


Powering the Next Generation of AMD Opteron™ Processors



Presented by
George Schuellein – IR
David Bates - AMD

©2007. Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Opteron, and combinations thereof, are trademarks of Advanced Micro Devices. Other names are for informational purposes only and may be trademarks of their respective owners.

The Data Center Challenge



- Data Centers consumed 1.2% of total US electricity in 2005¹
- This number grows at 4x the rate of server spending²
- HP plans to reduce energy in servers by 20% by 2010¹
 - Target to raise MIPS / W by 10x in 3 years (double every yr)¹

Reference: 1. Business Week, March 2007, 2. IDC WW Server Power & Cooling Expense Report, 9/06

Data Center Value Drivers

- ***MIPS / W (Efficiency)***
 - MIPS / Ft² (density)
 - MIPS / \$
 - Total cost of ownership
 - Reliability / Up-time
 - Serviceability (remote control)
 - Modularity / Scalability
 - Configurability
- Efficiency savings add up quickly
 - A watt saved in power consumption saves at least a watt in cooling .
 - A watt saved in power consumption saves a dollar in data center costs
 - Energy Efficient CPUs and Voltage Regulators increase MIPS / Watt

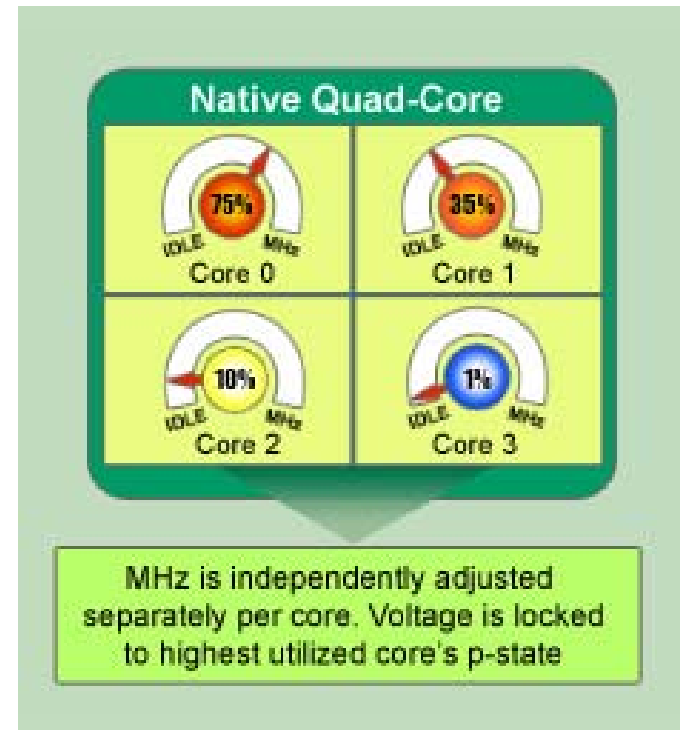
Reference: IBM & HP Websites

Next-generation AMD Opteron™ Processors

- “Barcelona” is the first x86 CPU to integrate four processing cores on a single silicon die
- New **Dual Dynamic Power Management (DDPM)**, provides an independent power supply to the cores and to the memory controller, allowing the cores and memory controller to operate at different voltages, determined by usage
- Significant performance and MIPS-per-watt improvements
- Backwards compatible with existing AMD Opteron™ processor-based platforms

New “Barcelona” Power Management Techniques

- Dynamic adjustment of individual core frequencies
- Processor cores can reduce their voltage level even while the on-chip memory controller runs at full speed.
 - Service external memory requests independent from core p-state transitions
 - Increased p-state opportunities lead to additional power savings
- “Clock gating” enables automatic shut-down of areas of logic not being utilized



AMD Opteron™ Processor Power Architecture

Current generation processors

- Single power plane (VDD) for Core and on-chip Northbridge
- 6 bit PVI (parallel voltage identification) used by the processor to program output voltage
 - 0.375V to 1.55V range
 - 25mV LSB for 1.55V to 0.8V
 - 12.5mV LSB for 0.8V to 0.375V
- Voltage regulator starts upon receiving an Enable signal
- $I_{VDD_{MAX}} = 90.4$ Amps
- DC Tolerance +/-50mV
- AC Tolerance (<5us) +/-100mV

