A Scalable Distributed Regulator Architecture for the POWER8™ Microprocessor

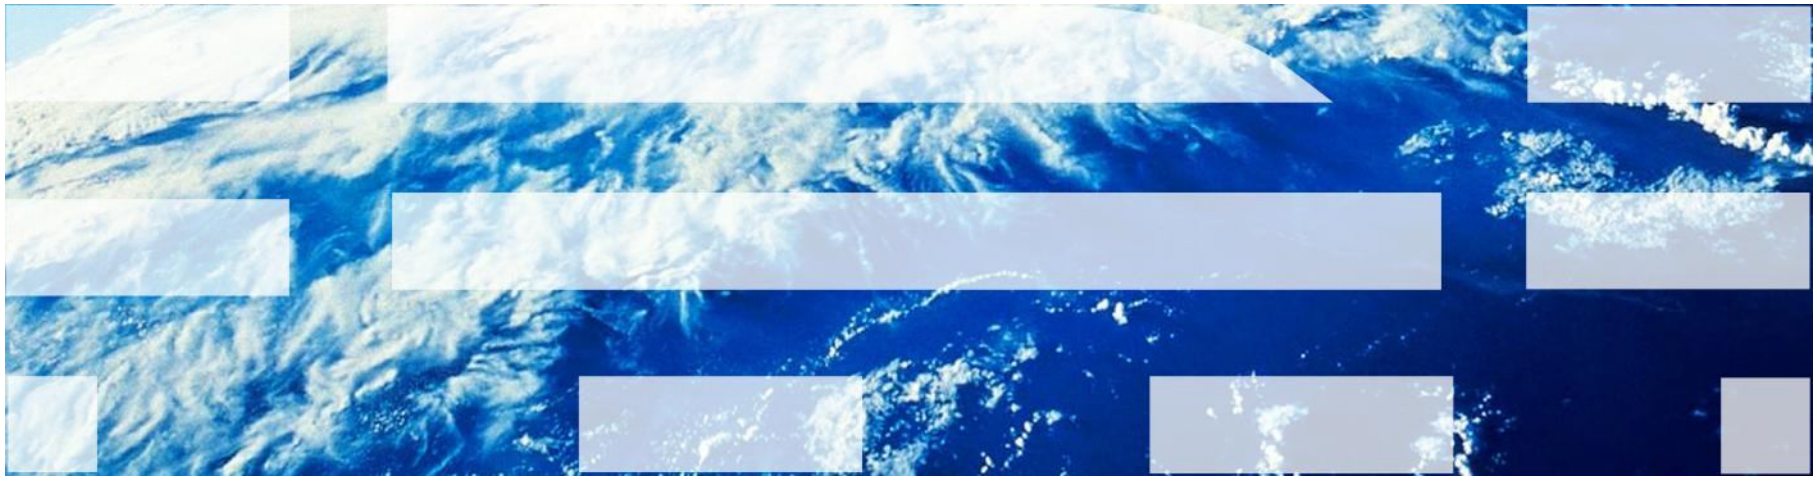

Zeynep Deniz et al.
IBM Research

Outline

- Motivation
- Distributed architecture for integrated voltage regulator (iVRM)
- Challenges of iVRM scaling
- POWER8™ iVRM design
- Experimental results
- Conclusions

Motivation

- Benefits of per-core DVFS
  - Performance optimization of multi-core processors (e.g., maximum total power constraint)
  - Minimize energy/operation for portable systems
- Effectiveness of DVFS depends on how fast voltage transitions are
  - Integrated Voltage Regulators (iVRMs) provide benefit of sub-µs voltage changes at low cost
- Key challenge for iVRM: load current can change from min to max in only a few hundred ps

Choice of Regulator Topology

- iVRM requirements:
  - High power efficiency
  - High current / power density
  - Fast output voltage changes
  - Fast load response time
- Inductor-based [1] and switched-capacitor converters [2] have not yet shown the power density needed for server-class microprocessor (i.e. I_{load} >100A)
- Linear “op-amp”-like regulators [3] have slow response time
  - Alternative: comparator-based (bang-bang) regulators

[1] E. A. Burton et al., Proc. IEEE APEC, Mar. 2014
[2] S. R. Sanders et al., IEEE Trans. Power Electron., Sept. 2013.
[3] M. Saint-Laurent et al., IEEE ISSCC, Feb. 2014.

Comparator-Based Regulator

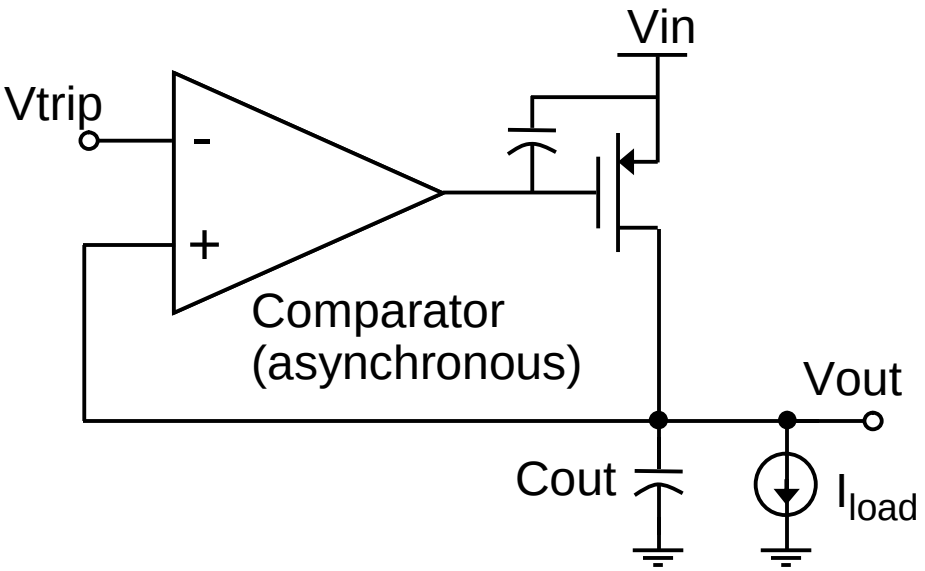

- **High-speed asynchronous comparator amplifies error signals to binary (rail-to-rail) levels**
- **"Bang-bang" control of PMOS passgate**

Benefits of high-gain bang-bang control loop:

- **DC load regulation better than 10 mV**
- **Fast load response time (T_R < 500 ps)**
- **Power-efficient pre-driver (CMOS inverters)**

Challenges/drawbacks:

- **Self-generated ripple at 1-2 GHz (need to minimize)**
- **DC offsets of fast comparators (need to compensate Vtrip)**

Earlier Work: Distributed iVRM for DDR3 I/O

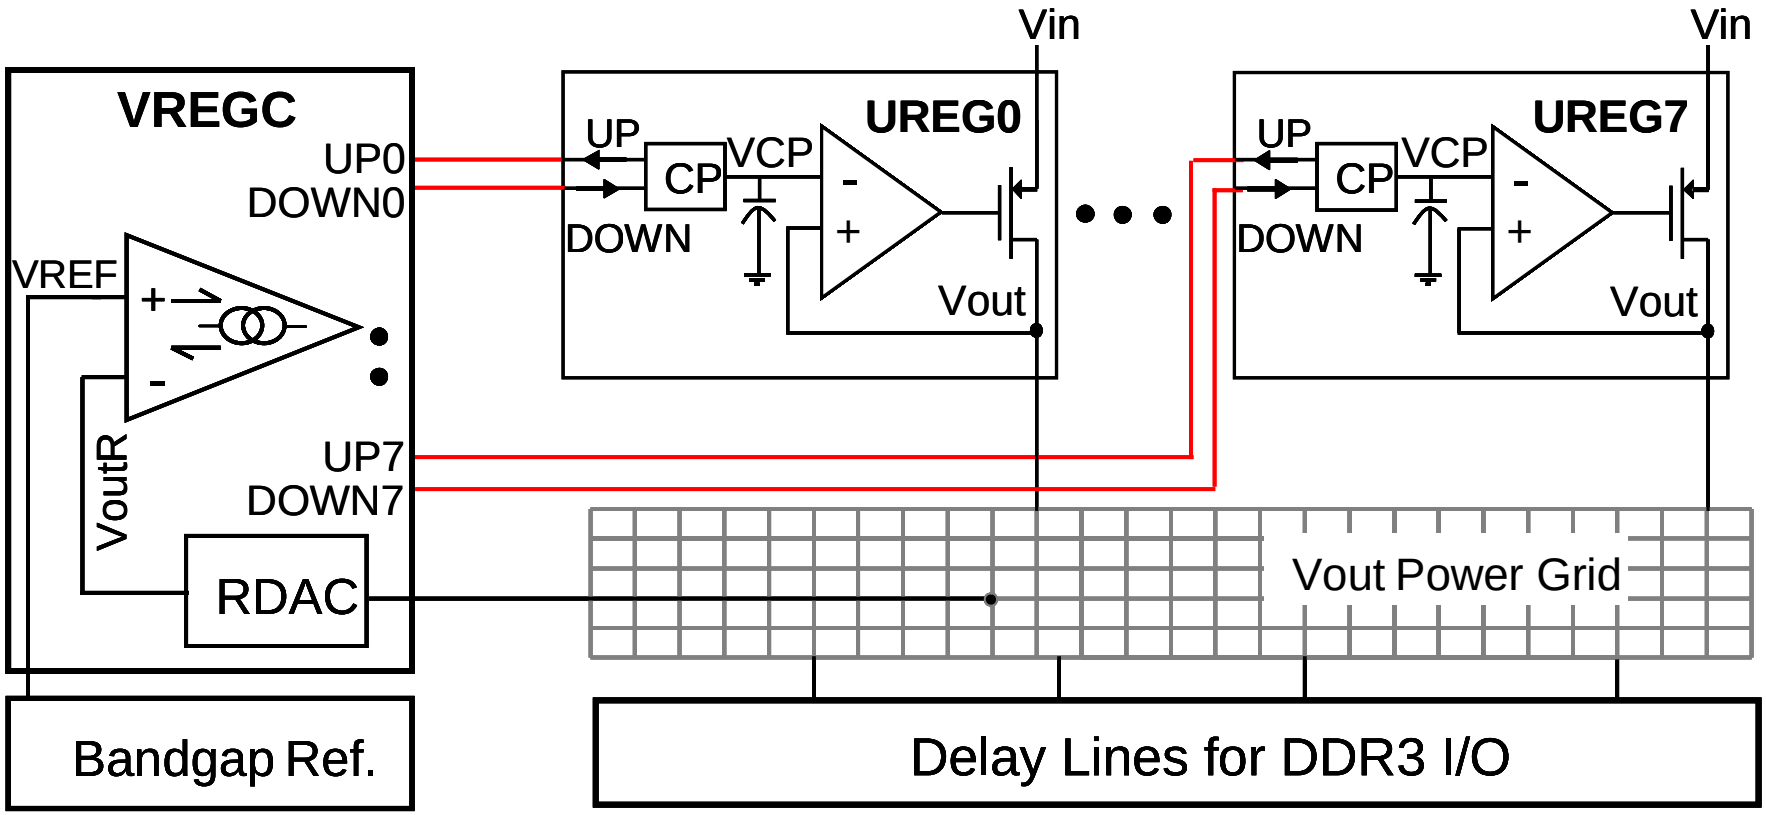

- **8 microregulators (UREGs) supply current to common power grid**
- **Trip point of UREG comparator tuned by local charge pump (CP)**
- **Central voltage regulator controller (VREGC) provides feedback to charge pumps in form of UP/DOWN currents to form outer loop**

[J. F. Bulzacchelli et al., *JSSC*, April. 2012]

