

Power and Temperature Control on a 90nm Itanium® Family Processor


Jim Ignowski, Richard McGowen,
Chris Poirier, Chris Bostak,
Sam Naffziger

Intel Corporation, Fort Collins, CO

Motivation


- Maximize performance per Watt
- Latest Itanium II (130nm) consumes **130 Watts**
 - **Single** core, **9MB** L3 Cache
- Montecito (90nm) consumes **100 Watts**
 - **Dual** core, **24MB** L3 Cache
- Simple port would be well over **200 Watts**
 - Need something more than low power circuit techniques to address power problem
- Take advantage of $P \propto V^2F$

Take Advantage of $P \propto V^2 F$




High Level View of System

10s of μs



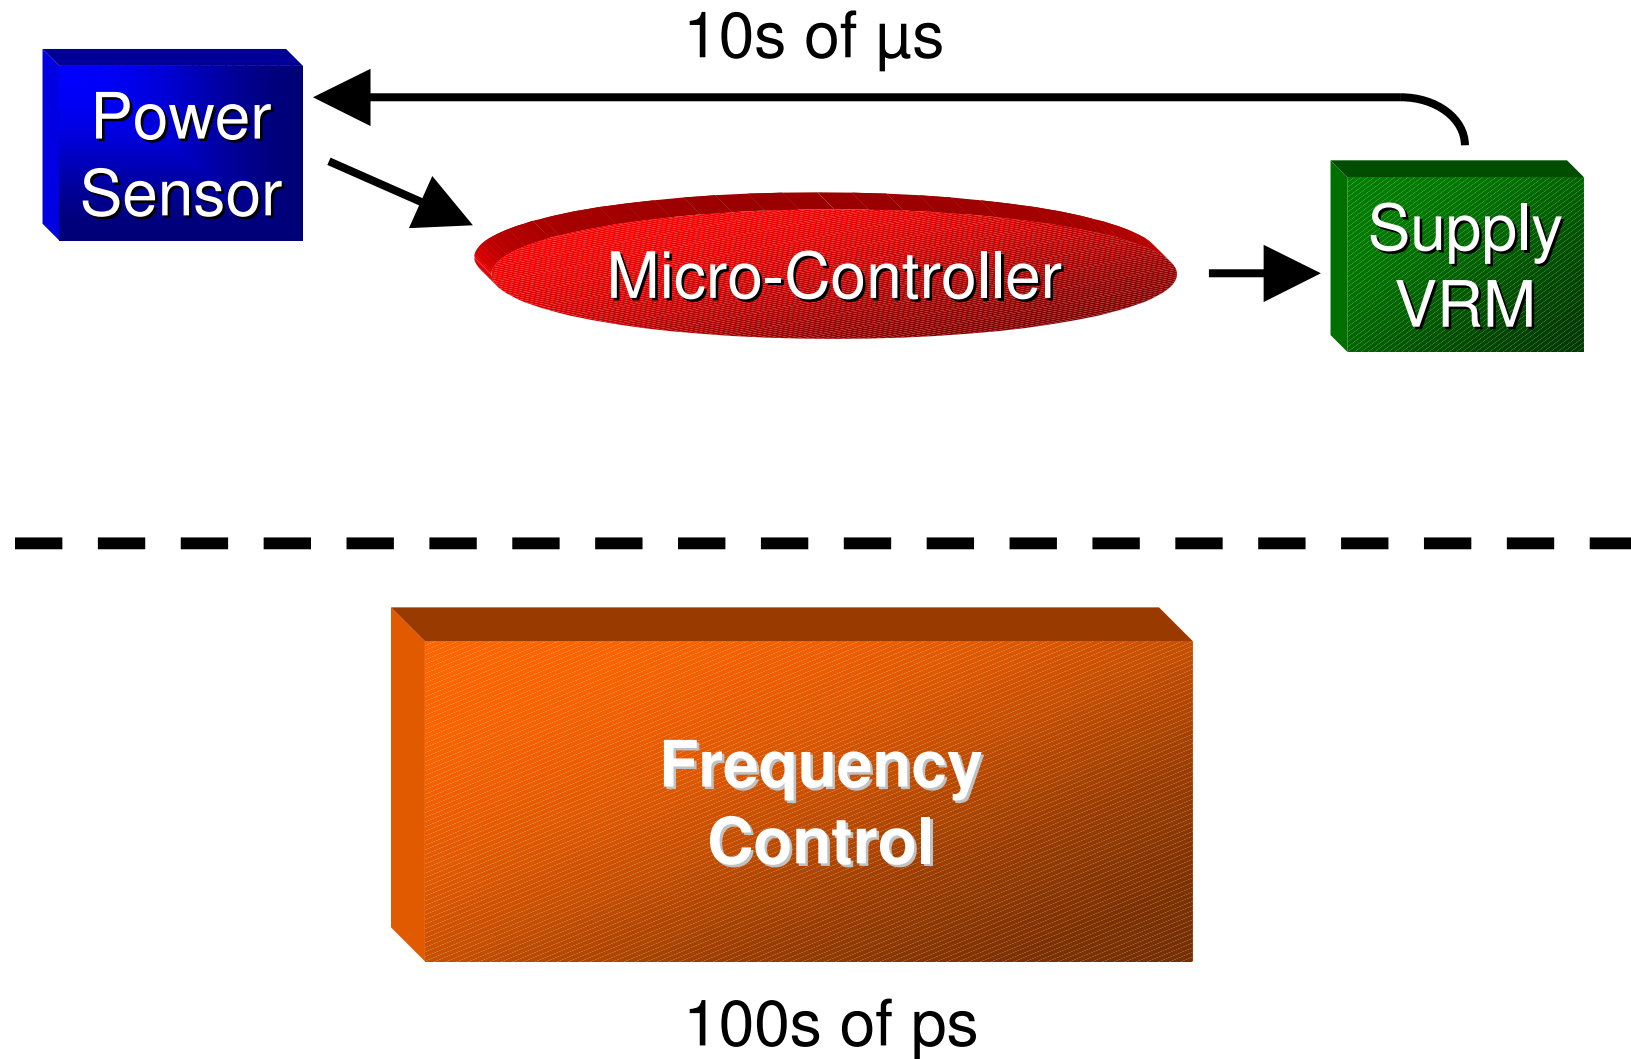
Voltage Control



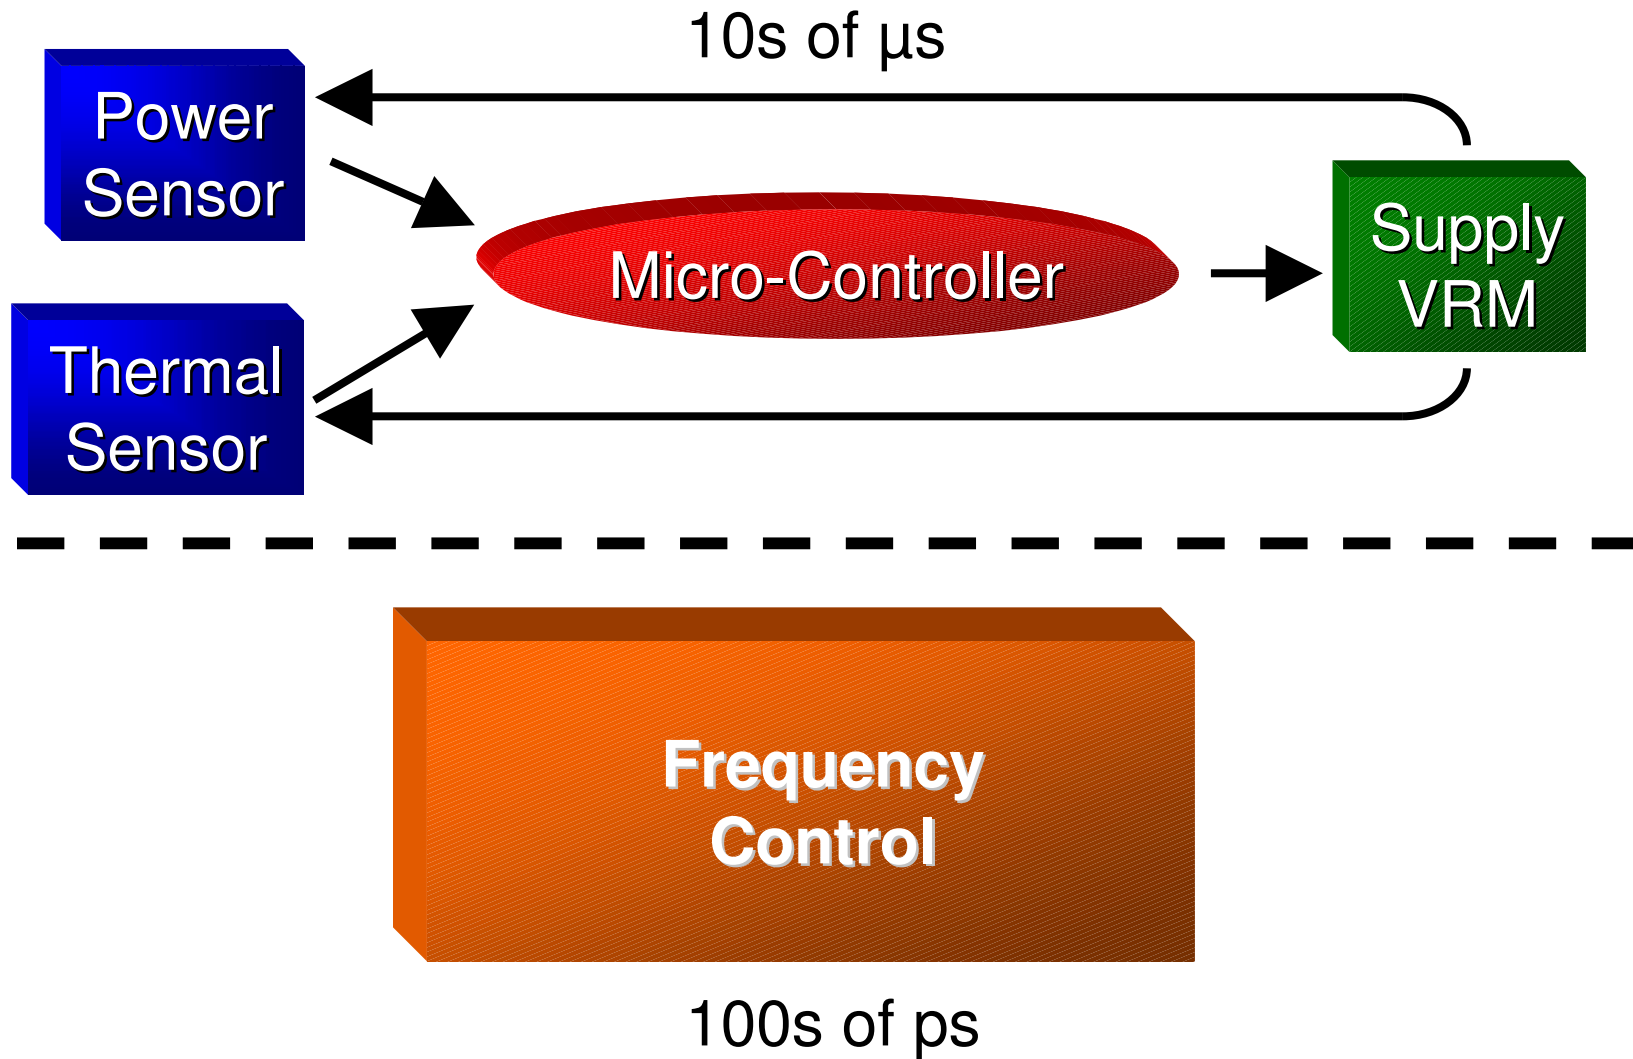
**Frequency
Control**

100s of ps

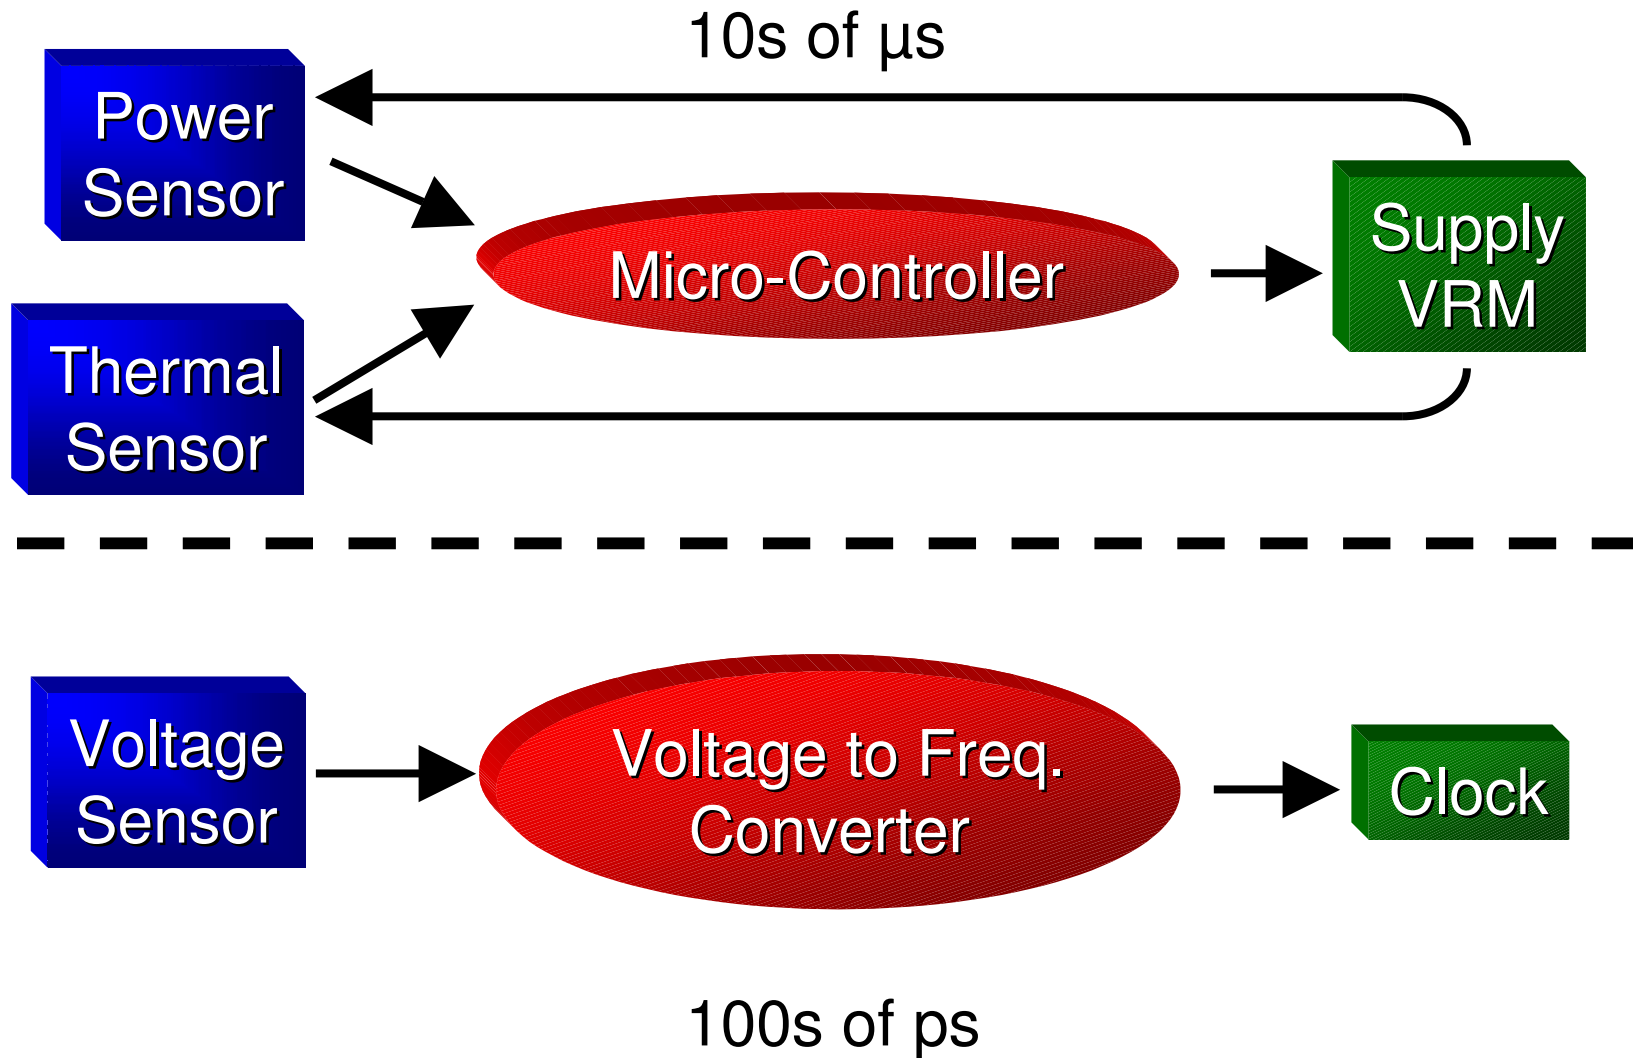
High Level View of System