AMD Opteron™ Processor Power Architecture

“Barcelona” processor

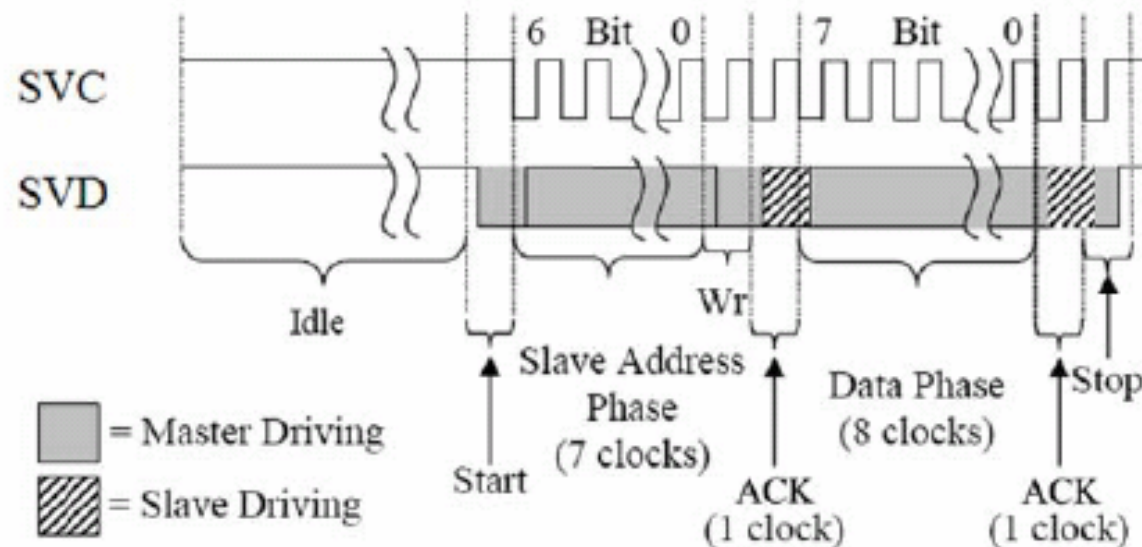
- 2 independent power planes
 - VDD supplies the Core
 - VDDNB supplies the Northbridge
- 7 bit SVI (serial VID interface)
 - 0.5 to 1.55V Range
 - 12.5mV LSB
- Voltage regulator reads a 2 bit parallel boot VID from the SVI inputs upon Enable, then starts
- Voltage regulator responds to SVI commands after receiving a PWROK signal indicating all system rails are within regulation
- $IVDD_{MAX} = 95$ Amps, $IVDDNB_{MAX} = 20$ Amps
- DC Tolerance +/-50mV
- AC Tolerance (<5us) +/-100mV

AMD SVI (Serial VID Interface)