Scaling Challenges: Distribution of UP/DN Currents

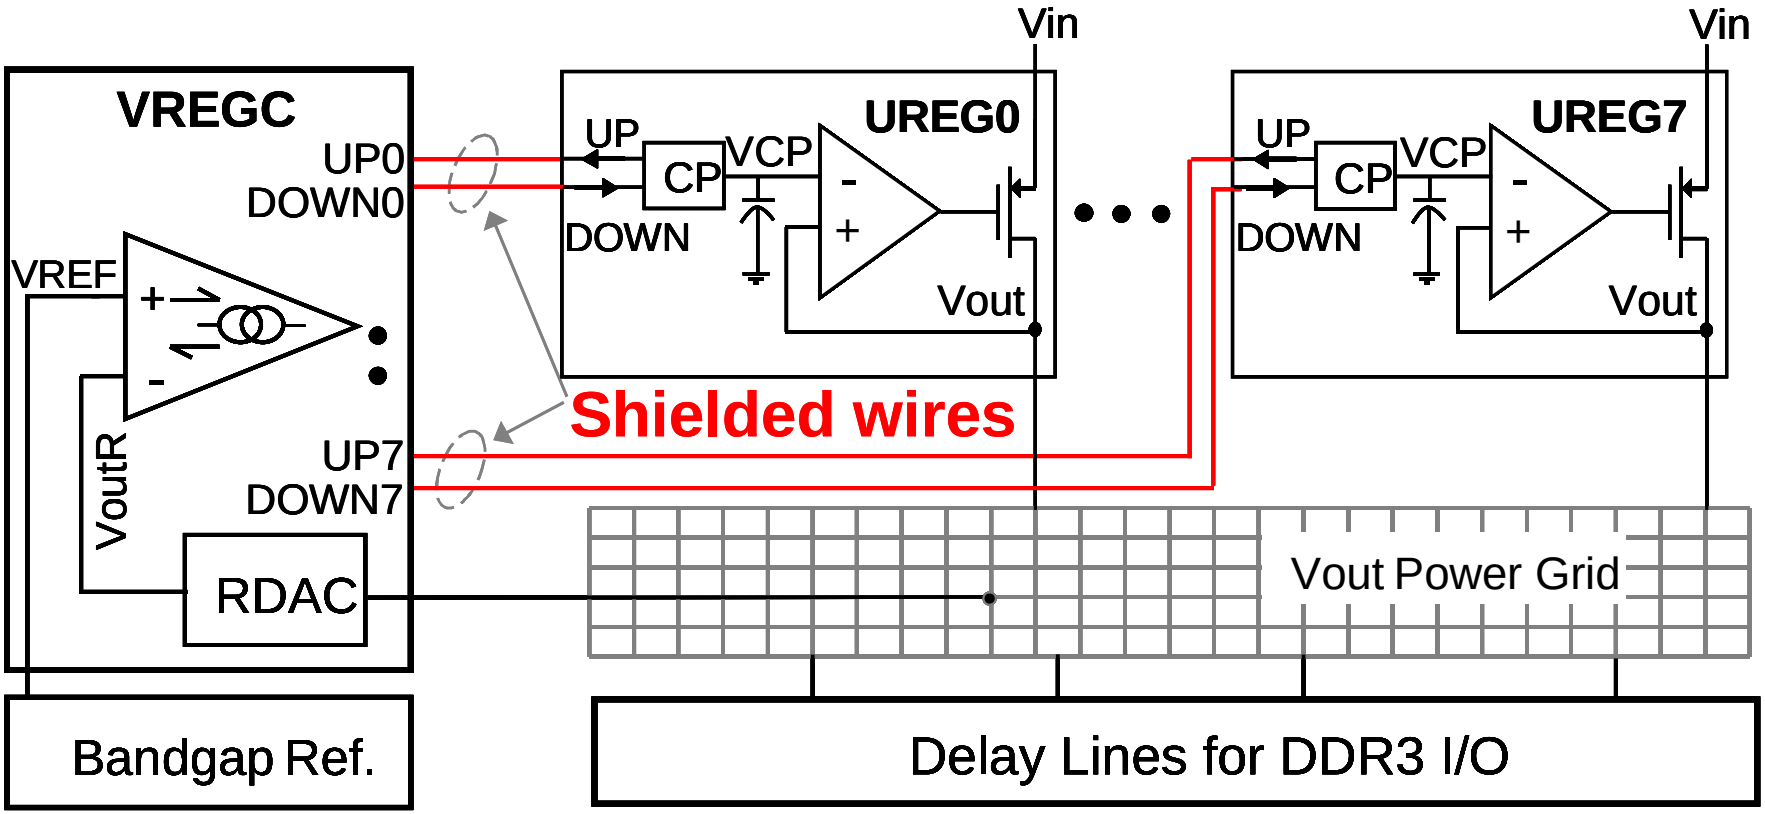

- **Point-to-point connection of analog UP/DN currents**
  - Shielded wiring is required
- **Manageable for 8 UREGs but not practical for >100 UREGs**

Scaling Challenges: GM-Amplifier

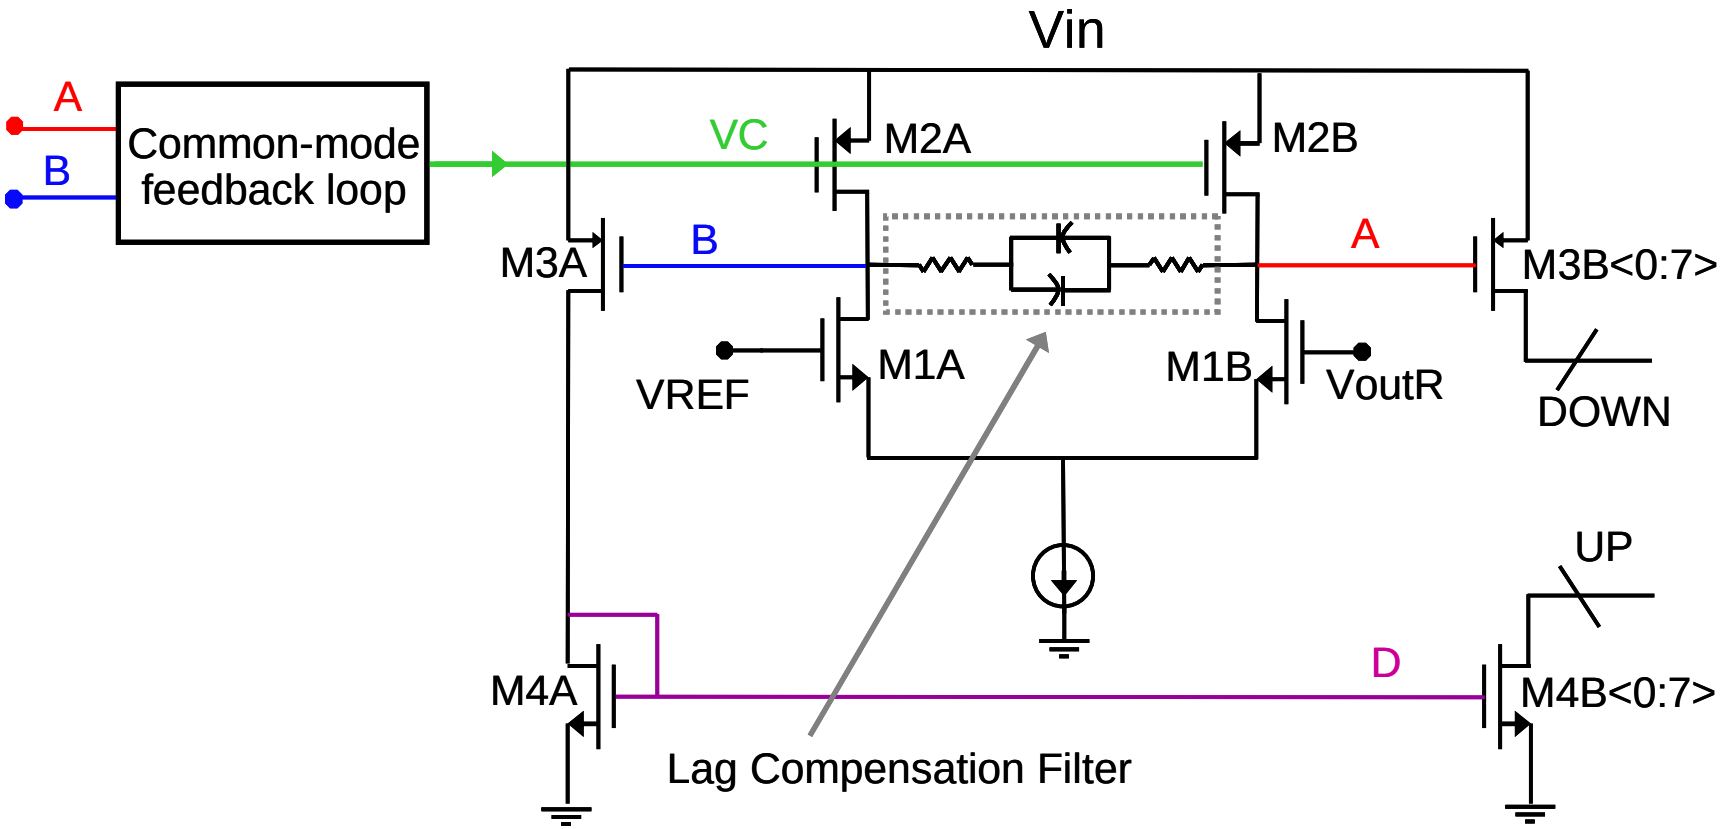

- No simple way of replicating/buffering analog UP/DN currents (w/o degrading accuracy)
- Need to redesign VREGC for each system (different number of outputs)

UREG with Hybrid Fast-Slow Passgate Control

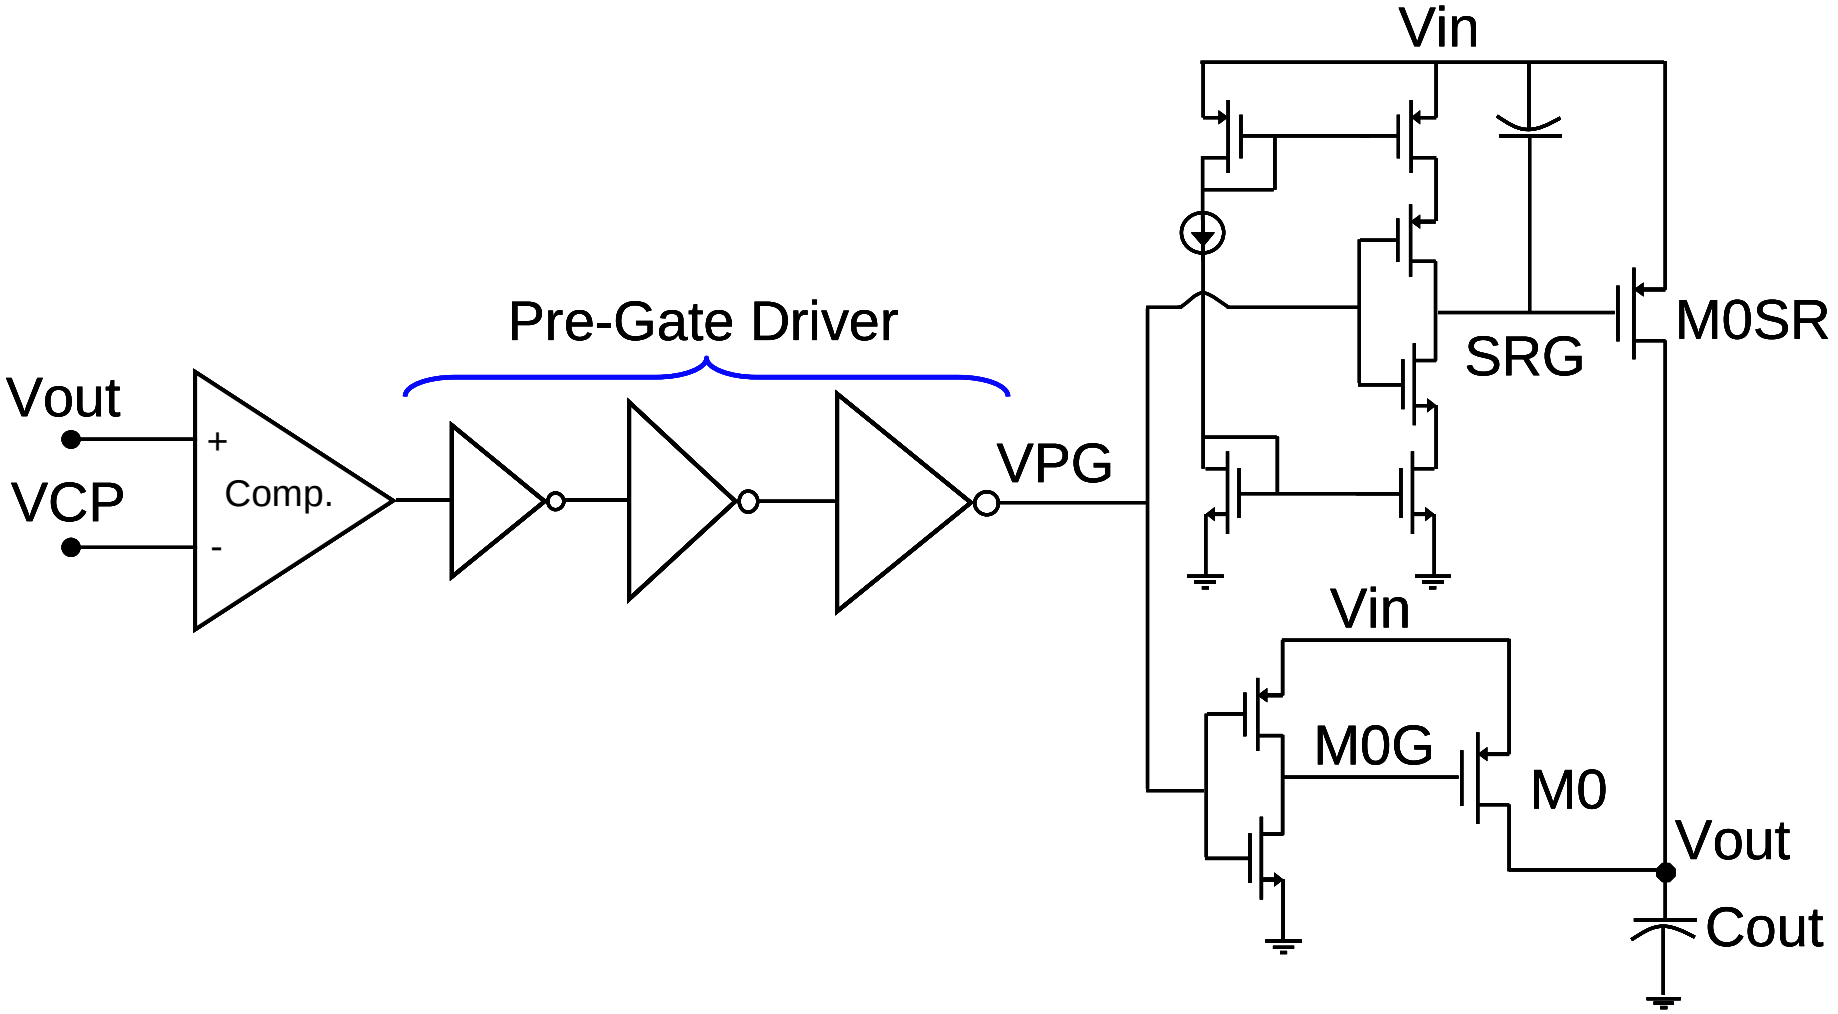

- **Issue:** excessive ripple if M0 is sized to supply total load current
- **Solution:** introduce M0SR to handle DC portion of load current

