Foxton Technology

HotChips 2005

Sam Naffziger
Intel Corp.

The Power Problem

- Power consumption is a primary limiter in today's processors and unfortunately, it varies a lot
  - Part to part (processing)
  - As a result of the application
  - Due to temperature

Measured data of multiple Montecito parts: Power vs. part speed

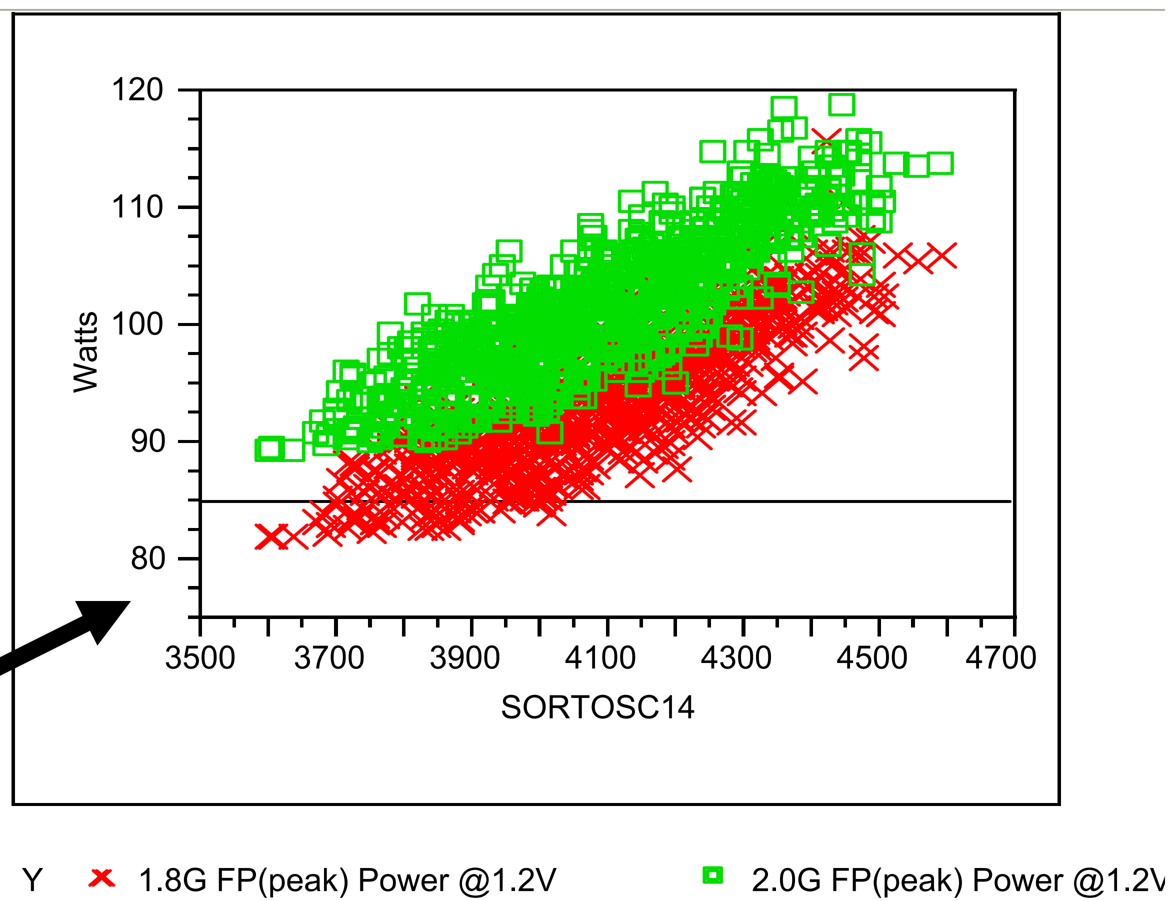

The Power Problem

- Power consumption is a primary limiter in today's processors and unfortunately, it varies a lot
  - Part to part (processing)
  - As a result of the application
  - Due to temperature

Measure Montecito core power vs. application

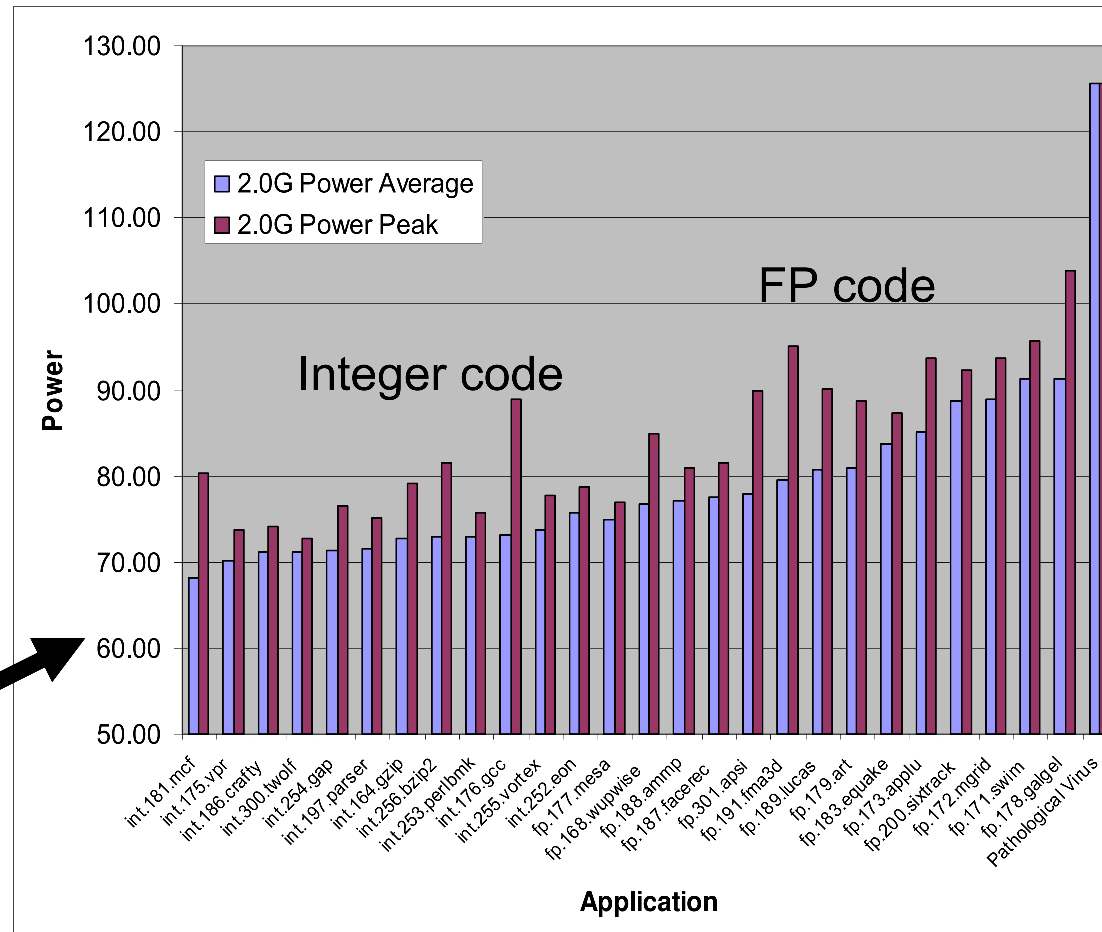

Current Approaches to Power Management and Reduction

- Split out the thermal power spec from the max electrical power spec
  - Use “Thermal Design Power” (TDP) to spec a sustained power that is lower than the true maximum (“electrical power”)
  - Counts on the rarity of very high power events
  - Relies on a thermal sensor to throttle the part if it’s too hot
  - Allows a lower cost thermal solution, but power supplies and power delivery must still handle the max electrical power
- Dynamic Voltage Scaling (C states/P states)
  - Conserve energy when the processor is under-utilized to reduce *average* power
- Fuse in a Vcc that is part-specific
  - Higher power but faster parts can use a lower voltage at the same frequency

The Ideal Power Management for Servers and Desktop

- We currently over-design our power supplies and thermal solutions for worst case parts and applications
- Most of the time the part isn't fully using the watts we've allocated for it
  - Lower power applications only run as fast as the highest power ones

➔We want to maximize performance / Watt for all situations

➔We want a processor to adapt operating point dynamically to it's situation

This is what Foxton Technology does

What is Foxton?

