Ultra-Low Power Design Approaches for IoT



Massimo Alioto



National University of Singapore (NUS) ECE Department Green IC group



Outline

- IoT: the context
- Ultra-low voltage operation
- Design Issues and Solutions at Ultra-Low Voltages
 - performance
 - leakage
 - variations and resiliency

Conclusions

Green IC

NUS

Internet of Things: The Context





Nodes for the Internet of Things: Peculiarities
Specific features of nodes for IoT



Prof. Massimo Alioto

~1-100 mm³ (battery, energy scavenger) small size





Recenic University Greenic







Green IC

NUS National University

communication

<u>computation vs</u> <u>communication tradeoff</u> - data representation (compressive sensing, compression)

limit TX to critical events
 or significant changes
 (critical event monitoring)





Power vs Energy

Duty cycled systems with limited power

• active only periodically (or on demand) for a short time



 partition into always-on block (timers, retentive memory) and duty cycled blocks (all others, active 0.1-1% of the time)

$$P_{avg} = P_{always-on} + E_{active} / T_{wkup}$$

[RJA12] M. Alioto, et Al., "Active RFID: A Perpetual Wireless Communications Platform for Sensors," ESSCIRC 2012

Prof. Massimo Alioto

• minimum power is the goal for always-on block

minimize power, essentially

leakage (dynamic power very small

– little active and slow)



- minimum energy per operation is the goal for dutycycled blocks
 - minimize energy per operation
 (dynamic + leakage energy)
 - 1/X duty cycling increases energy budget by X

Green IC



in both cases, ultra-low voltage operation is absolutely needed

Ultra-Low Voltage Operation



Operation at Ultra-Low Voltages (ULV)
Voltage scaling is a powerful knob to improve energy efficiency

- ◆ quadratic benefit, if dynamic energy CV² dominates
- performance degradation
- How aggressively should we scale V_{DD} ?

- energy-performance tradeoff
- NT: relatively good speed nearly min. energy,
- **ST**: low speed, min. power



• Example: first mm³ system

[CGH11] G. Chen, et Al., "A 1 Cubic Millimeter Energy-Autonomous Wireless Intraocular Pressure Monitor," ISSCC 2011



transistors in weak (ST) or moderate (NT) inversion



Performance @ ULV: Transistor On-Current



♦ *I_{on}* defines performance



◆ *I*_{ON} EKV model (transregional) [EV06]

Inversion Coefficient vs gate overdrive

$$I_{0} = 2 \cdot n \cdot \mu \cdot C_{0X} \frac{W}{L} v_{t}^{2}$$

$$IC = \frac{I_{on}}{I_{0}} = [\ln(e^{v} + 1)]^{2}$$

$$V = \frac{v_{DD} - v_{TH}}{2 \cdot n \cdot v_{t}}$$

$$IC < 0.1 \quad 0.1 \le IC \le 10 \quad IC > 10$$

$$Weak$$
inversion
$$moderate$$
inversion
$$Strong$$
inversion

[EV06] C. Enz, E. Vittoz, Charge-Based MOS Transistor Modeling (...), Wiley, 2006 Prof. Massimo Alioto 17

17

region	V _{DD}	on current
weak inversion	sub threshold	$I_{on} \approx I_0 \cdot e^{\frac{V_{DD} - V_{TH}}{n \cdot v_t}}$
moderate inversion	near threshold	$I_{on} \approx 0.54 \cdot I_0 \cdot (\nu + 0.88)^2$
strong inversion	above threshold	$I_{on} \approx I_0 \cdot [v]^{\alpha} (\alpha \sim 1)$

• I_{ON} vs V_{DD} : steep (superlinear) in **NT** and **ST**



👺 NUS

Green IC

16

◆ Performance: fan-out-4 delay (FO4)

NUS National University Green IC

• speed of technology and (μ) architecture



Prof. Massimo Alioto

FO4

• Leveraging high sensitivity of I_{on} to V_{DD} @ NT

- V_{DD} powerful knob to dynamically improve performance
 - performance improvement due to $\Delta V=100$ -mV boosting

V_{DD}	FO4 improvement	
200 mV	7.5X	ST: 8X improvement/100 mV
400 mV	3.6X	
600 mV	1.6X	N1: 2-4X improvement/100 mV
800 mV	1.2X	
1 V	1.1X	

- **ST**: more effective, pretty useless (low performance anyway)
- **NT**: effective and low energy cost (ΔV is small fraction of V_{DD})