Master-Slave Control of M0SR Passgates

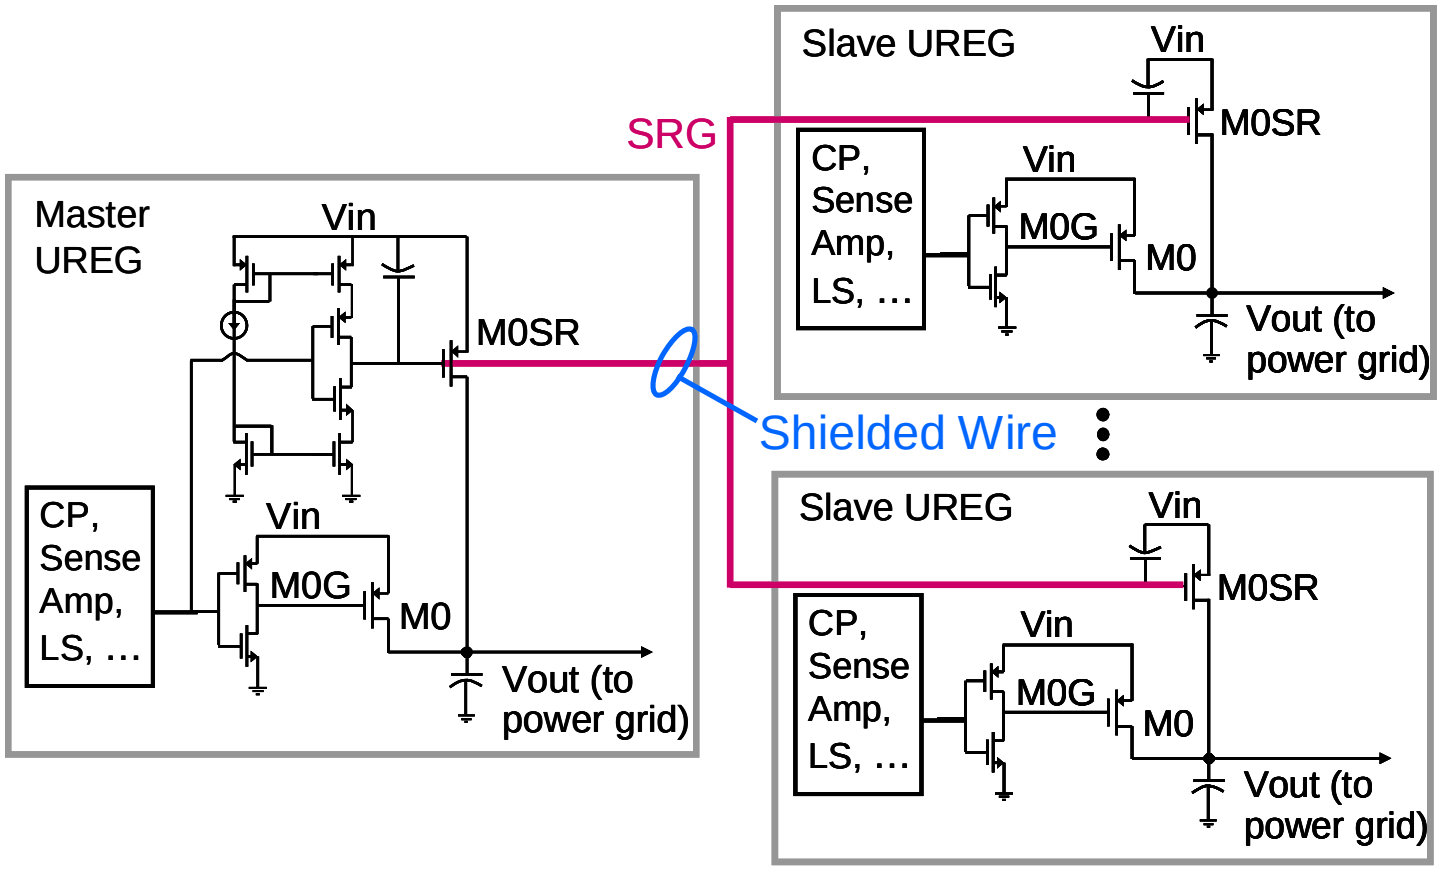

- **Issue:** current-starved inverters can amplify small mismatches in UREG duty cycles into large mismatches in M0SR currents
- **Solution:** use only one current-starved inverter (in "Master")
  - Not attractive for large-scale system, e.g. processor

Scaling Limitations of Previous Design

- **Point-to-point distribution of analog UP/DN signals**
  - Requires shielded wiring for low noise
- **Custom VREGC design optimized for given number of outputs**
  - Requires re-design with different number of UREGs
- **Master-Slave configuration for slow PFET control is not viable for large scale-ups**
  - Low noise distribution of analog gate voltage is difficult in processor environment

Outline

- Motivation
- Distributed architecture for integrated voltage regulator (iVRM)
- Challenges of iVRM scaling
- POWER8™ iVRM Design
- Experimental results
- Conclusions

POWER8™ Voltage Domains

- 4 regulated domains per chiplet (48 total)
- Each iVRM features a Bypass mode (BP)

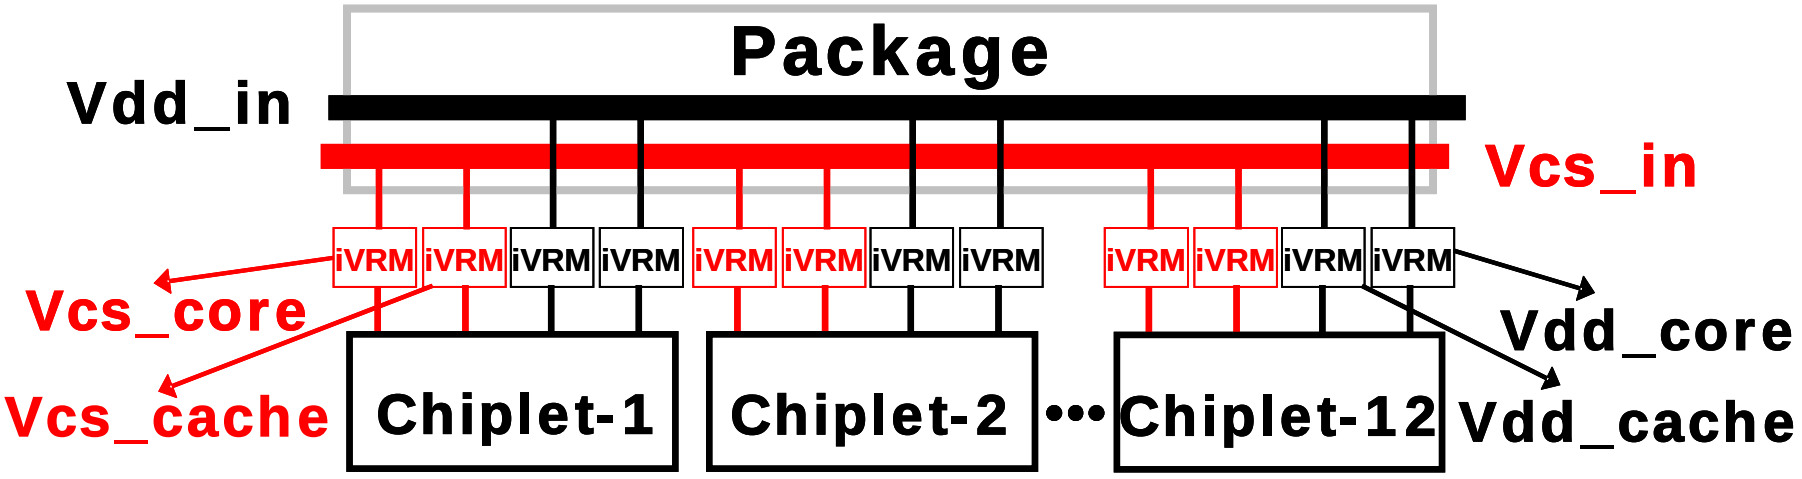

Scalable Distributed iVRM Architecture

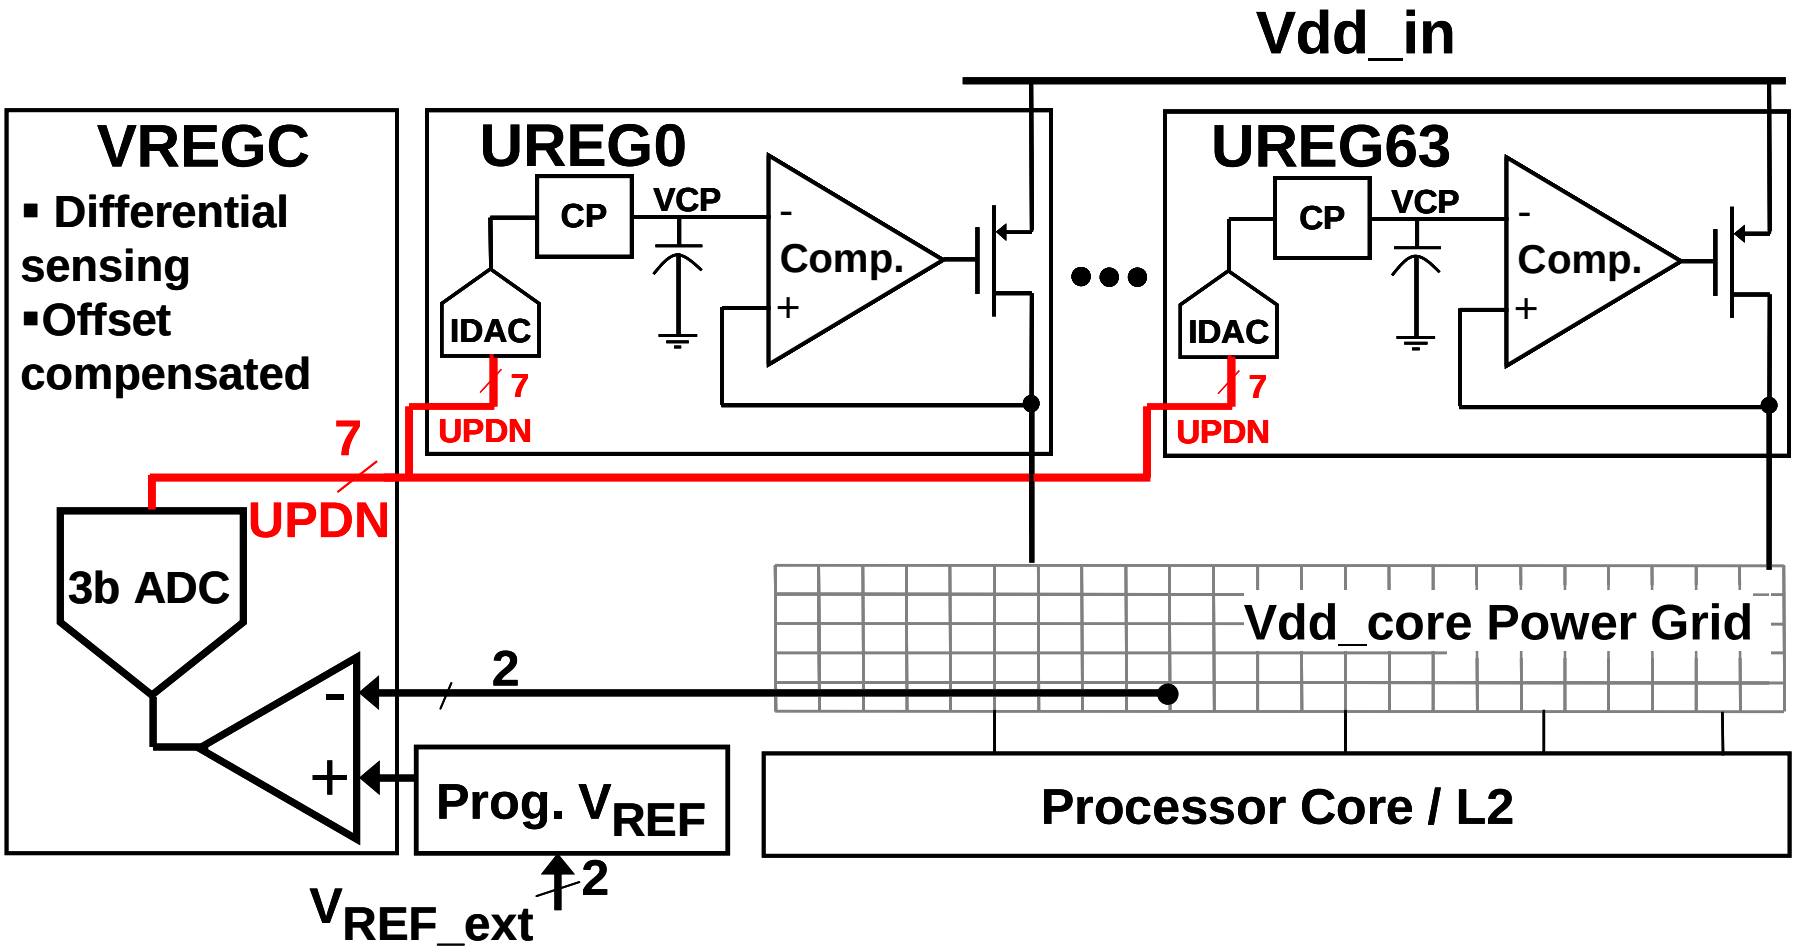

- Analog UP/DN currents replaced by digital UPDN codes
- External reference voltage used for high precision