- An integrated system that dynamically maximizes performance per watt including
  - Accurate, integrated power measurement
  - Integrated temperature measurement
  - Frequency control to maximize hertz/volt
  - A microcontroller to incorporate instantaneous {power, temperature, voltage, frequency} and optimize the operating point
- The result is processor cores doing their computation at optimal power efficiency

High Level View of System

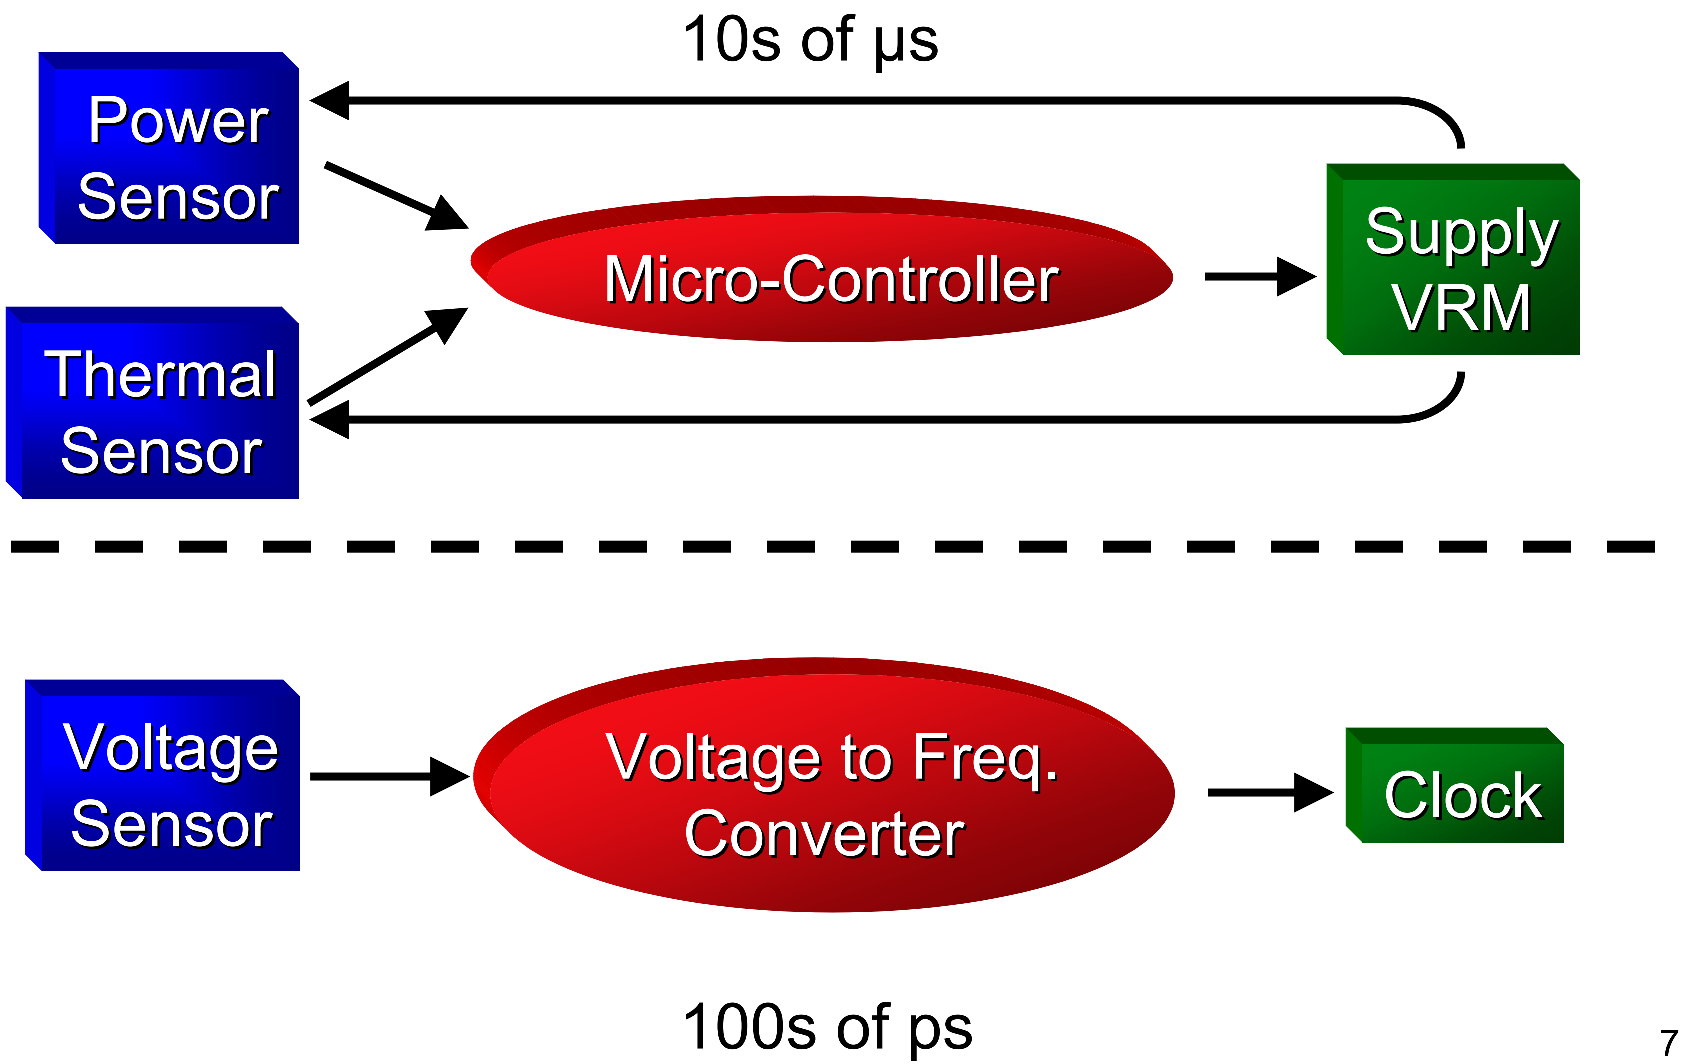

Power Consumption Contour

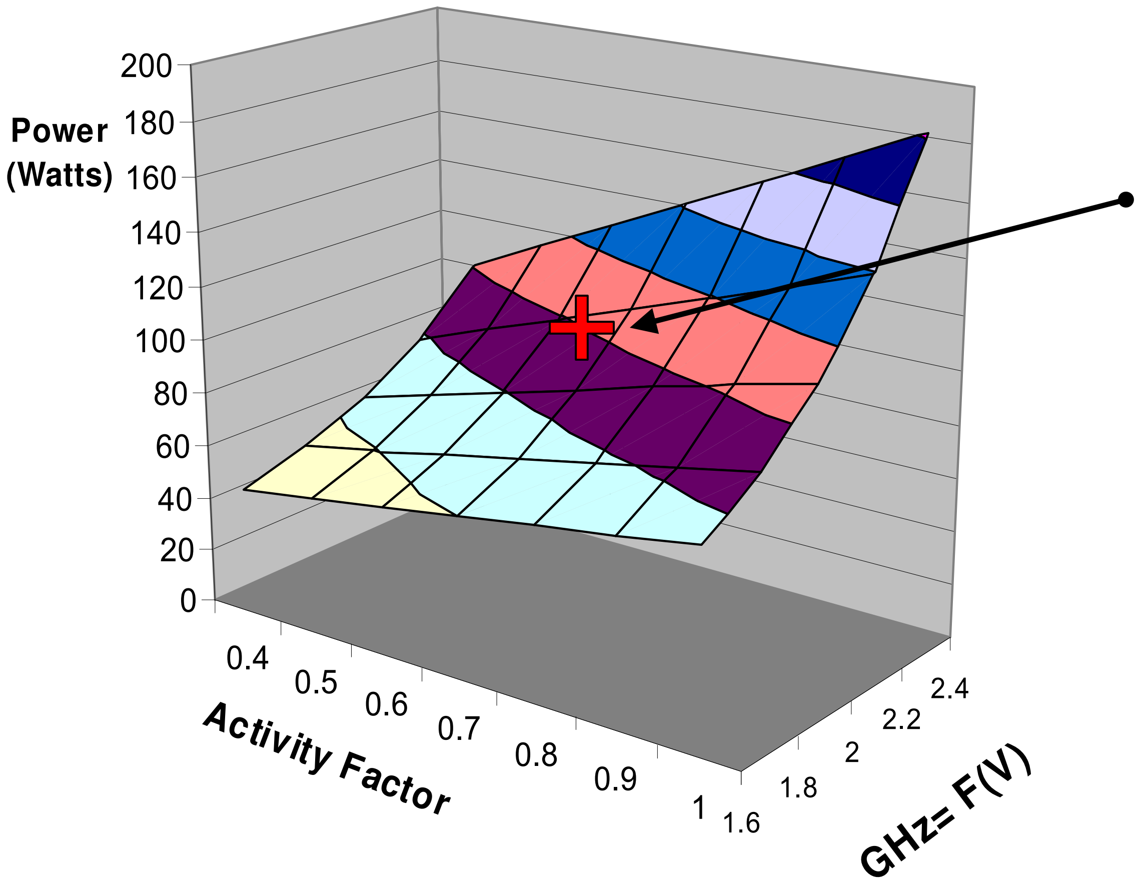

Optimization point is for typical *integer* applications which have .6X the switching power of the worst case
➔ Amdahl's law

Manufacturing test is accomplished by observing the *self measured power*, and the *self-generated frequency* for typical code at the power limit

