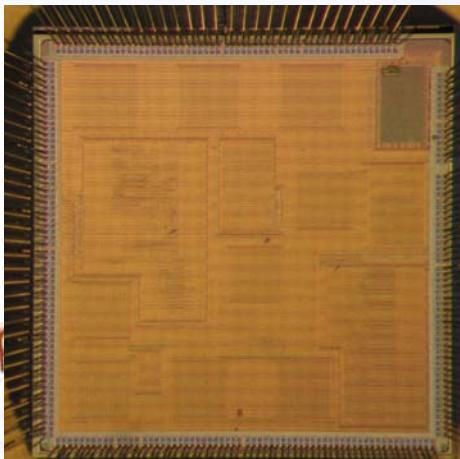
microsystems

ambient intelligence
biology and health

image chain



ALPIN :an “Asynchronous Low Power Innovative NoC”



A close-up photograph of a microchip (silicon wafer) showing a complex grid of gold-colored interconnects and circuitry. The chip is mounted on a black substrate with gold-plated pins.

Loyalty
Entrepreneurship
Team work
Loyalty Innovation
Entrepreneurship
Team work
Innovation

Team work
Innovation

Leti

MINATEC

INSTITUT CARNOT
CEA LETI



Four small inset images showing people in laboratory or office environments:

- Top-left: A person in a white lab coat and glasses looking at a computer screen.
- Top-right: A person in a white lab coat and glasses working with a piece of equipment.
- Bottom-left: A person in a white lab coat and glasses looking at a computer screen.
- Bottom-right: A person in a white lab coat and glasses looking directly at the camera.