[Z. Deniz et al., *ISSCC*, Feb. 2014]

Voltage Regulator Controller (VREGC)

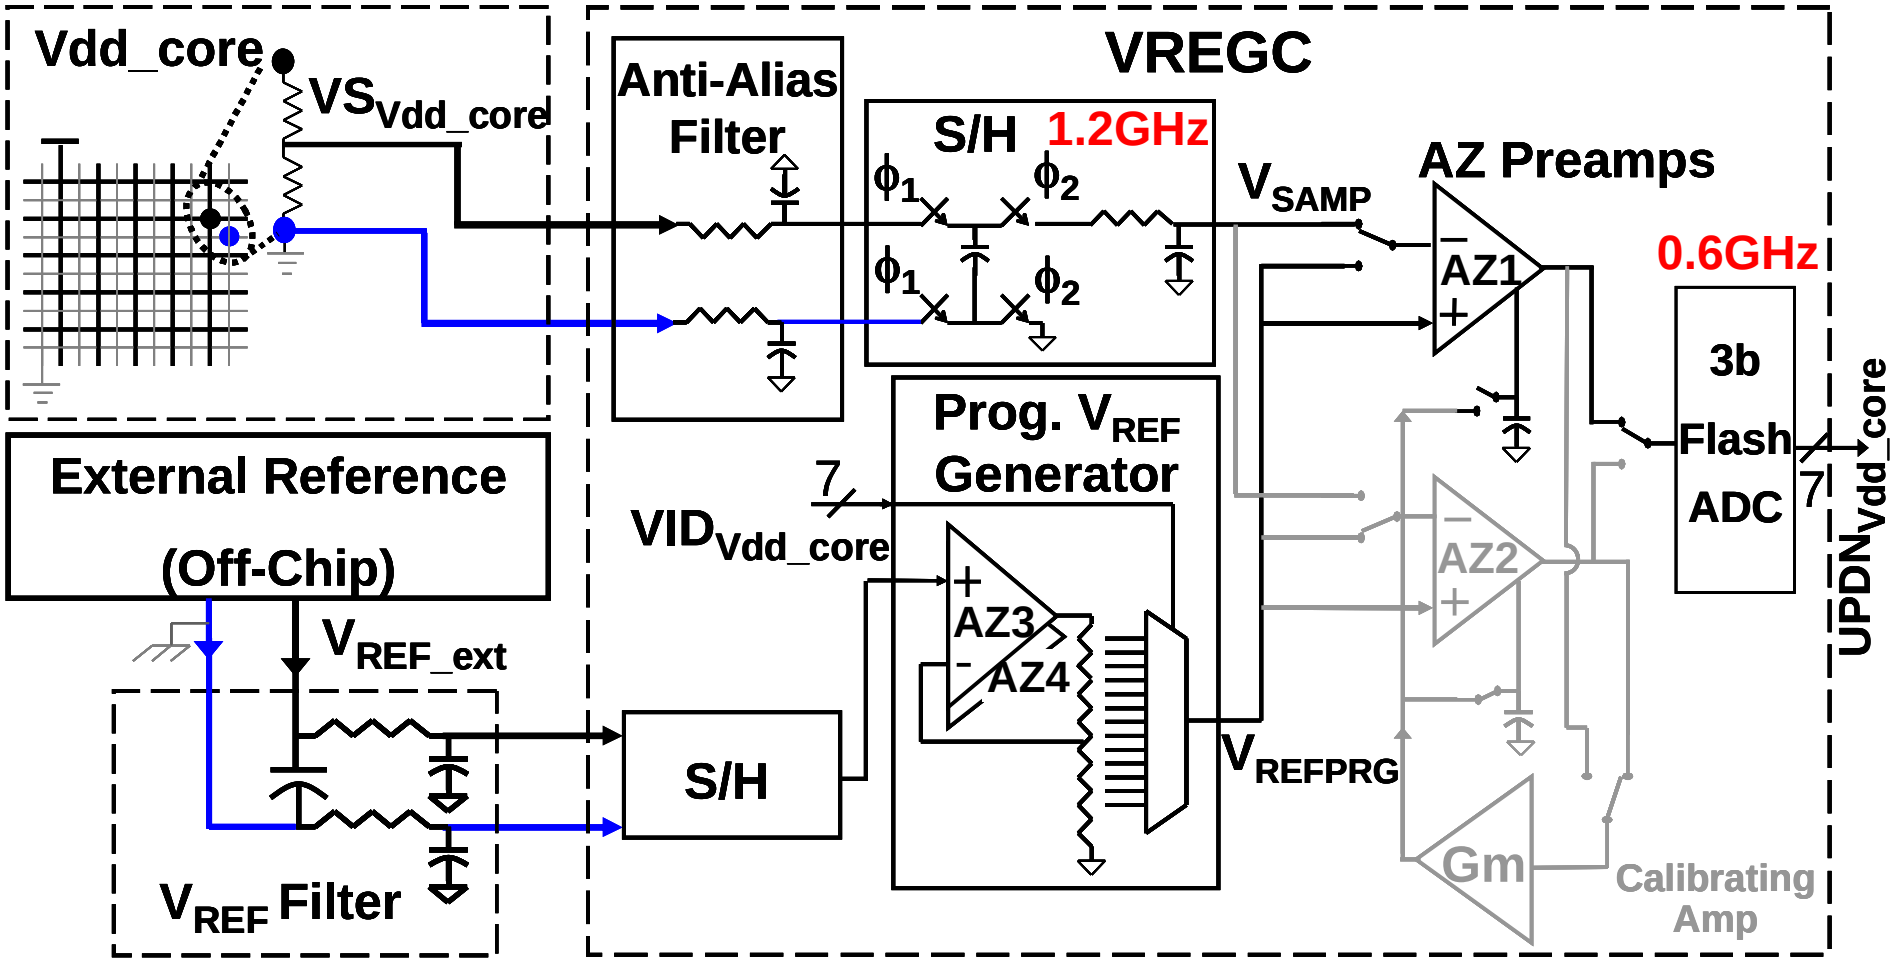

- Regulated voltage is filtered before sampling to avoid aliasing of high-frequency noise/ripple
- S/H also accomplishes differential to single-ended conversion

Voltage Regulator Controller (VREGC)