Measuring Power

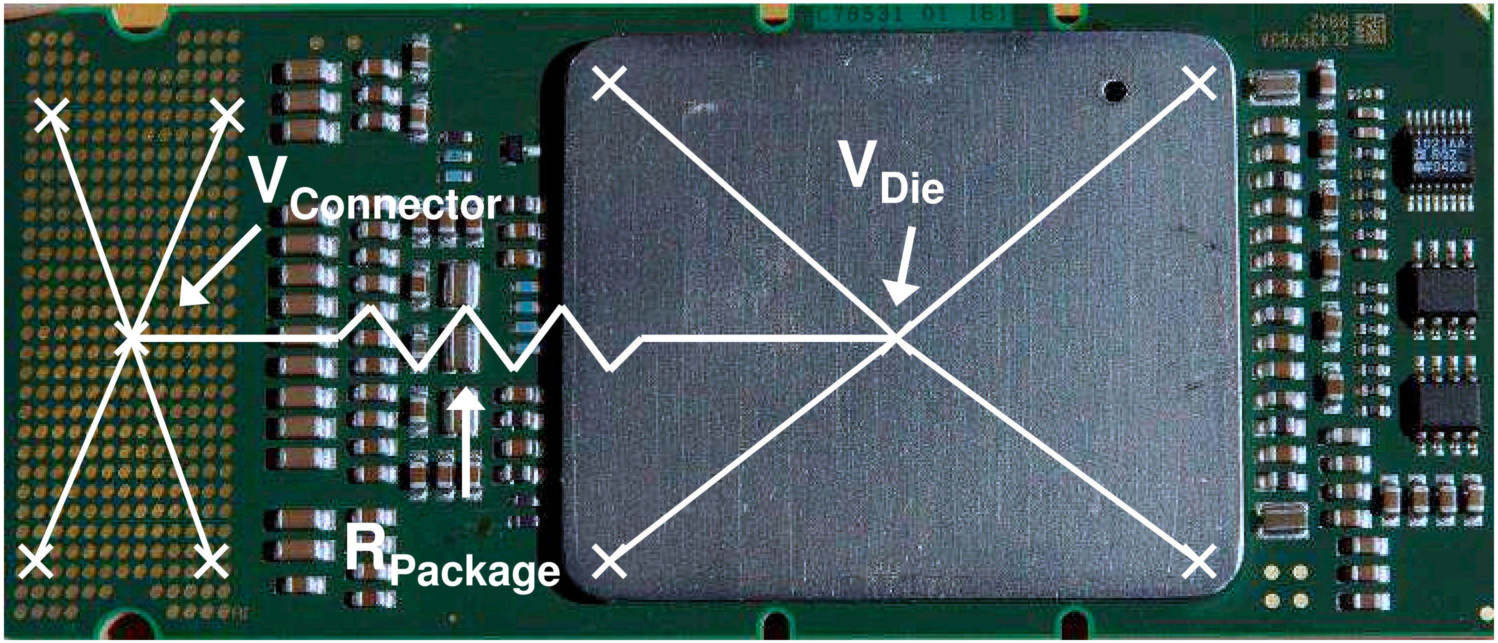

- Use package resistance to measure power
- Avoids burning extra power in measurement
- Portable, self-contained solution
  - No dependence on external power supply