- Selective voltage boosting largely improves area and energy efficiency at NT and ST
 - sizing has a linear impact on transistor strength
 - boosting: superlinear \Rightarrow smaller transistor at iso-strength
 - example at NT (V_{DD} =500 mV)





Design Issues and Solutions at Ultra-Low Voltages: Performance



Performance Degradation

Performance at NT/ST worse than nominal voltage

- can be acceptable in practical situations
 - tasks of IoT nodes are often relatively simple
 - example of typical throughput: few hundreds of MOPS (e.g., video processing) down to kOPS (e.g., temper. monitoring)
- Wire delay at NT/ST
 - gate delay much less critical
 - than nominal V_{DD}
 - design as in the "good old days"





• If higher performance is sometimes needed...

wide voltage scaling

(e.g., from 400 mV to 1.2 V)

- large performance improvement (10X)
- energy/op increases quadratically (9X)
- tolerable if occasional
- optimized for NT

Green IC

 \Rightarrow energy/performance tradeoff

at 1.2 V is degraded by ~25%

[J12] S. Jain, et Al., IEEE ISSCC Dig. Tech. Papers, pp. 66–67, Feb. 2012



Design Corner Evaluations

$E_{dyn} \propto C \cdot V_{DD}^2$







- trading off area for performance: parallelism
 - intra-chip communication limits perf./energy gains
 - specificity of task can be leveraged for better balance of computation/communication cost
 - across-level design is required to manage these tradeoffs





NT



- trading off area for performance: specialized HW
 - example: FFT, Java processor, AES, MPEG decoder...
 - specialized HW has better performance/energy tradeoff than general-purpose

Green IC



• larger benefit for recurrent and specific tasks

Design Issues and Solutions at Ultra-Low Voltages: Leakage



Energy vs V_{DD}

• If dynamic energy per operation dominates:

$$E_{dyn} = \alpha_{SW} \cdot C \cdot V_{DD}^2$$

- reduce V_{DD} as much as possible
 - energy reduction limited by $V_{DD,min}$ (defined by robustness issues, very different for logic and memory)





Minimum Energy Point

- Combining dynamic and leakage energy
 - minimum energy point (MEP)
 - relatively **flat** (V_{DD} mainly set by performance target)
 - can lie in either **NT** or **ST**
 - time varying: depends on temperature, data set...





- Leakage energy takes up increasingly larger fraction of total energy at lower V_{DD}
 - at low V_{DD} , leakage energy increases exponentially, dynamic energy decreases quadratically
 - ex.: processor with L1 cache ($V_{DDcache,min}=0.55$ V)

[J12] S. Jain, et Al., IEEE ISSCC Dig. Tech. Papers, pp. 66–67, Feb. 2012

Green IC



🛛 Logic Leakage Power

Memory Leakage Power



Traditional Techniques for Low Leakage

- Leakage is truly critical (process not enough)
 - large, limits energy reduction
- Several traditional circuit techniques do not work...
 - transistor stacking is ineffective



• **power gating** is much less effective $(I_{on}/I_{off}$ degradation)

typical leakage reduction: 10-100X

Green IC

🔡 NUS

- NT: small leakage reduction, ST: no leakage reduction at all
- solution: boost gate voltage of sleep transistor (increases I_{on}/I_{off})
 selective voltage boosting

• multi- V_{TH} actually degrades energy efficiency

• delay sensitive to $V_{TH} \Rightarrow$ critical path changes at scaled V_{DD}

Counteracting Leakage at ULV

Leakage reduction at ULV: alternative approaches

- fine-grain power gating
 - disable unused blocks at runtime
 - small \Rightarrow quickly and frequently
 - overhead: multiple sleep transistors, isolation, control
 - lower control overhead via ckt/architectural support for SW
- fine-grain voltage domains

Green IC

- E_{lkg} reduced at lower V_{DD} (e.g., 2X/100 mV)
- selectively reduce V_{DD} wherever possible (slower)
- similar considerations as power gating

 $V_{\underline{DD},H}V_{\underline{DD},L}V_{\underline{DD},L}$

microarchitecture-circuit co-design: pipelining

 $E_{lkg} = V_{DD} \cdot I_{off} \cdot FO4 \cdot LD_{eff} \cdot CPO$

 use deep pipelines + refined circuit techniques/methodologies to deal with clocking overhead

• example: 17FO4/stage in FFT engine

(30MHz @ 0.27V, 4X less energy than state of the art) [AJC11] M. Seok, et Al. "A 0.27V, 30MHz, 17.7nJ/transform 1024-pt complex FFT core with super-pipelining," ISSCC 2011

body biasing

• makes sense in FDSOI (V_{TH} sensitivity to V_{BB} large enough)

Design Issues and Solutions at Ultra-Low Voltages: Variability and Resiliency

Variations: Why do They Matter?