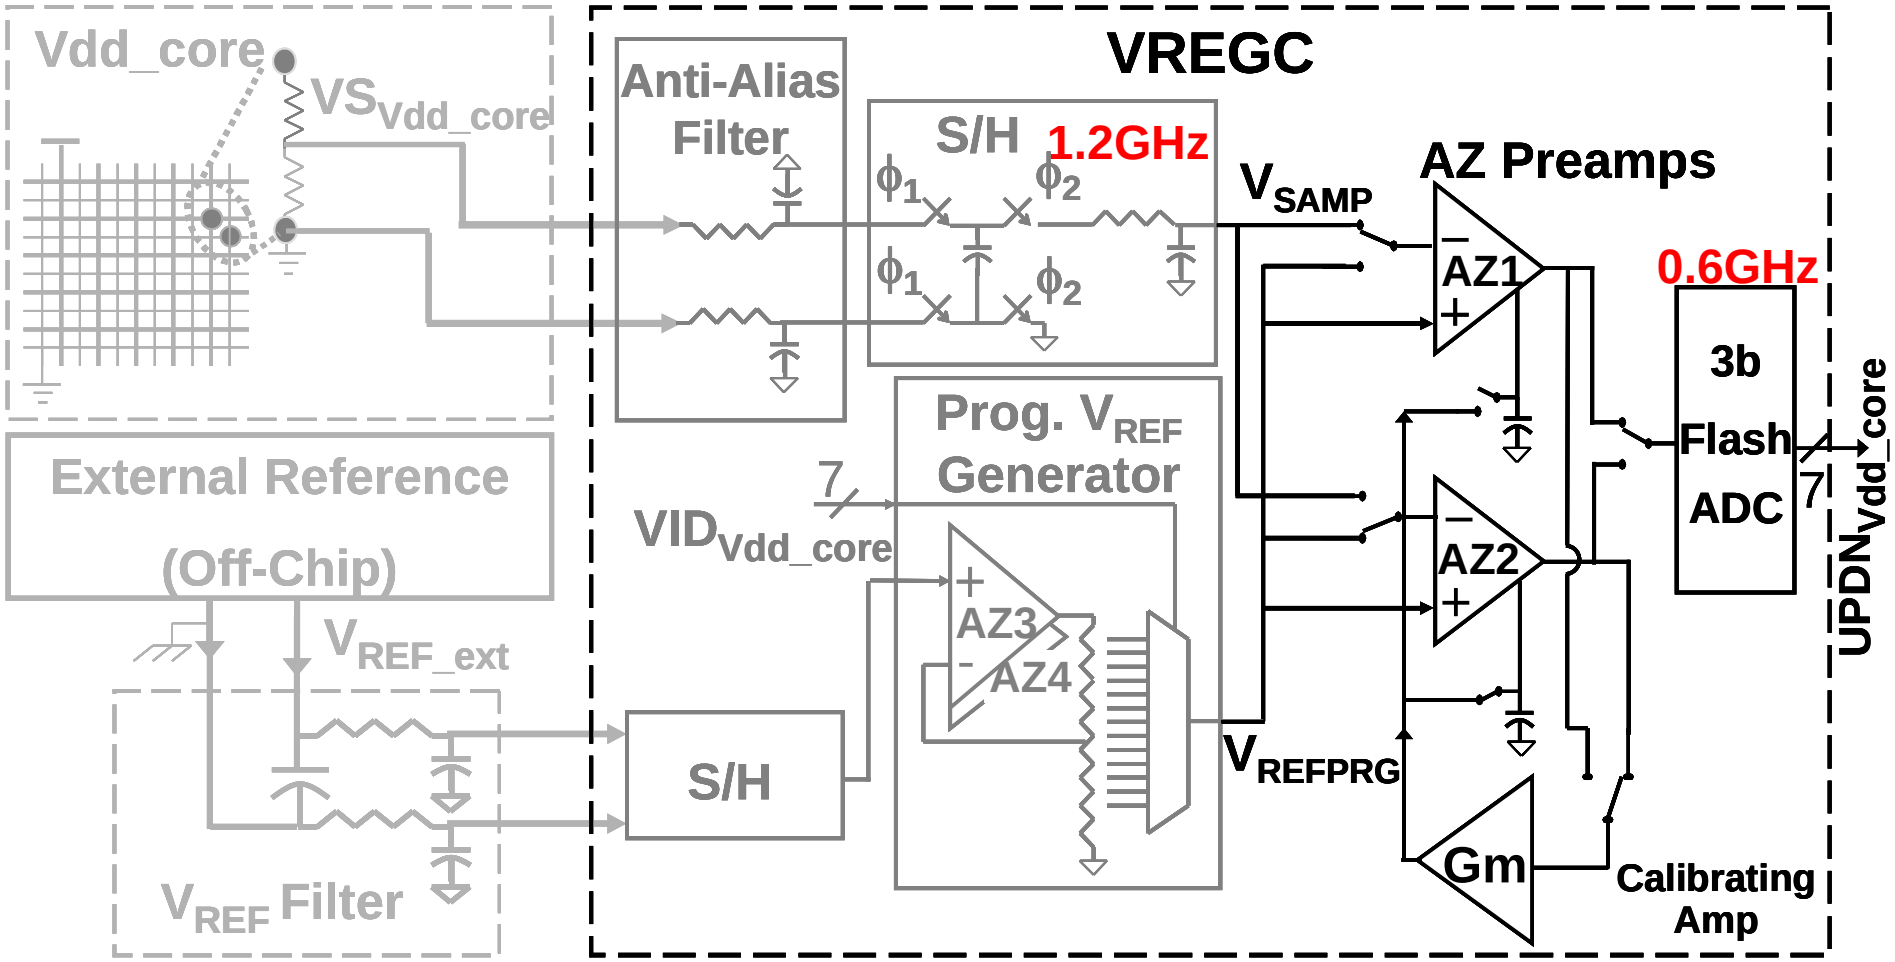

- **Ping-pong architecture for auto-zeroing preamplifiers**

Schematic of Basic UREG

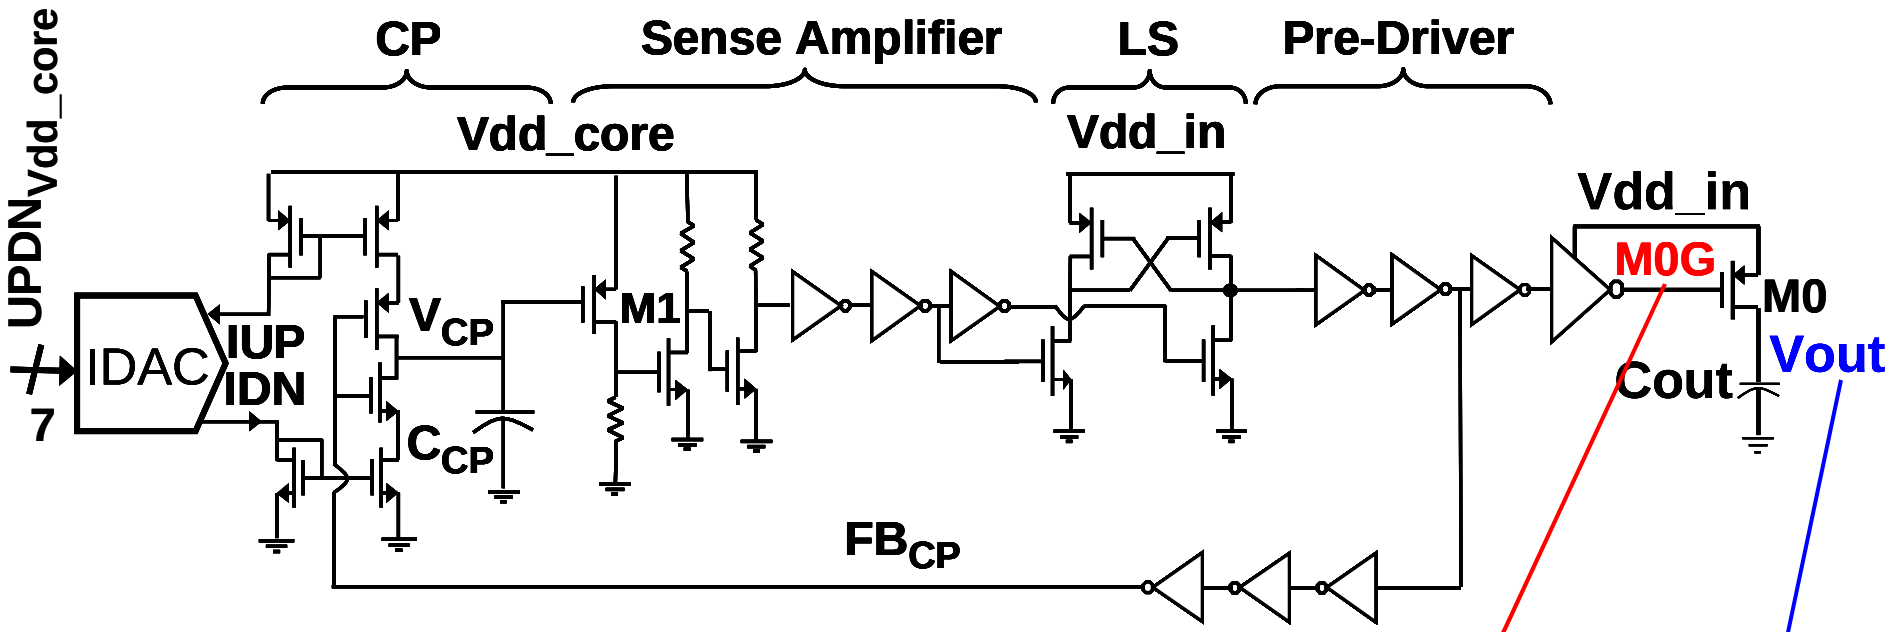

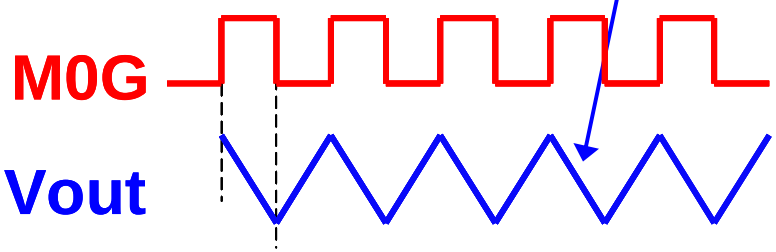

- **Sense amplifier chain powered off Vout ➔ excellent PSRR**
- **Charge pump is balanced when IUP/IDN=D/(1-D)➔ equal load sharing**

Slow Passgate Driver Implementation

- **Gate control for slow passgate is generated by local RC filter**
- **Active PFET (both fast & slow) strength calibration**

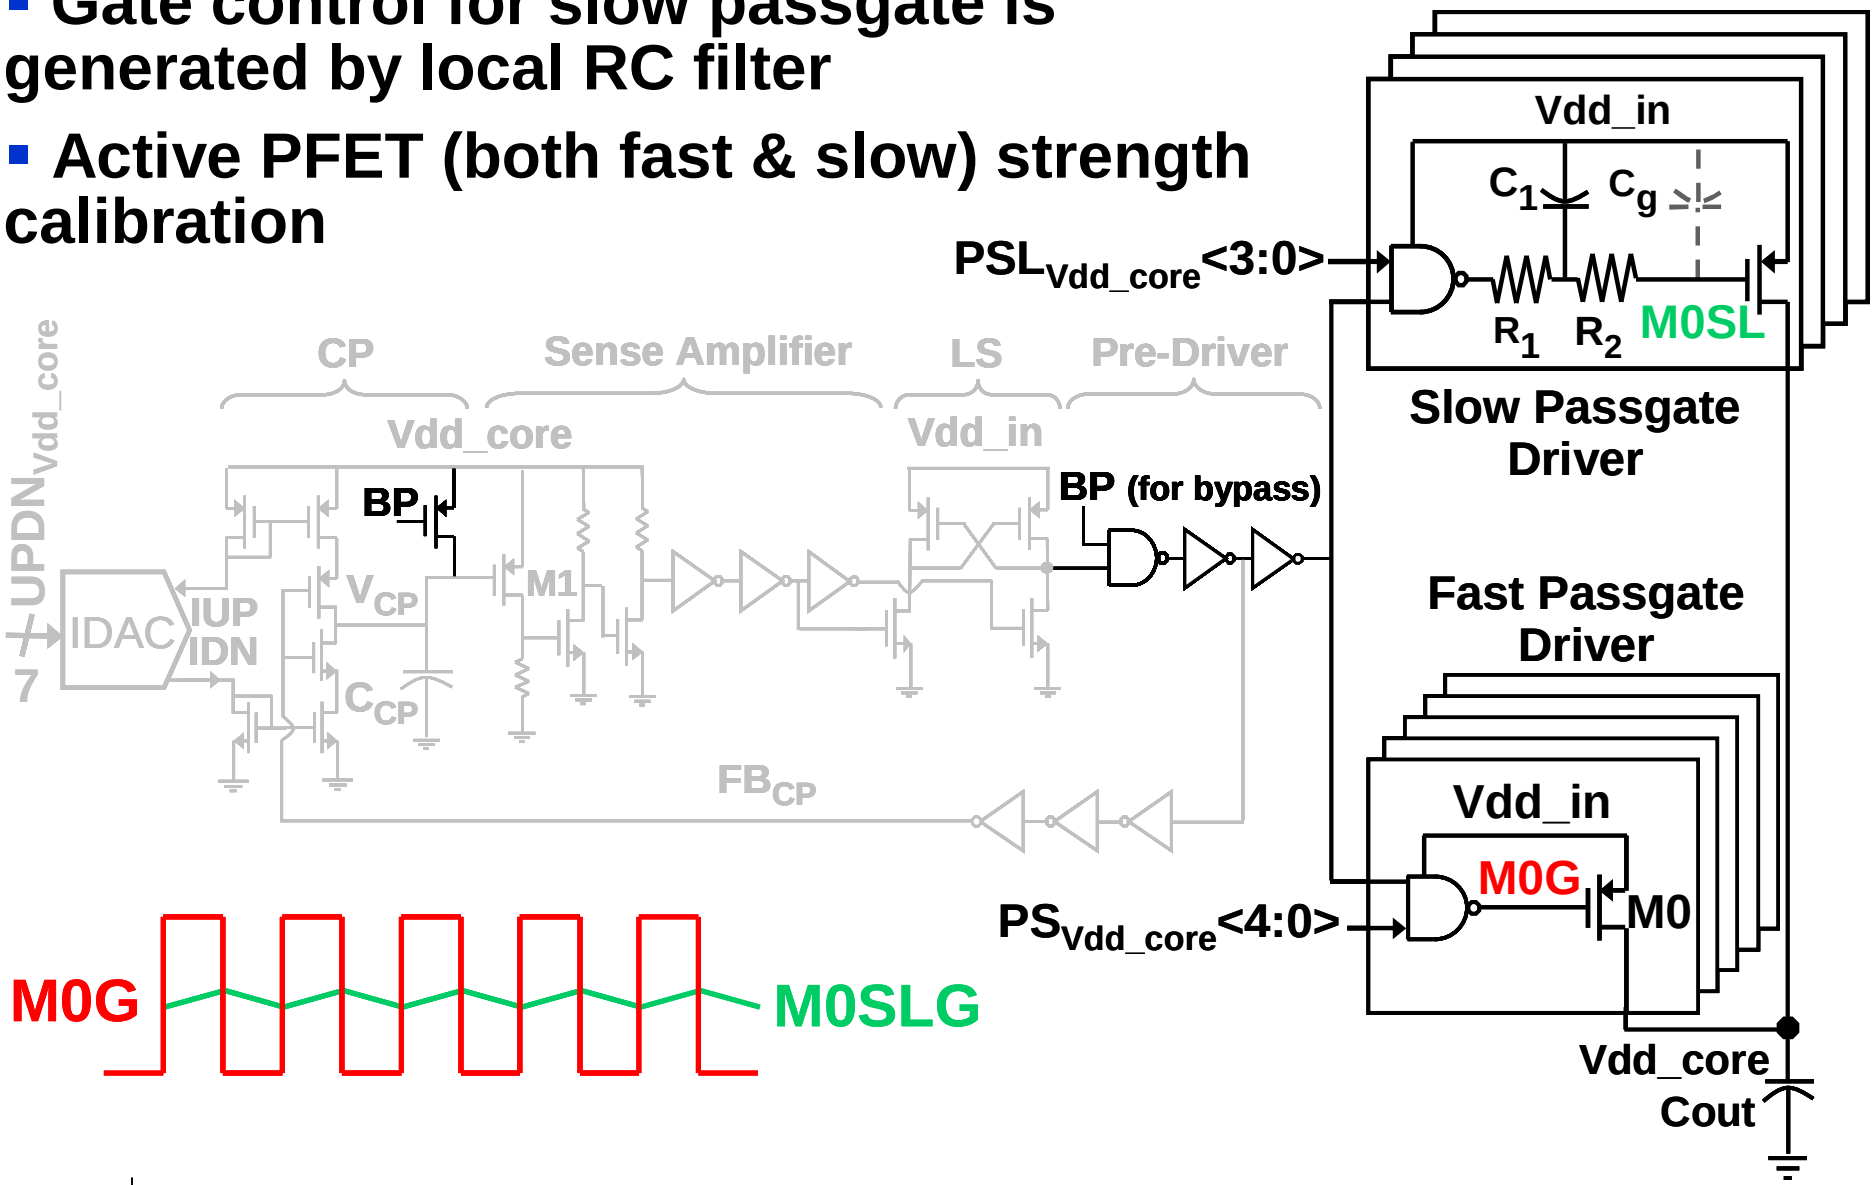

Predictive PFET Strength Calibration

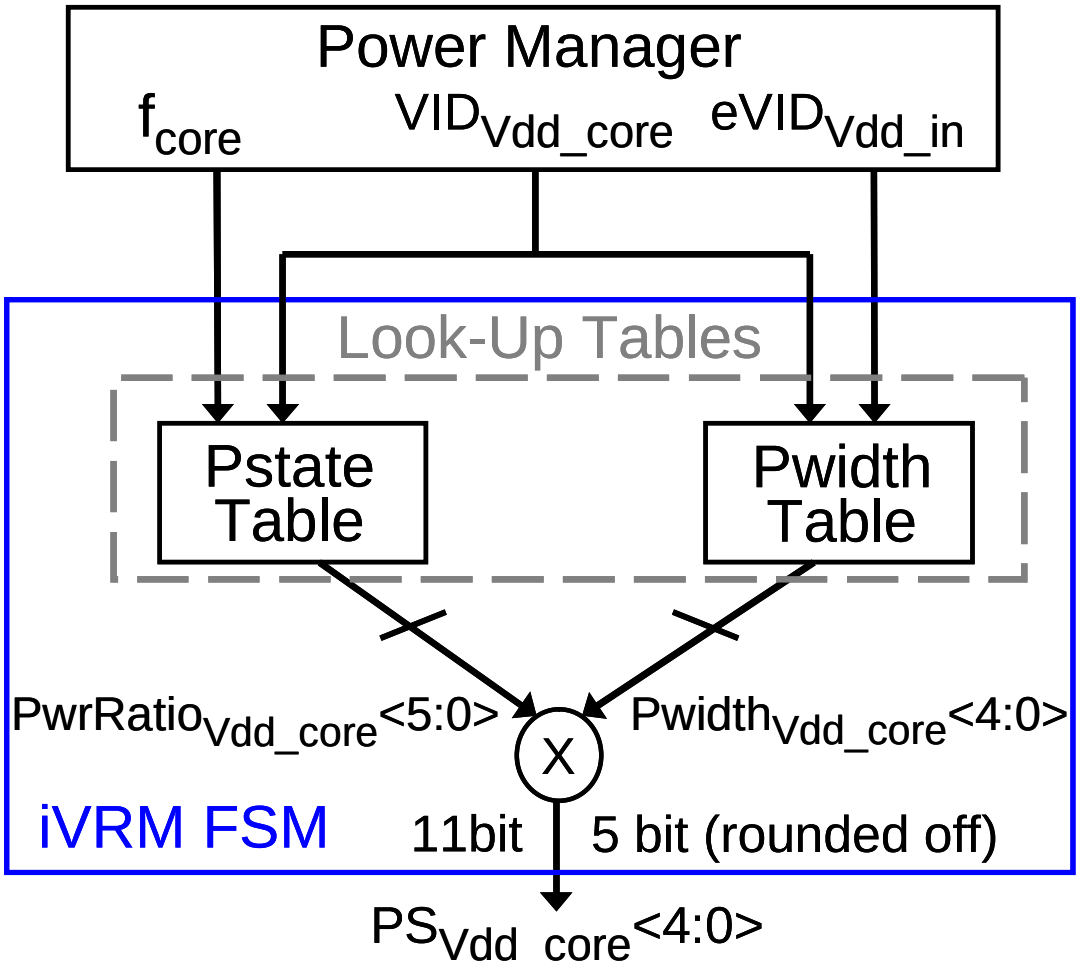

- **Pstate Table: stores max. load current information as function of output voltage and f_{core}**
- **Pwidth Table: records passgate width needed to support given drain current as function of input and output voltages**