Power Control System

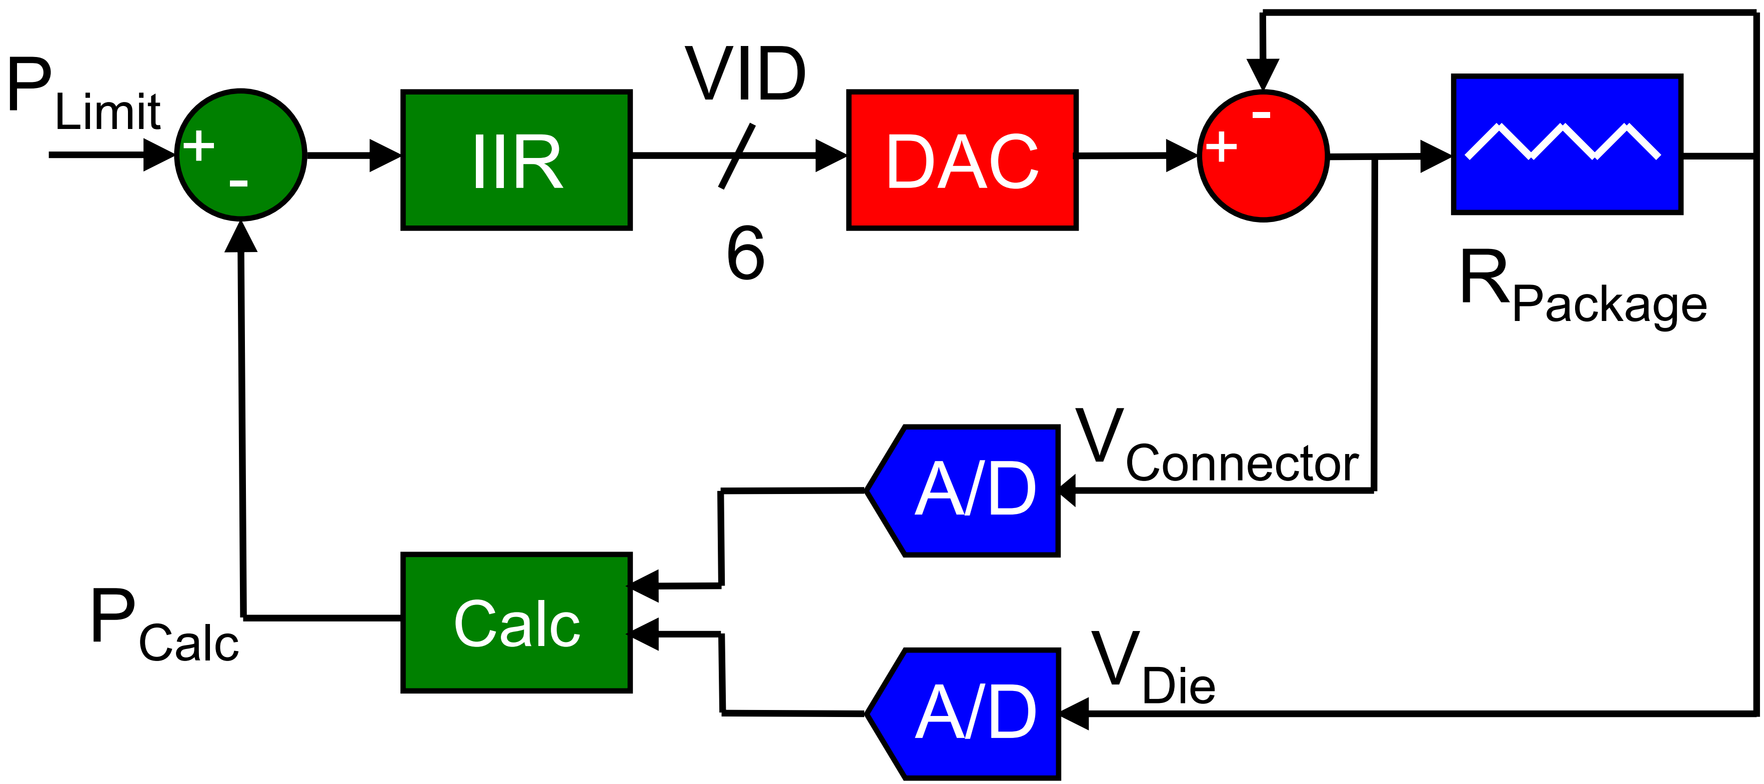

Micro-Controller | Power Supply | Package/Die

Temperature Measurement

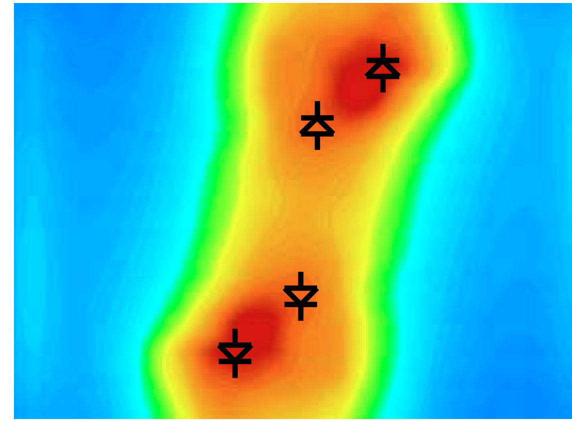

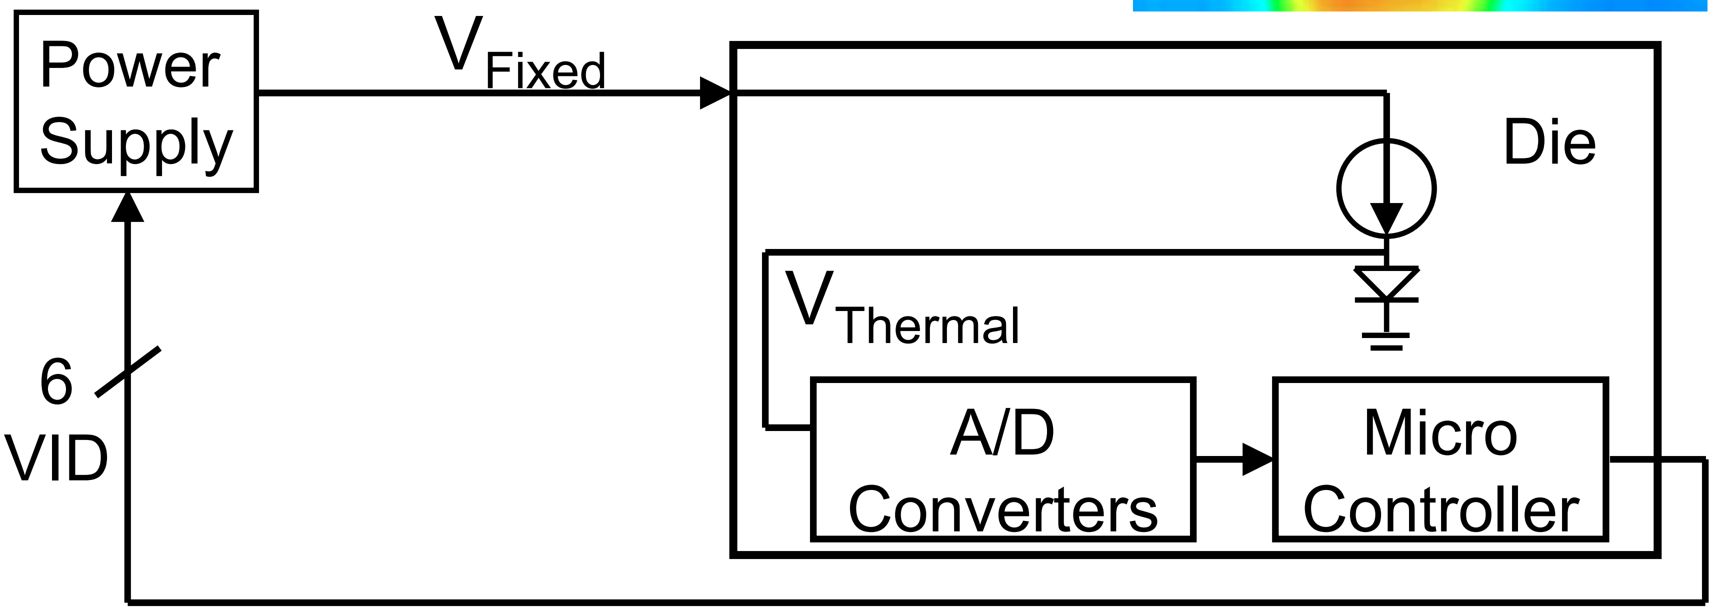

- Calibrate the voltage drop at test to T_J target (90C)
- Use the known -1.7mV per degree C temperature coefficient to calculate die temperature
- Measure the voltage drop across the diodes every 20ms

Package Resistance Calibration

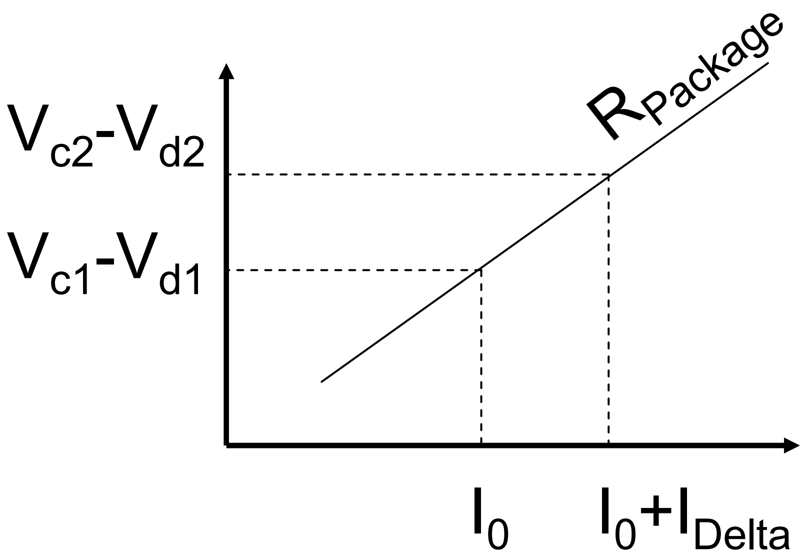

$$R_{Package} = \frac{(V_{c2} - V_{d2}) - (V_{c1} - V_{d1})}{I_{Delta}}$$