- Resiliency degraded at ULV
 - process/voltage/temperature
 - larger than at nominal V_{DD}
 - ageing, soft errors...

[A12] M. Alioto, "Ultra-Low Power VLSI Circuit Design Demystified and Explained: A Tutorial," IEEE TCAS-I, Jan. 2012.

- design margining
 - cycle margin
 (20-30% @ full V_{DD})
 - degrades performance
 AND energy efficiency
 - large energy penalty

Green IC 🖬 🗐

Prof. Massimo Alioto

subthreshold

400

near threshold

600

σ\μ of *FO*4 normalized *to the case* V_{DD}=1 V

5

0

200

above threshold

800

1000

Process Variations

• Variability of delay (mainly due to I_{on})

Random Dopant Fluctuation, Line Edge Roughness...

- random variations are dominant (area-mismatch tradeoff)
- two negative effects arise at NT [GIS11]

^WLarger variability σ/μ

• due to larger sensitivity of I_{on} to V_{TH}

PDF heavily non-Gaussian

- in subthreshold tends to lognormal
- right skewed \Rightarrow average > nominal
- larger no. of σ s for given yield

[GIS11] G. Gammie, et Al., "A 28nm 0.6V low-power DSP for mobile applications," *ISSCC* 2011

Variability/leakage tradeoff unavoidable

path delay variability from FO4 variability (min. sized)

$$n_{\sigma} \frac{\sigma_{pathdelay}}{\mu_{pathdelay}} = n_{\sigma} \sqrt{\left(\frac{\sigma_{FO4,D2D}}{\mu_{FO4,D2D}}\right)^{2} + \frac{\left(\frac{\sigma_{FO4,random}}{\mu_{FO4,random}}\right)^{2}}{N_{gate} \cdot strength \cdot N_{stacked}}}$$

[MWA10] M. Merrett, et Al., "Design Metrics for RTL Level Estimation of Delay Variability Due to Int<u>radie (Random) Variations," ISCAS 2010</u>

Green IC 😭

Voltage/Temperature Variations

- Delay is very sensitive to V_{DD} and temperature
 - voltage: up to 30-50% margin
 - temperature: up to 100% margin
 - NT systems require adaptive schemes
 - sense V_{DD} and T and adjust clock cycle
 - can compensate slow variations (margin needed for fast)

Clock Cycle Margin Reduction/Elimination

Compensation of variations at different times

- tradeoff between energy cost of margining vs energy cost to reduce amount of margin
- large variations at NT/ST: runtime compensation usually required
- detect timing errors/correct: minimum or no margin at all

NUS

Margin Elimination: Timing Error Detection

 Reduce/eliminate worst-case margin by catching delay faults

correct at run-time, tune to compensate actual variations

In-situ monitoring

- P, V, T, aging, fast variations
- no margin
- invasive, limited tuning

Fault prediction (Tunable Replica Circuit)

- partially: P, V, T, aging, fast (not soft errors)
- needs some margin (false positives, mimics only critical path)
- ted tuning

 little invasive, tuning required, low overhead
 prof. Massimo Alioto
 43

Margin Elimination: Error Correction

Faults can be corrected at various levels

across-level design/optimization/control needed
 faster correction

energy overhead

$$E \neq E_{correction} \cdot error \ rate \neq f \cdot C \cdot V_{DD}^2$$

testing is painful (long, tuning)

Approximate Computing: Negative Margining

• Some apps do not need to have perfect computation

- approximate computing (deterministic, voltage overscaling)
 - ex.: multimedia, sensor fusion
 - avg error rate kept within bound (slow correction loop)
 - ex.: our first SRAM with Dynamically Adjustable Error-Quality

[FKB14] F. Frustaci, et Al., "A 32kb SRAM for Error-Free and Error-Tolerant Applications with Dynamic Energy-Quality Management in 28nm CMOS," ISSCC 2014

Functional Failures and V_{DD,min}

• Functional failures occur at very low V_{DD}

- flip-flops/latches prone to such failures (highest $V_{DD,min}$)
 - due to multiple connected outputs (current contention due to low

$$I_{on}/I_{off}$$

Green IC 🔛

[A12a] M. Alioto, "Ultra-Low Power VLSI Circuit Design Demystified and Explained: A Tutorial," IEEE TCAS-I, Jan. 2012.