PFETStrength = PstateTable • PwidthTable

Micrograph of POWER8™ Chiplet

- **22nm SOI CMOS**
- **UREG: 20µm X 64µm (each)**
- **VREGC: 150µm X 150µm (each)**
- **Vdd_core iVRM overhead <1% of chiplet area**
  - **Includes Cin (e.g., 90nF for Vdd_in)**
  - **Excludes Cout (e.g., 750nF for Vdd_core)**

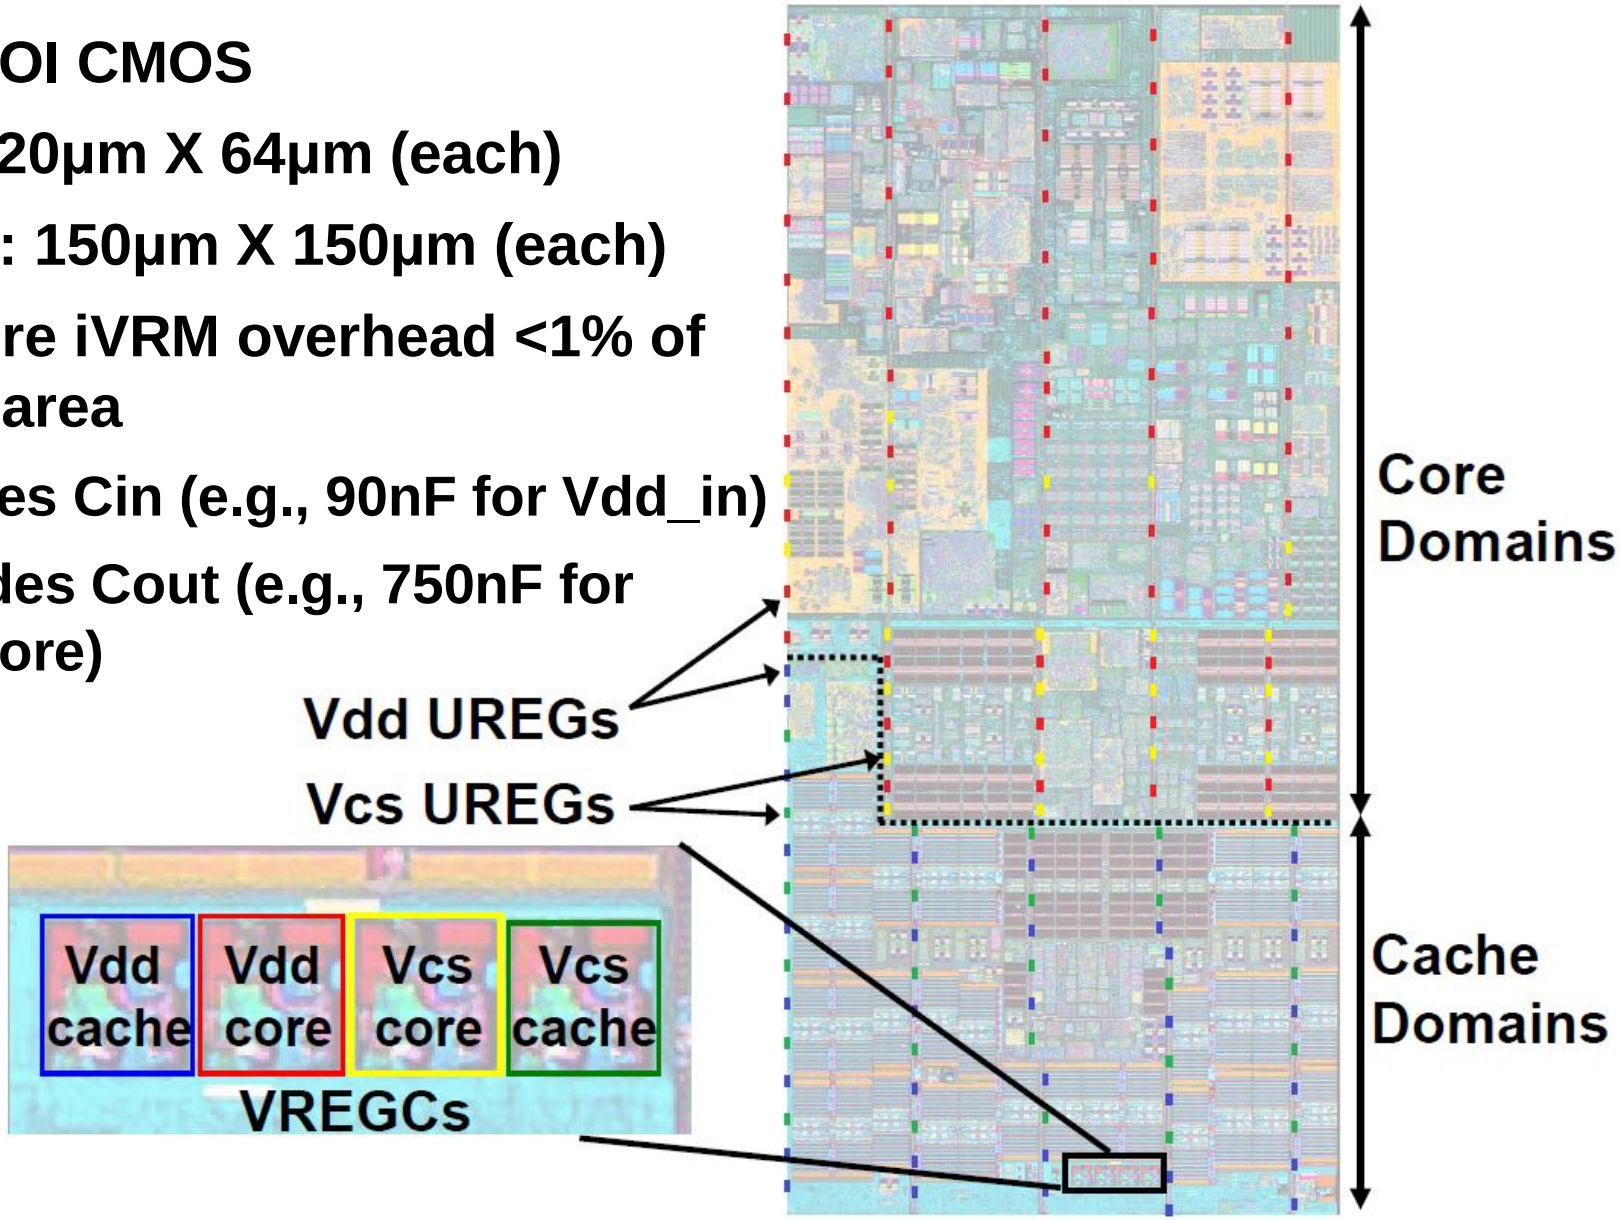

Measured DC Voltage versus VID

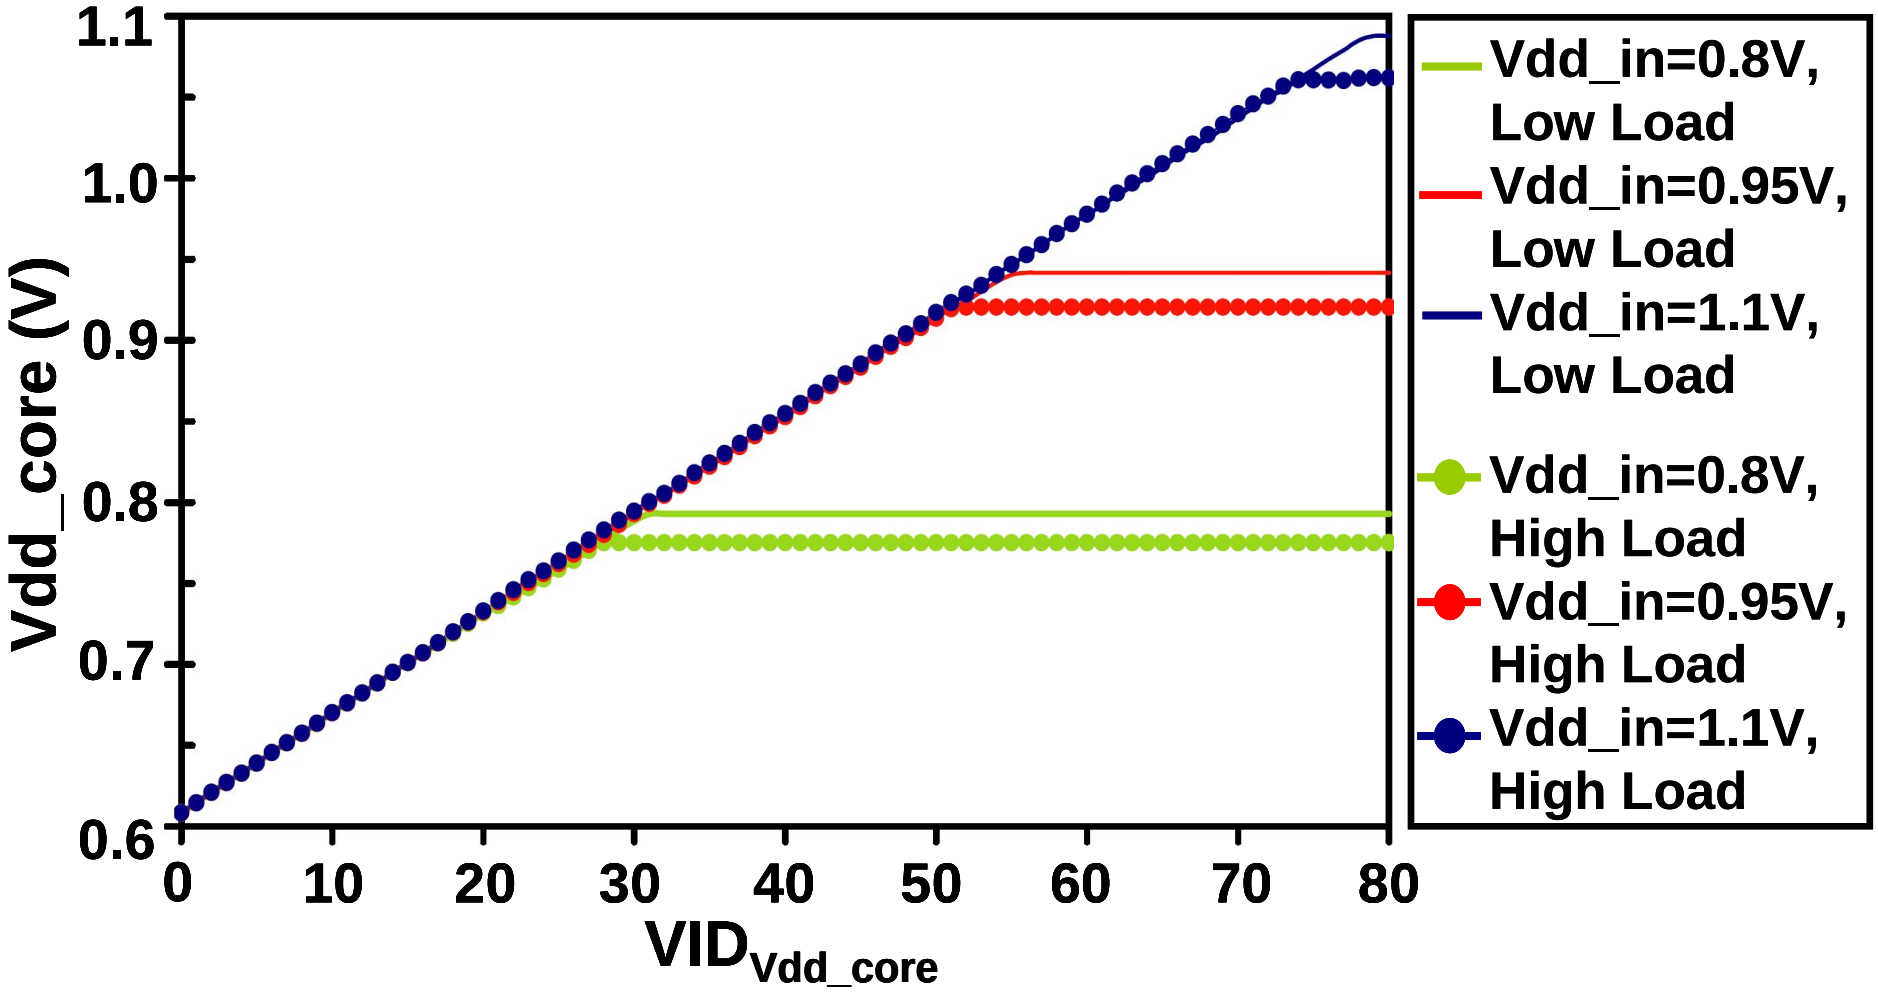

- **Dropout voltage < 50mV**

Measured Absolute Voltage Error