- Package resistance can be computed with two voltage measurements with processor stalled
  - Pulling quiescent current I_0
  - Pulling I_0 + a precision, on-die generated current I_{Delta}
- On-package precision R for consistent I_{Delta}
- 66ms recalibration rate

Frequency vs. Power Limit Measured Data

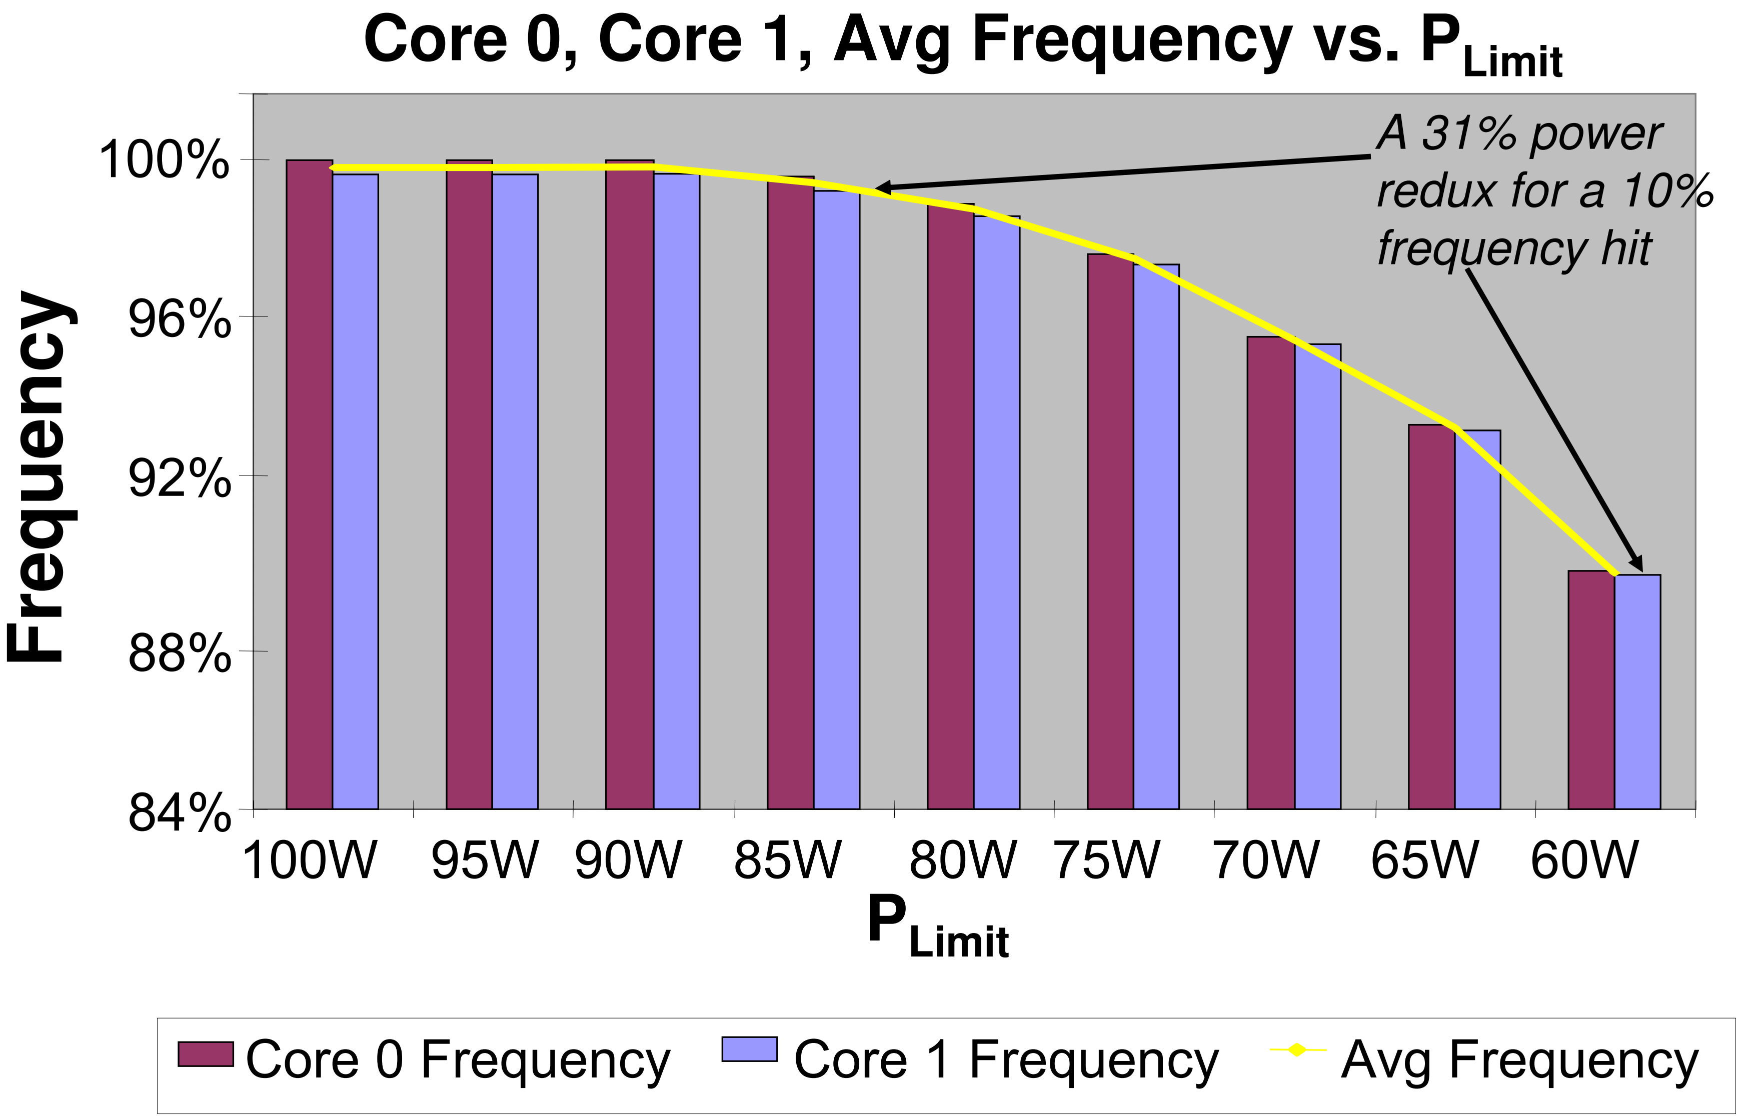

Managing Frequency

- Voltage variability costs frequency and hence performance/watt
- A clock system that can track rapid voltage changes will both maximize hertz/volt and provide smooth response to micro-controller induced voltage changes