- large worst-case variations
 due to high number of FFs
 (millions)
- V_{DD,min} dominated by variation induced p/n imbalance [A12a] 0.5 v_t

$$V_{DD,min} = n \cdot v_t \cdot \left[1 + \ln\left(\frac{2}{n}\right) + \ln(pn) \right] \Big|_{2 v_t}^{2.5 v_t}$$

• SRAM more vulnerable than logic

- read/write/hold margins set by strength ratio
- no averaging across multiple cells, as opposed to logic

• reducing $V_{DD,min}$ of SRAMs at different levels

- within the cell: V_{TH} adjustment, lithography-friendly layout, circuit (sizing, more robust topologies)
- **outside** the cell: array architecture, assist techniques

Green IC

Fine-Grain Adaptation

- Multi- V_{DD} with small area/energy overhead
 - avoid level shifters altogether [MYN11]
 - small voltage domains dynamically assigned
 - no level shifters (small voltage difference)

[MYN11] A. Muramatsu, et Al. ''12% Power Reduction by Within-Functional-Block Fine-Grained Adaptive Dual Supply Voltage Control (...),'' ESSCIRC 2011

Panoptic Dynamic Voltage Scaling [PDC09]

spatial and temporal fine granularity

Renus

Green IC

- sleep transistors (re)used to dynamically →
 select V_{DD} (workload)
- 34% (44%) energy saving over multi- V_{DD} (single V_{DD})

[PDC09] M. Putic, et Al., "Panoptic DVS: A Fine-Grained Dynamic Voltage Scaling Framework for Energy Scalable CMOS Design," ICCD 2009

С

V_{DD1}V_{DD2}V_{DD3}

в

Α

Selective boosting, state-retentive sleep [TRA14]

- Graphics Execution Core
- V_{DD} of register file/ROM boosted by 270 mV to improve $V_{DD,min}$
- adaptive clocking reacts to first V_{DD} droop: senses and divides f_{CLK} for a fixed time to recover (margin reduction)
- ◆ 4-20X register file leakage reduction in sleep mode through

voltage reduction down to Data Retention Voltage of bitcells [TRA14] C. Tokunaga, et Al. "A Graphics Execution Core in 22nm CMOS Featuring Adaptive Clocking, Selective Boosting and State-Retentive Sleep," ISSCC 2014

• Exploiting variations via cherry picking [RTG13]

- different cores: different energy-performance tradeoffs
 - redundant cores, choose most energy efficient

• 22% better performance at 33% dark silicon [RTG13] B. Raghunathan, et Al., "Cherry-Picking: Exploiting Process Variations in Dark-Silicon Homogeneous Chip Multi-Processors," DATE 2013

Conclusions

IoT Naturally Follows Historical Trends...

- Size is a technology driver for the IoT
 - Bell's law: 10-100X size reduction every 10 years
 - IoT should happen in this decade
- Energy is the bottleneck for size
 - Koomey's law [KBS10]: 2X / 1.6 years
 - ♦ 75X in 10 years (~4X from technology scaling)
 - rest of it must come from ckt/architecture/system
 - quicker development, more aggressive reduction:
 more innovation

Challenges and Ideas for IoT

- IoT needs ultra-low voltage operation
 - energy vs power, NT vs ST
- Issues and solutions at ULV
 - performance

- fine-grain selective voltage boosting, across-boundary design
- leakage
 - fine-grain VDD/power gating domains, across-boundary design
- variations and resiliency
 - run-time adaptation to eliminate cycle margin and compensate variations

THANKS FOR YOUR ATTENTION

Massured

BACKUP SLIDES

Typical performance and power specs

throughput: tens-hundreds
 of MOPS down to kOPS
 (10X slower or more than mobile)
 [UAK12] K. Uchiyama, et Al., Heterogeneous Multicore
 Technologies for Embedded Systems, Springer, 2012

- achievable through ULV (5X), architecture tailoring (5-10X), specialized HW (10X) w.r.t. high performance
- not much gain from technology scaling...