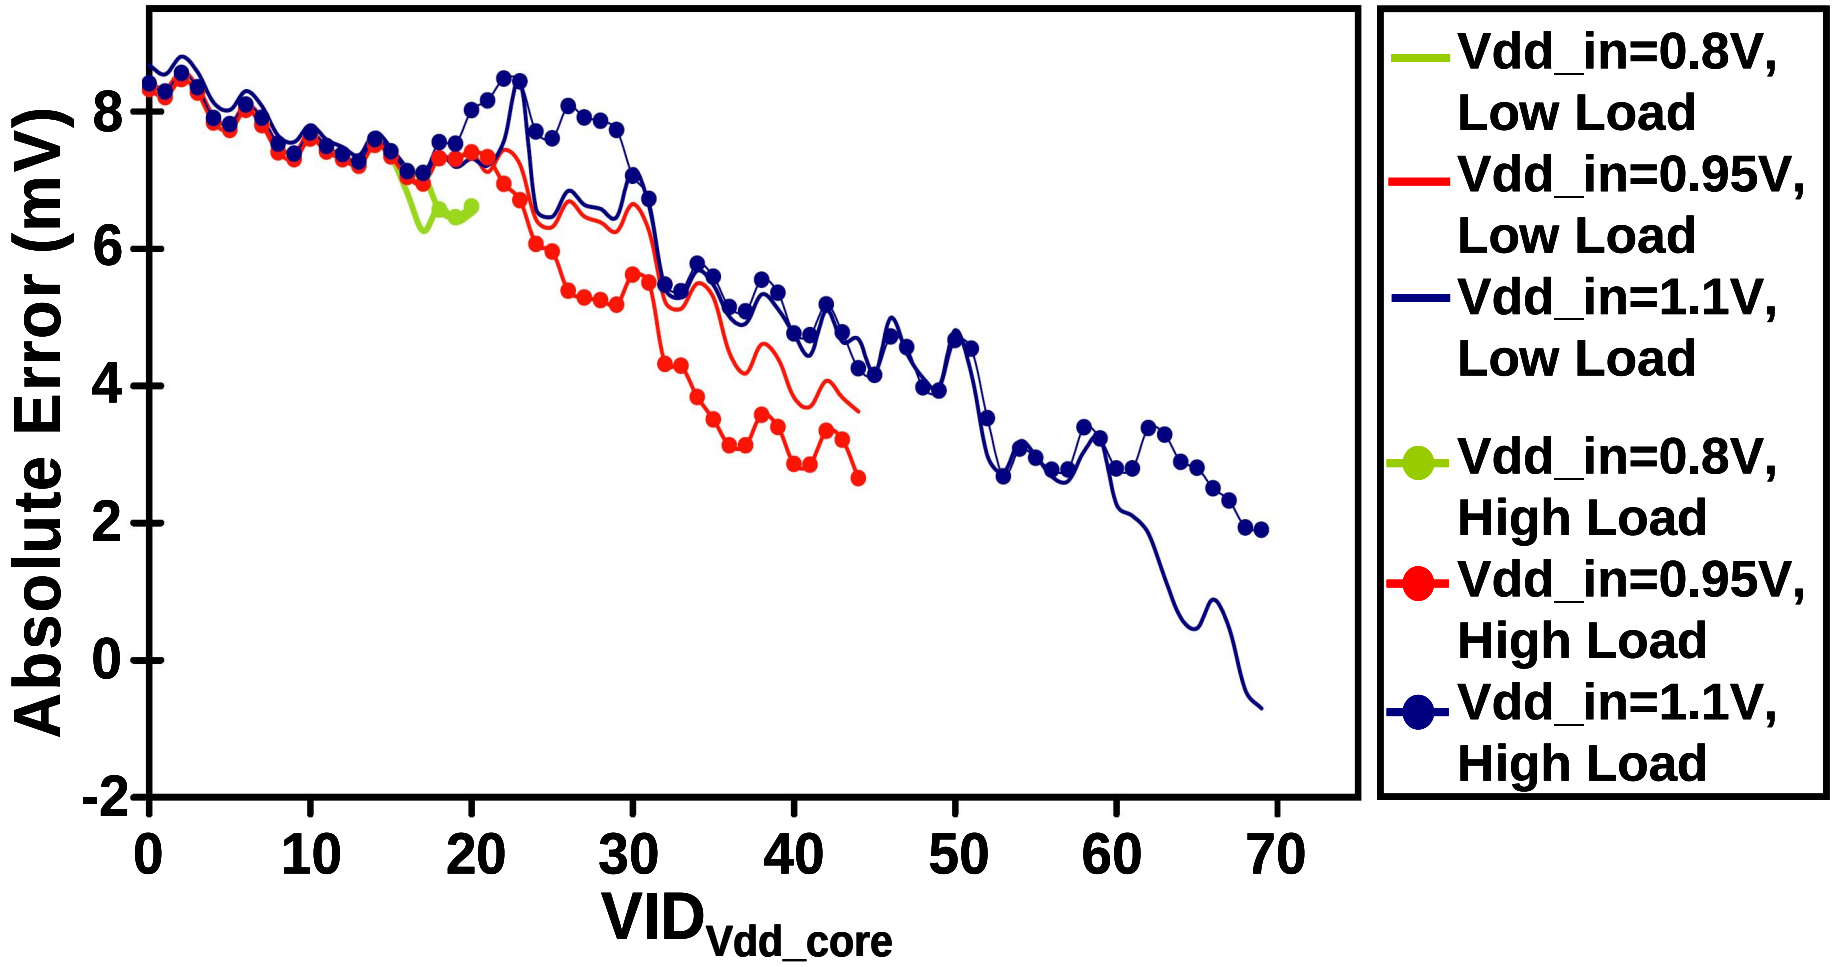

- Absolute voltage error < 9mV
- Load regulation error < 3mV
- Variation with Vdd_in < 5mV

Output Voltage Stepping

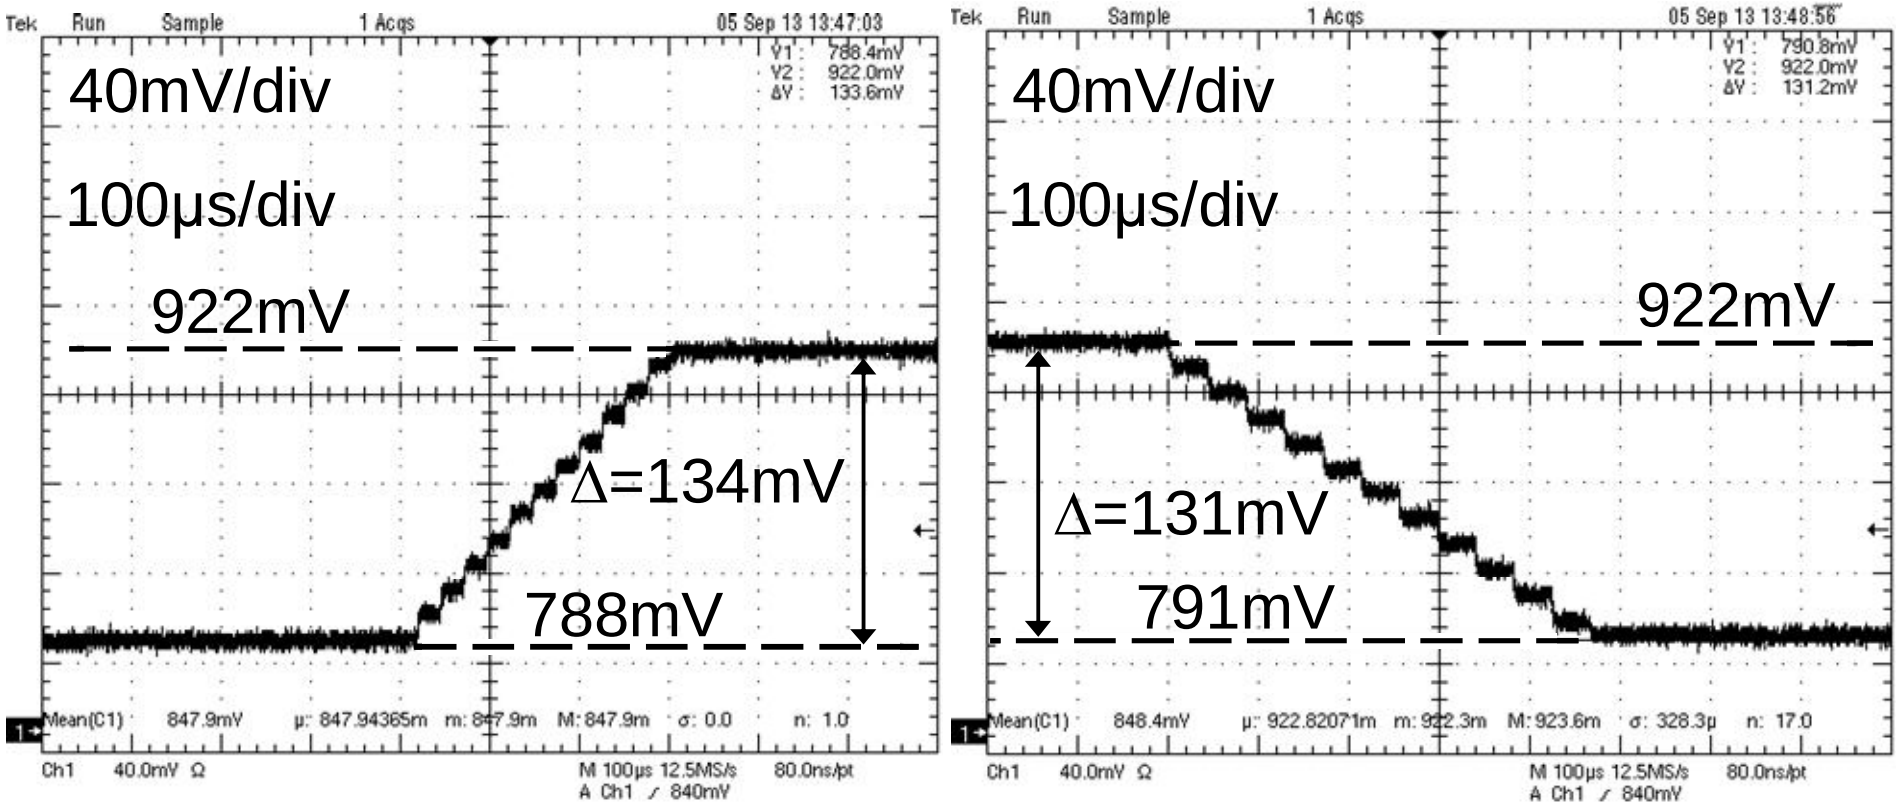

- **Small steps (e.g., 12.5mV) ensure tracking between domains**

Power Efficiency (under High Load)

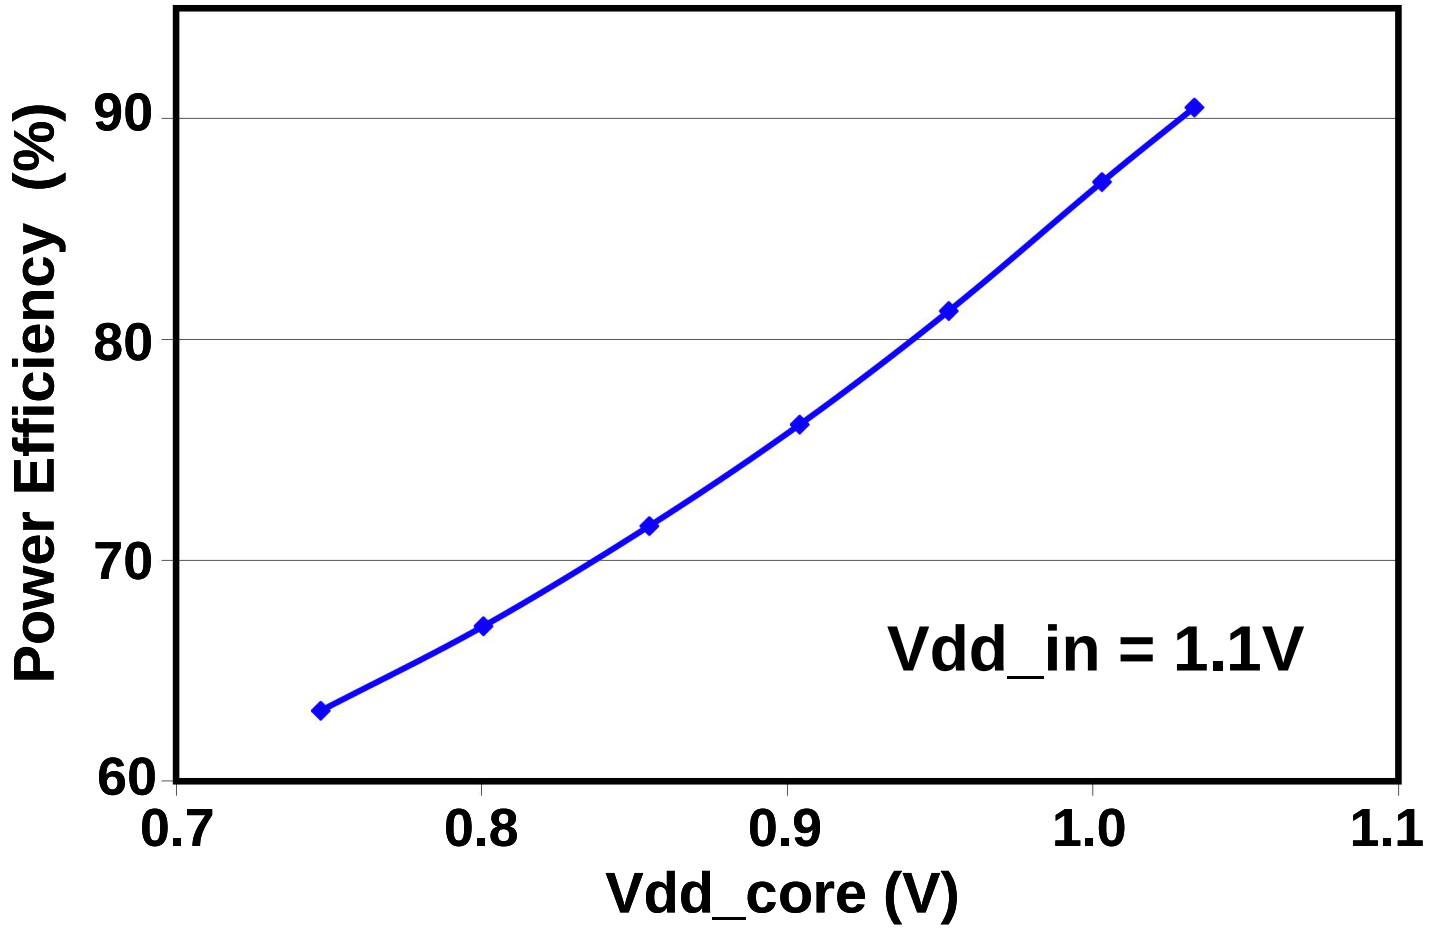

- **Peak power efficiency of 90.5% at Vdd_core =1.03V**
  - **Load current = 11.9A**
  - **iVRM power density = 34.5W/mm^2**