High Level View of System



High Level View of System

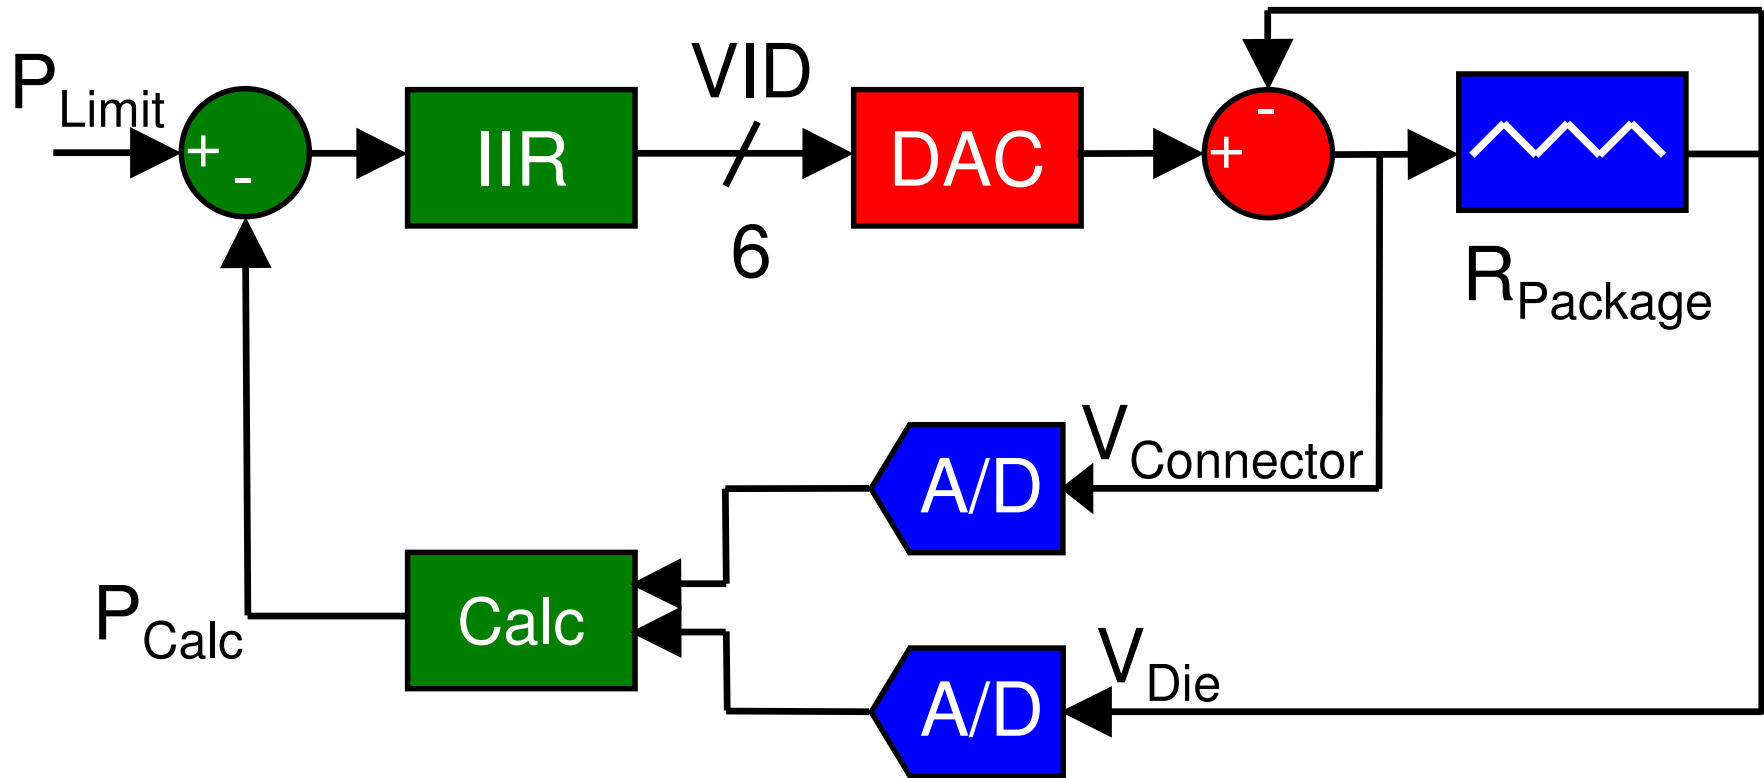


Measuring Power



- 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

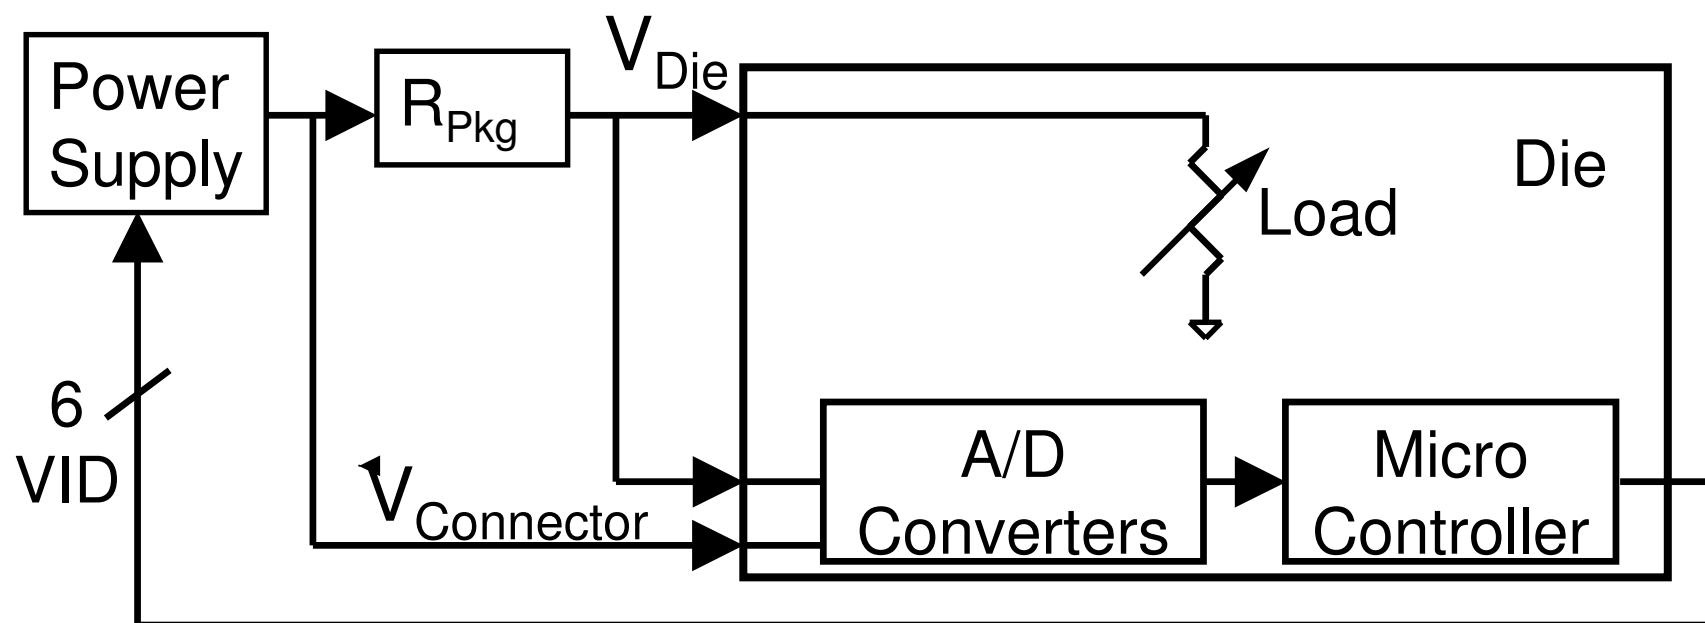


Micro-Controller

Power Supply

Package/Die


System Architecture – Power Measurement



- Measure voltages at both sides of the package to determine current and compute power (8 μ s sample period)