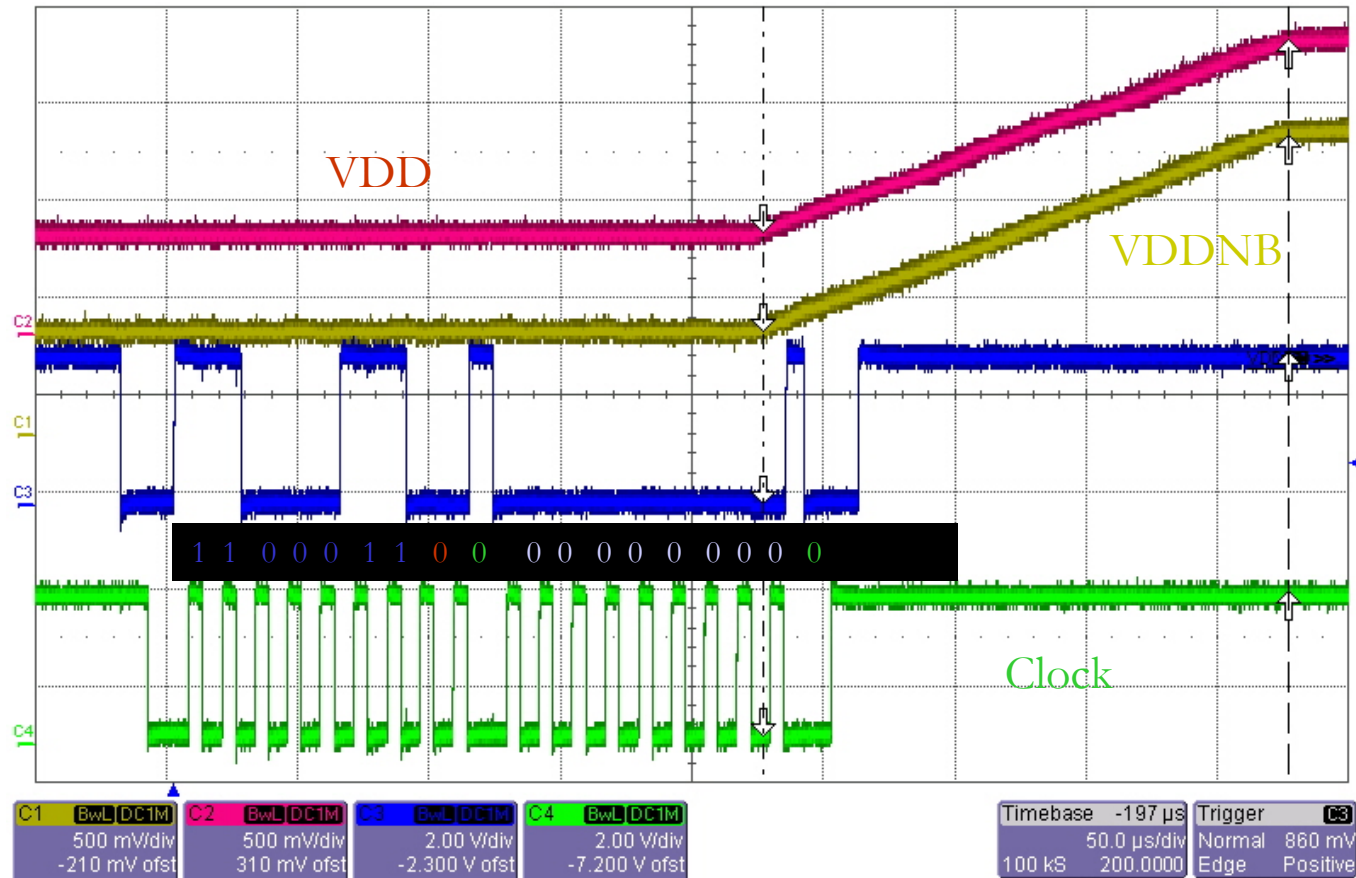
- What is SVI?
 - SVI is a two wire (clock and data) bus that connects a single master (processor) to one or more slaves (voltage regulators)
 - Based on fast-mode I²C/SMBus interface
 - Programs voltage regulator output voltage
 - Voltage regulators do not transmit data to processor
- Why SVI?
 - Replaces the ever-growing parallel VID interface
 - Allows independent VDD power rails
 - Lower pin count/smaller package sizes for voltage regulator suppliers
- Specification owned and developed by AMD
 - Increased flexibility for AMD and its technology partners
 - Allows for quick enhancements for future processors

SMBus Send Byte Protocol

- 8 bit words
- Processor starts send byte by pulling SVD low
- Address for VDD or VDDNB next 7 clocks
- Voltage regulator sends ACK bit
- Processor sends VID byte
- Voltage regulator sends acknowledge bit (ACK) and moves output voltage to the new value



0.5 to 1.55V SVI Commands




Address = 1100011 = Set VID on both Output 1 and Output 2

Write = 0, ACK = 0, Data = 000000 = 1.55V, ACK = 0

PVI/SVI Hybrid Voltage Regulator

- Allows a motherboard to accept either PVI or SVI processors
- Includes six I/Os for PVI
- 2 of 6 I/Os are clock and data inputs in SVI mode
- Reads VID1 upon Enable
 - VID1 = 1 = SVI mode
 - VID1 = 0 = PVI mode
- VDDNB voltage regulator remains off in PVI mode with it's output in a high impedance state

The Power Design Triangle



- *New Technology can improve all 3 goals and minimize the power triangle!*
- With a given Technology;
 - Must make tradeoffs between the 3 goals
 - 2 of 3 goals can be optimized; 3rd suffers
 - Focus on 1 goal; 2nd & 3rd suffer

Goal Oriented Design Approach


▪ Low Cost/Size

- 750kHz Fsw
- “Value” MOSFETs
- VDD Output
 - 4 Phase
 - 120nH Inductors
 - 34 x 22uF MLCC output caps
- VDDNB Output
 - 1 Phase
 - 150nH Inductor
 - 1 x 220uF SP + 10 x 22uF MLCC Output Caps
- Heat sink

▪ High Efficiency




- 300kHz Fsw
- “Performance” MOSFETs
- VDD Output
 - 5 Phase
 - 220nH Inductors
 - 5 x 330uF SP + 10 x 1206 22uF + 20 x 0805 22uF + 8 x 0402 0.1uF MLCC Output Caps
- VDDNB Output
 - 1 Phase
 - 220nH Inductor
 - 1 x 330uF SP + 10 x 1206 22uF + 10 x 0805 22uF MLCC Caps
- No heat sink