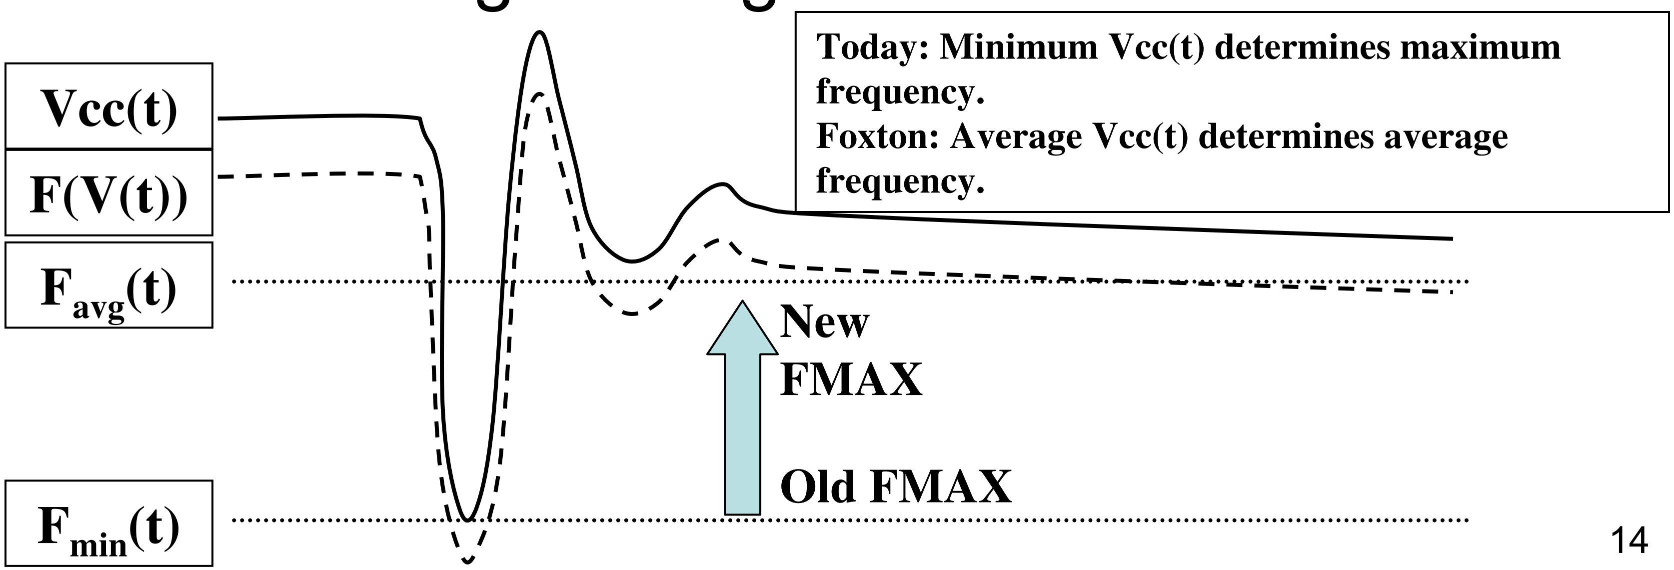

A Variable Frequency Clock System

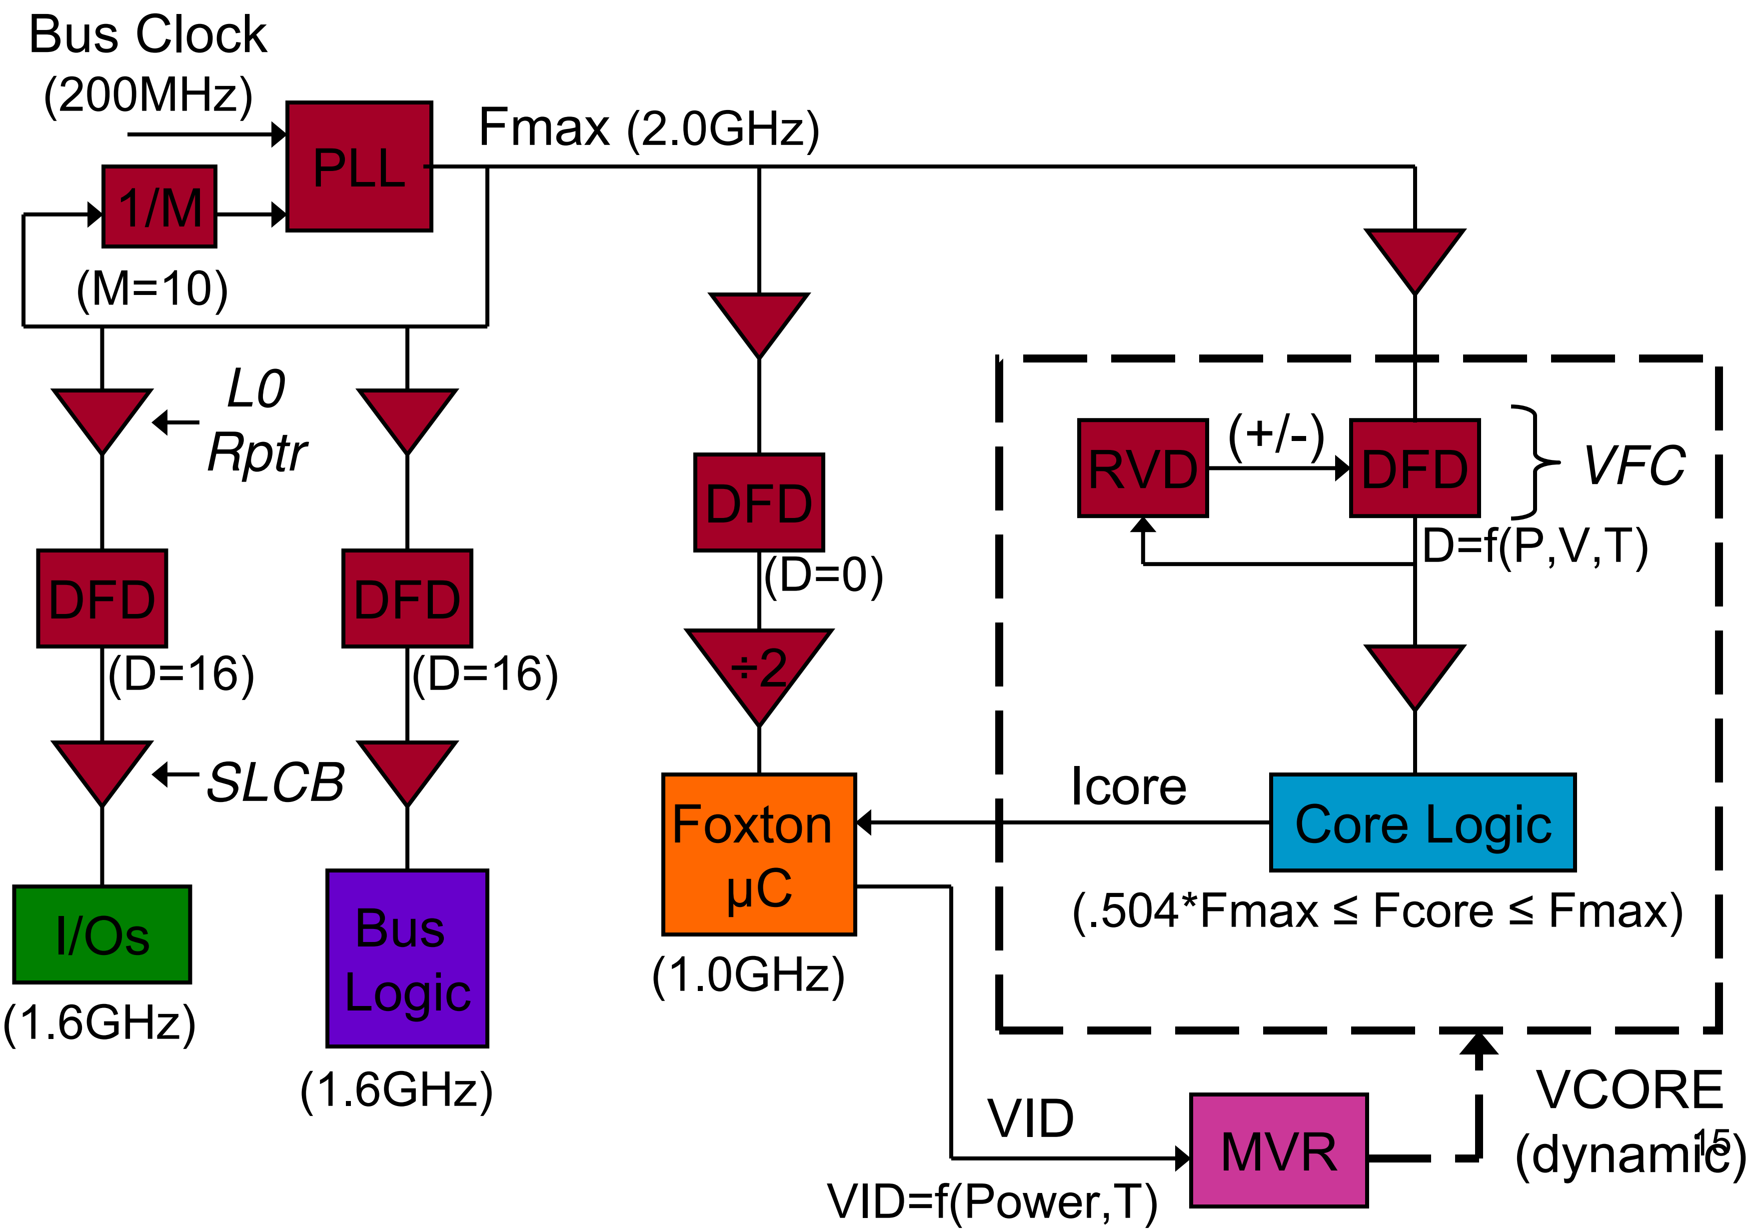

Montecito Clock System Floorplan

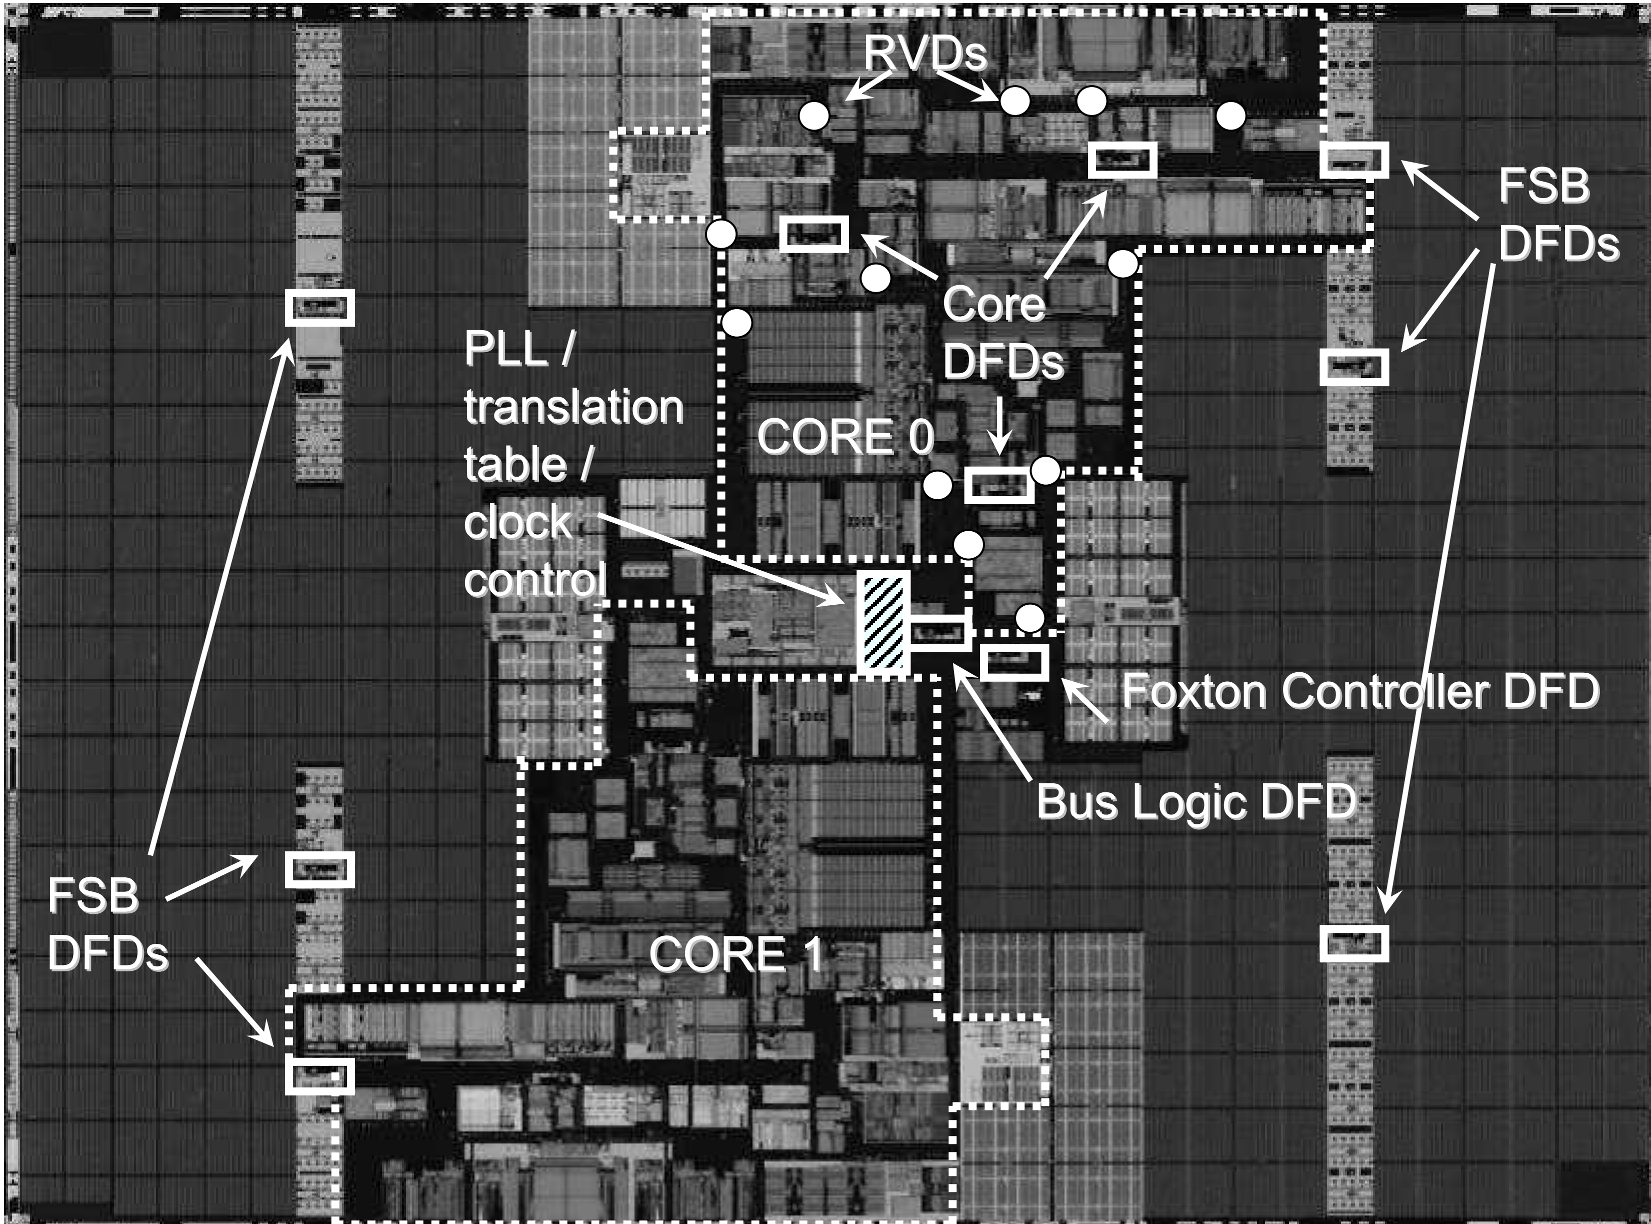

Clock System Modes

- Fixed Frequency (FFM)
  - Cores/Uncore are frequency and phase aligned
  - Cores/Uncore interfaces synchronous
- Variable Frequency (VFM)
  - Core supply modulated by Foxton Controller to manage power envelope
  - Core frequencies track Vcore via Regional Voltage Detector (RVD) V-F curves
    - Respond to Foxton modulation and local transients
    - V-F curves match worst-scaling paths on chip
  - Core/Uncore interfaces asynchronous