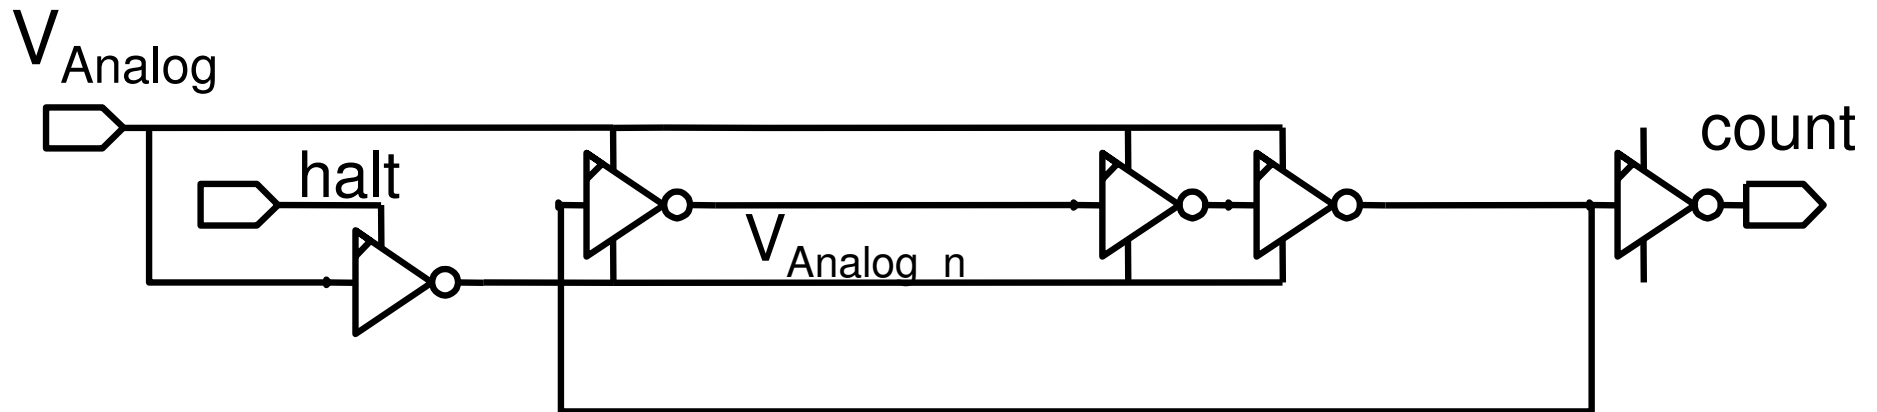
$$Power = V_{Die} I_{Die} = V_{Die} \frac{(V_{Connector} - V_{Die})}{R_{Package}}$$

Voltage Measurement




- By using an analog mux, we can reuse the same VCO to measure voltages from multiple sources
 - On die voltmeter
- High speed counter, $\gg 10\text{GHz}$ for $\sim 100\mu\text{V}$ measurement granularity
- $8\mu\text{s}$ counting interval for filtering and resolution

VCO Detail




- Short loop for high speed
- Analog “inverter”: bias circuit for complementary control signal, also starts/stops oscillation
- Performs averaging function by running at over 10GHz for 8 μ S