PVI/SVI Voltage Regulator Down





Top

Bottom




Load Line (droop) on VDD Reduces Output Capacitors

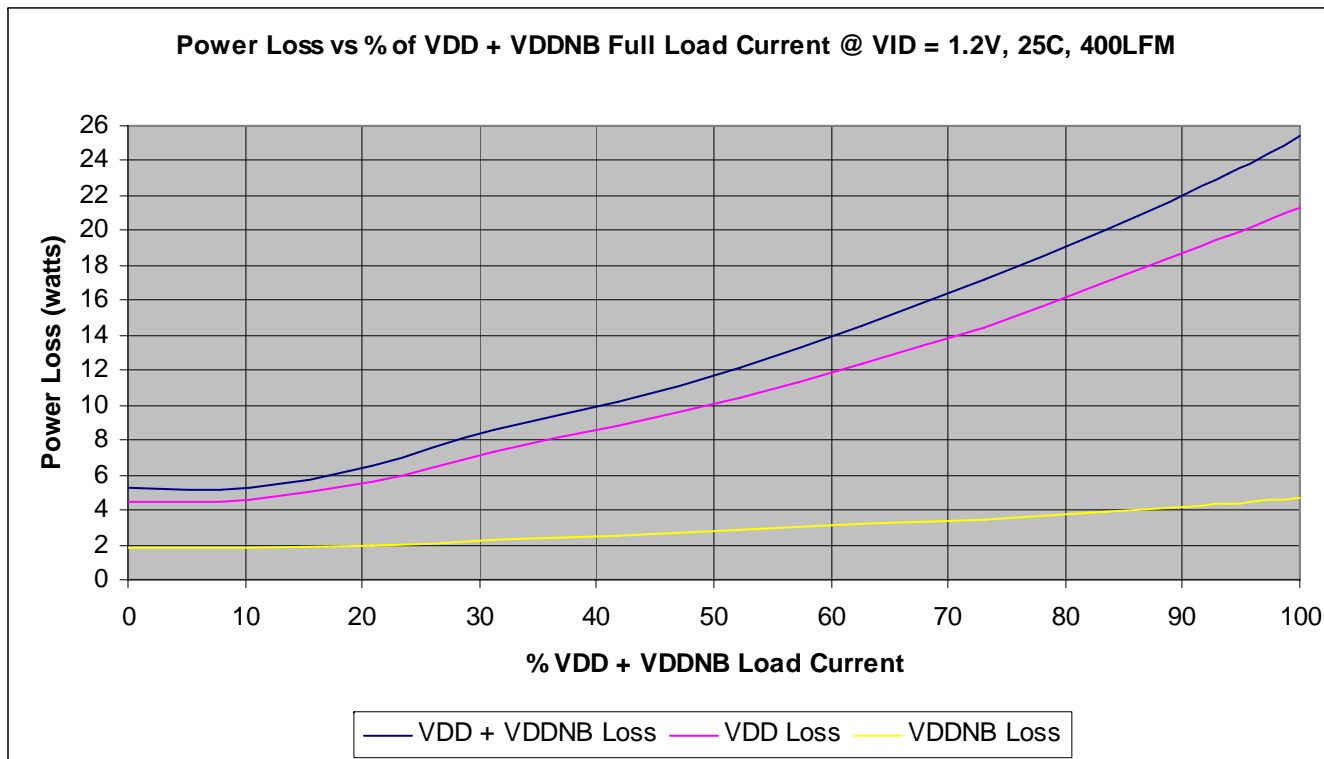


Test Data for Low Cost/Size Design

VDD & VDDNB Efficiency




VDD + VDDNB Power Loss




■ 8.3 Watts / inch²

VDD Thermals (with heat sink)