Normalized Fmax versus Vdd_core

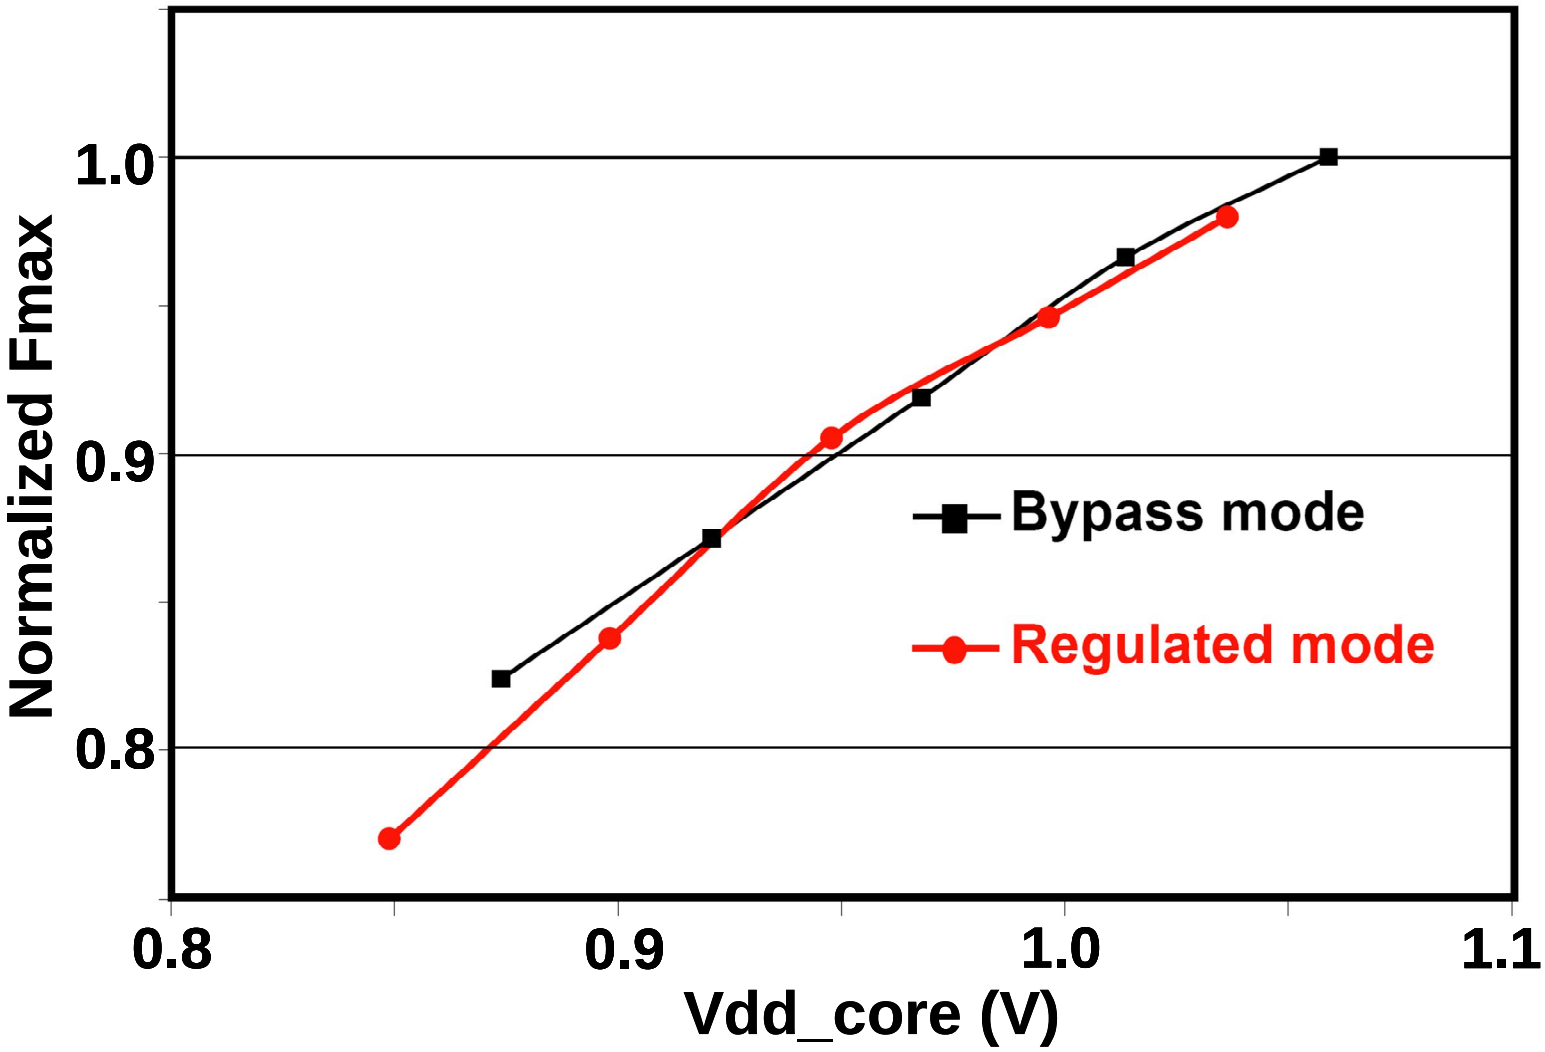

- **Fmax normalized to its value at Vdd_core =1.05V in bypass mode**

DVFS vs. DFS Savings

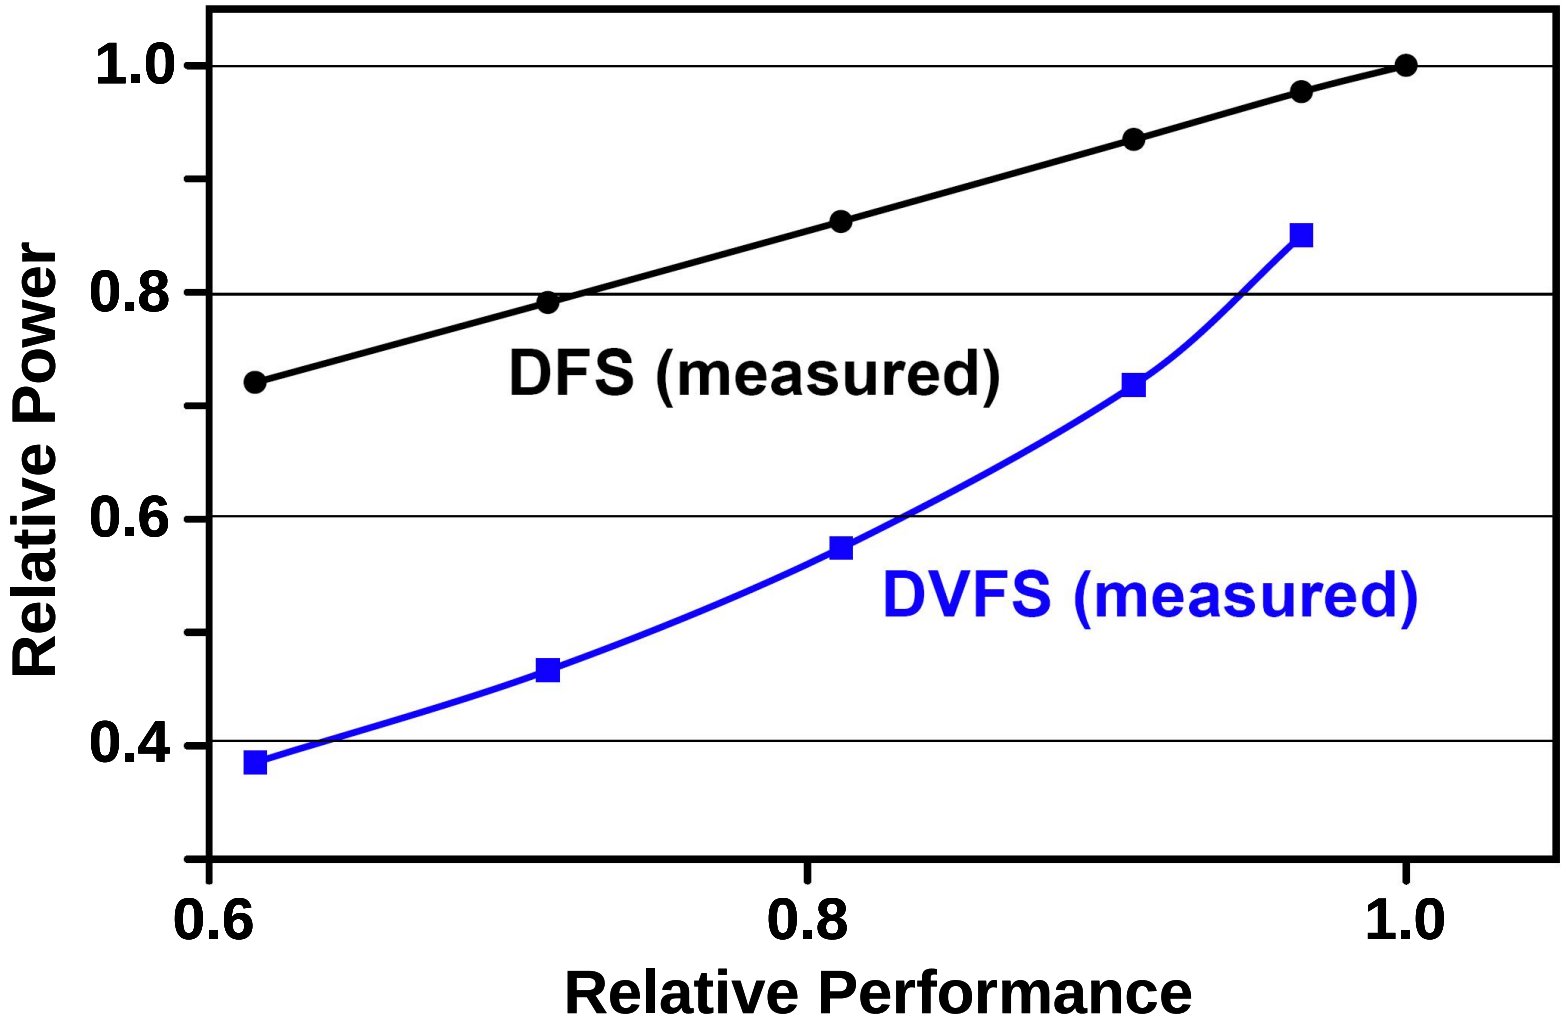

- **At 62% of max. core frequency, power with DVFS is 1.8X lower than with DFS**

Conclusions

- **iVRMs enable per-core DVFS for POWER8™ processor**
  - **First reported integrated voltage regulator system with no off-chip components for server-class processor**
- **Scalable distributed dual loop architecture**
  - **UREG building blocks (total of 1764 UREGs per chip)**
  - **Key features for large scale system**
    - **Digital distribution of UPDN codes**
    - **Local RC filtering for slow passgate control**
- **Measured performance**
  - **Peak power efficiency of 90.5%**
  - **Power density of 34.5W/mm^2**
    - **at least 3.5X higher than that of inductor-based or SC converters**

Acknowledgements

- **Members of IBM Research and IBM Systems and Technology Group for essential design, test, and managerial support**

  - **Michael Sperling**
  - **John Bulzacchelli**
  - **Gregory Still**
  - **Ryan Kruse**
  - **Seongwon Kim**
  - **David Boerstler**
  - **Tilman Gloekler**
  - **Raphael Robertazzi**
  - **Kevin Stawiasz**
  - **Timothy Diemoz**
  - **George English**
  - **David Hui**
  - **Paul Muench**
  - **Joshua Friedrich**