VCO Detail– Delay Element



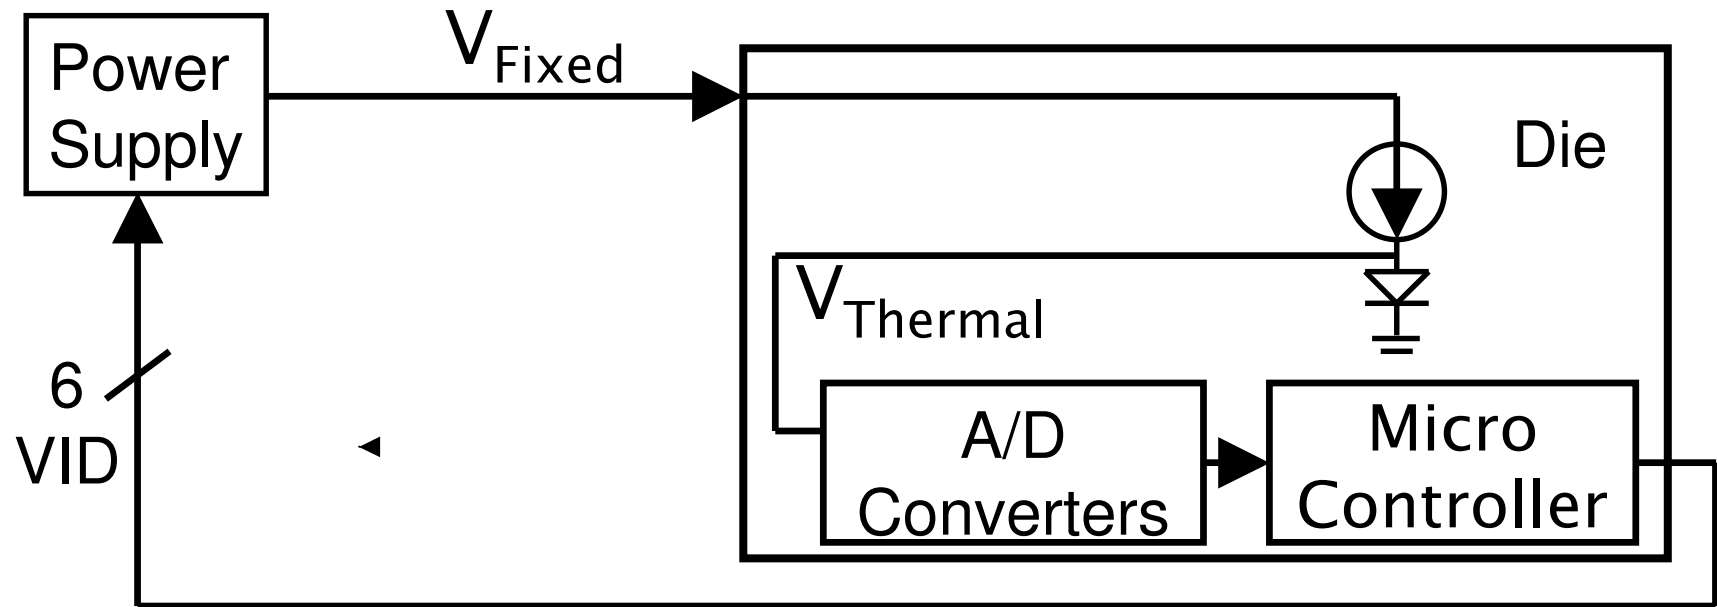
- Modulate output current via series output resistance combined with controlled drive fight
 - Large input voltage range with very linear response

Measuring Temperature



- Two thermal sensors per core
 - Mux thermal diodes into VCOs to measure temperature

System Architecture – 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

Sources of Error and Inaccuracy

- R_{Package} variation with manufacturing process
- R_{Package} variation with temperature
- $<25\text{mV}$ DC voltage drop across package
 - 1% accuracy target requires $250\mu\text{V}$
- Noise on the measurement system supply
 - VCO designed for speed, not power supply rejection
- VCO gain and frequency variation over PVT
- Inherent inaccuracies in the reference components

Requires frequent calibration


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}
- Recalibrate every 66ms (<0.2% perf. hit)


System Overhead



- Burns less than 0.5W, less than $<0.5\%$ area


Frequency vs. Power Limit

Core 0, Core 1, Avg Frequency vs. P_{Limit}



Core 0 Frequency Core 1 Frequency Avg Frequency

Measured Data



Conclusions

- Power and temperature control system using voltage and frequency modulation
 - Maximizes perf./Watt for at least a 25% efficiency gain
 - 90°C on-die temperature control with graceful throttle for improved reliability and performance
 - Programmable power limit
- Voltmeter and package resistance calibration enables
 - 3% average power measurement accuracy
 - 1°C temperature control accuracy
- Embedded on die and package, consuming <0.5W and occupying <0.5% of the die