RVD Delay Line

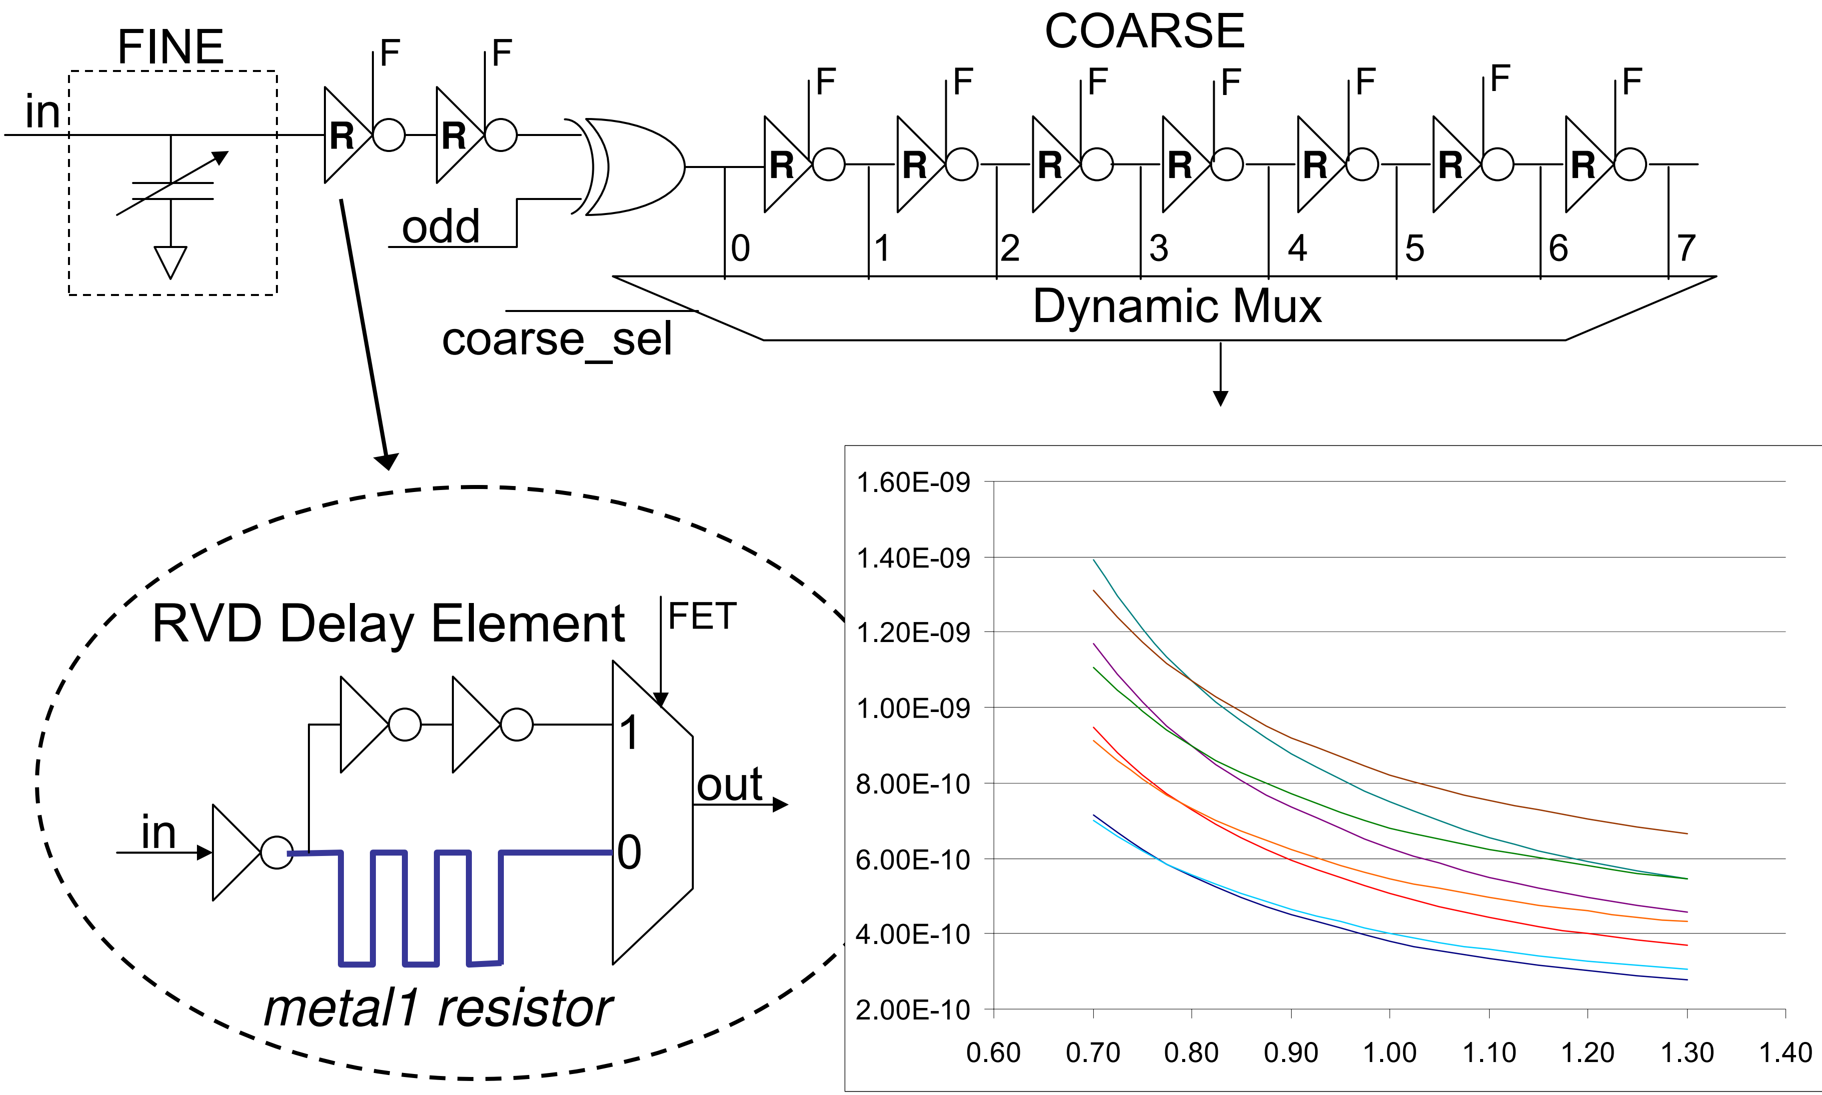

Example VFC Supply Droop Response

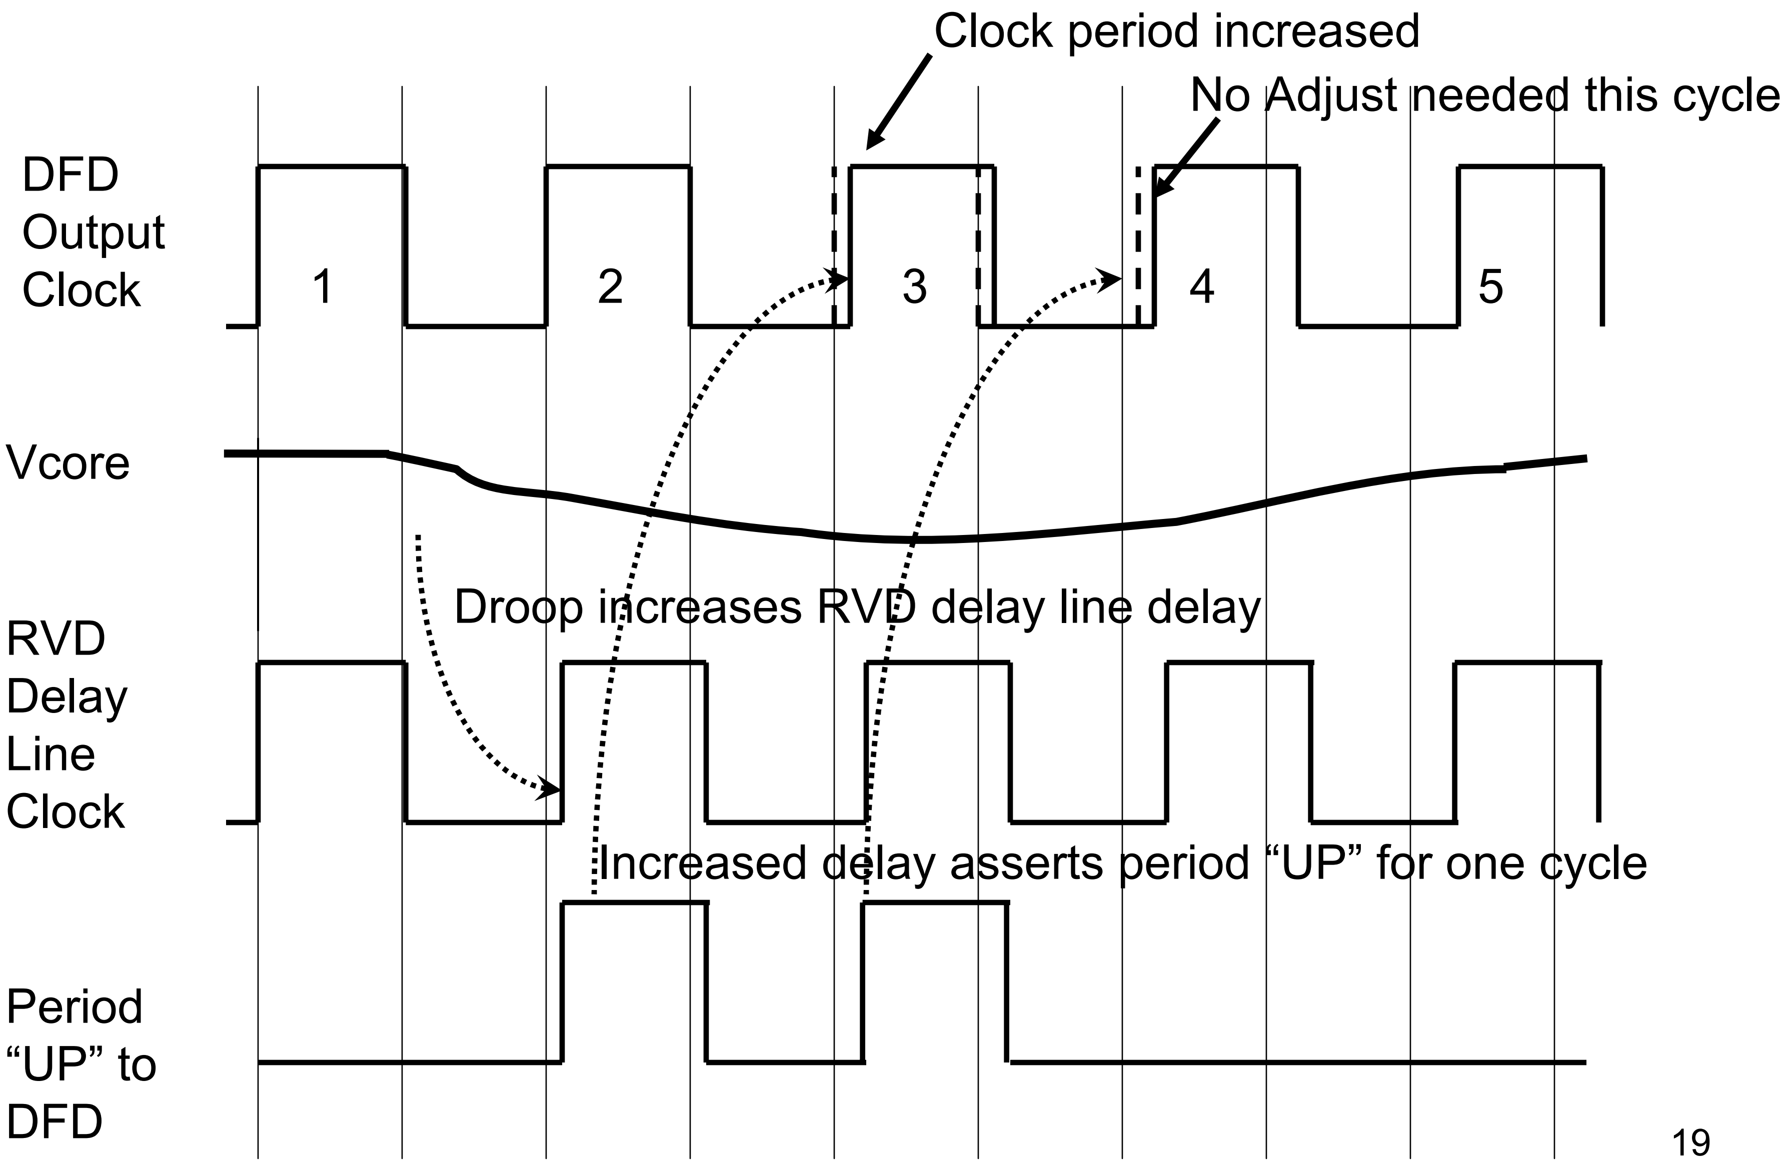

Speed gains from Adaptive Frequency

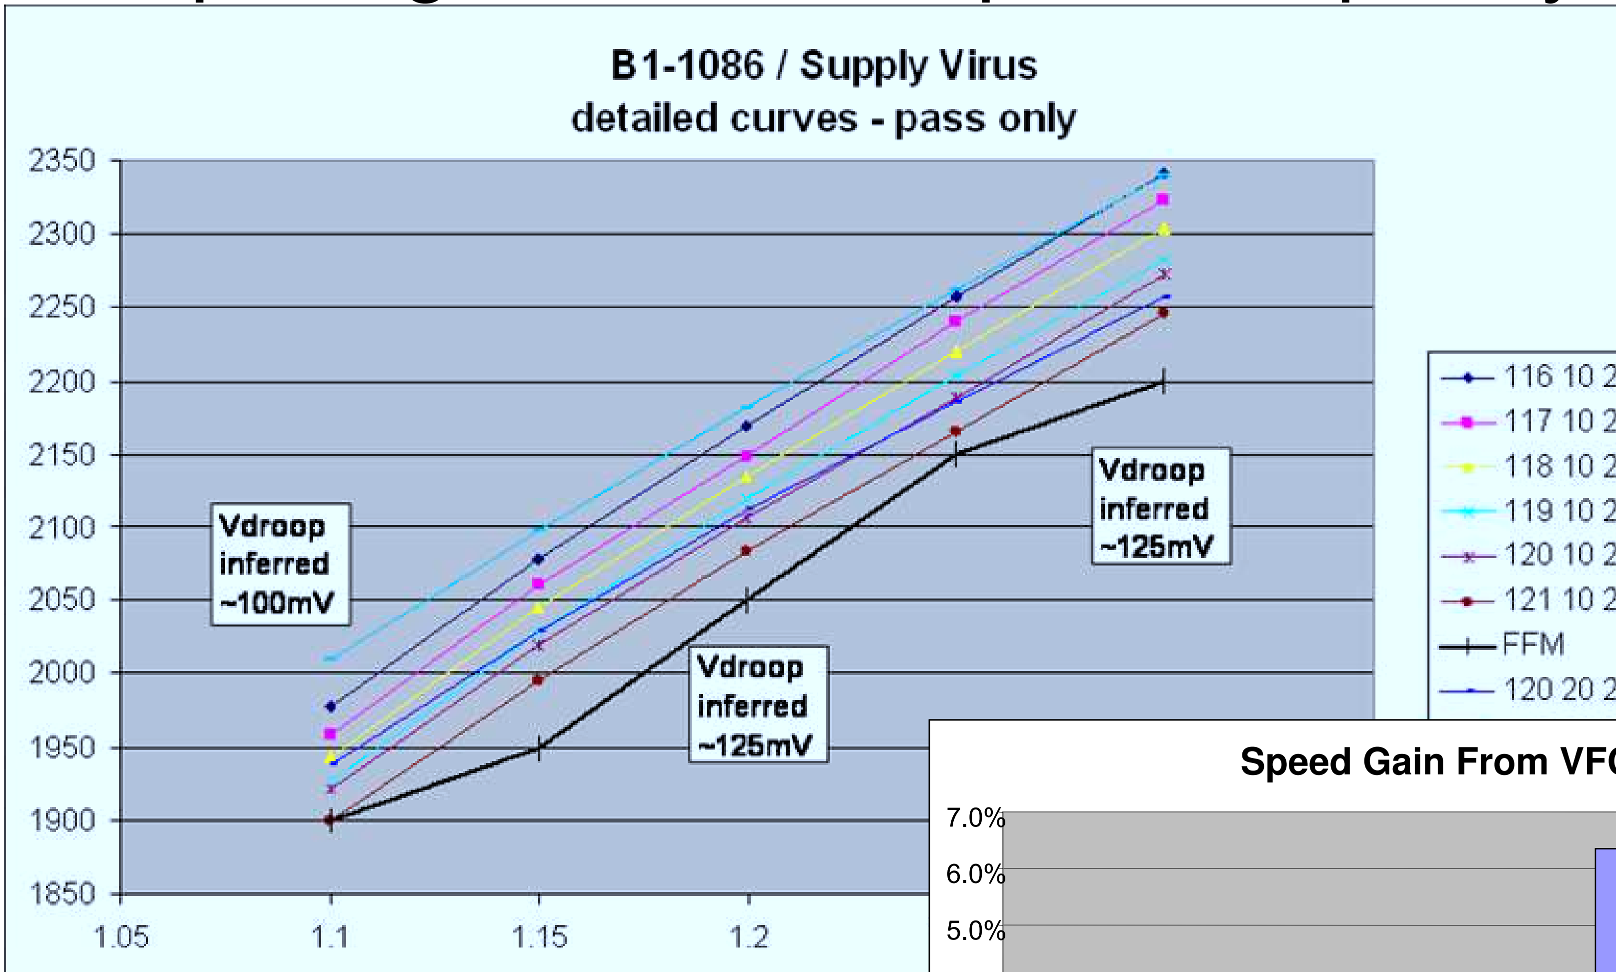

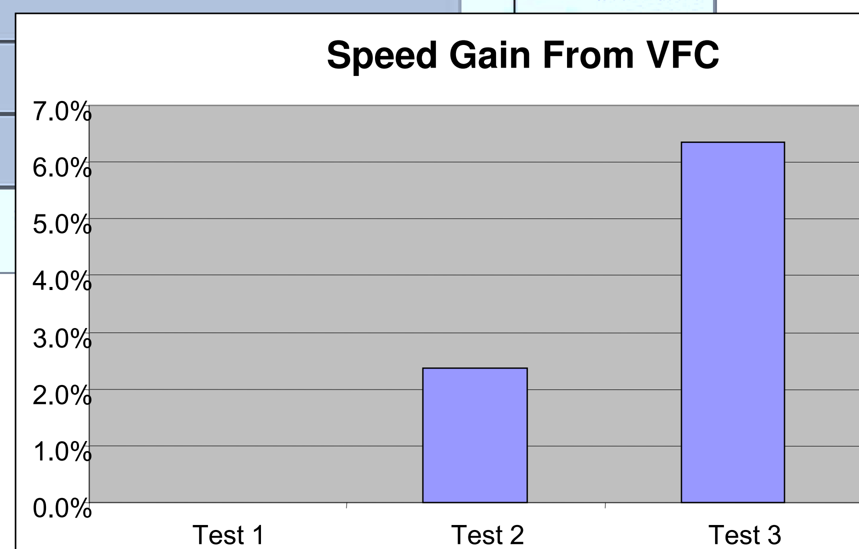

Summary

- Foxton is a system comprised of several key components
  - Accurate power and temperature measurement
  - Fine grained voltage control
  - Dynamic fast-response frequency control
  - A micro-controller to manage the system
- It can be wrapped around any processor or ASIC which can be virtually unchanged except:
  - An asynchronous interface to the rest of the system
  - Must support a wider range of operating voltages
- The result is a self-optimizing chip dynamically delivering greatly improved performance/watt