Bode Plots @ VID = 1.1V, No Load



Data	M1
Frequency	179.38 kHz
Magnitude	-0.007 dB
Phase	64.237 deg



- VDD
BW = 180kHz



Data	M1
Frequency	126.74 kHz
Magnitude	-0.002 dB
Phase	60.680 deg

- VDDNB
BW = 125kHz


VDD 50-100A & VDDNB 8 – 20A Load Transient



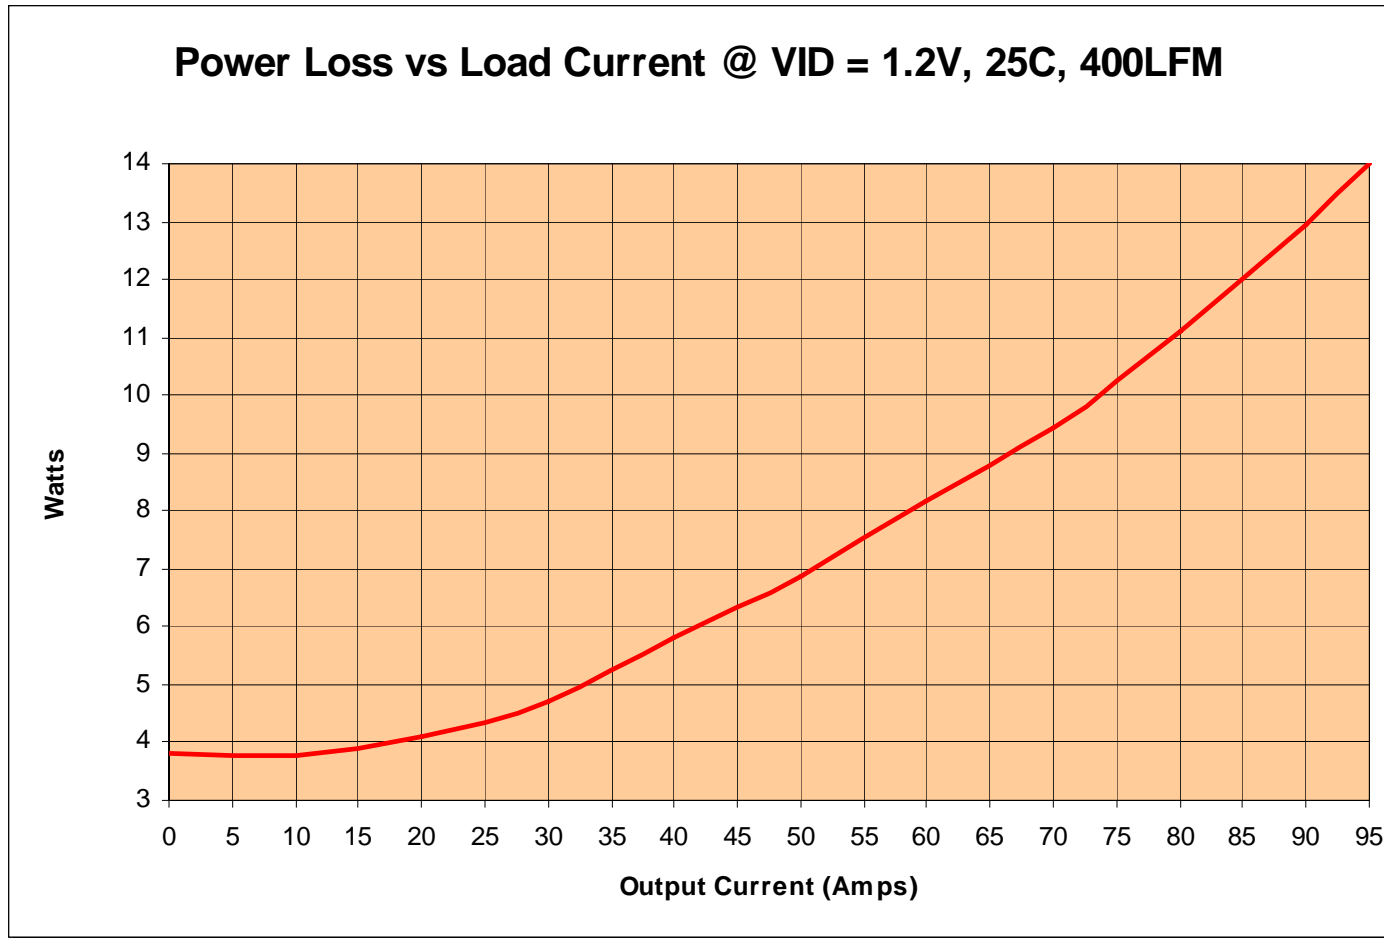
Test Data

High Efficiency Design

VDD Efficiency




VDD Power Loss





- 5.6 Watts / inch²

VDD Thermals (no heat sink)



SVI Voltage Regulator Module


- AMD defined standard VRM
- Dimensions, connector, and output capacitors specified by AMD
- Up to 6 Phase VDD
- 1 Phase VDDNB



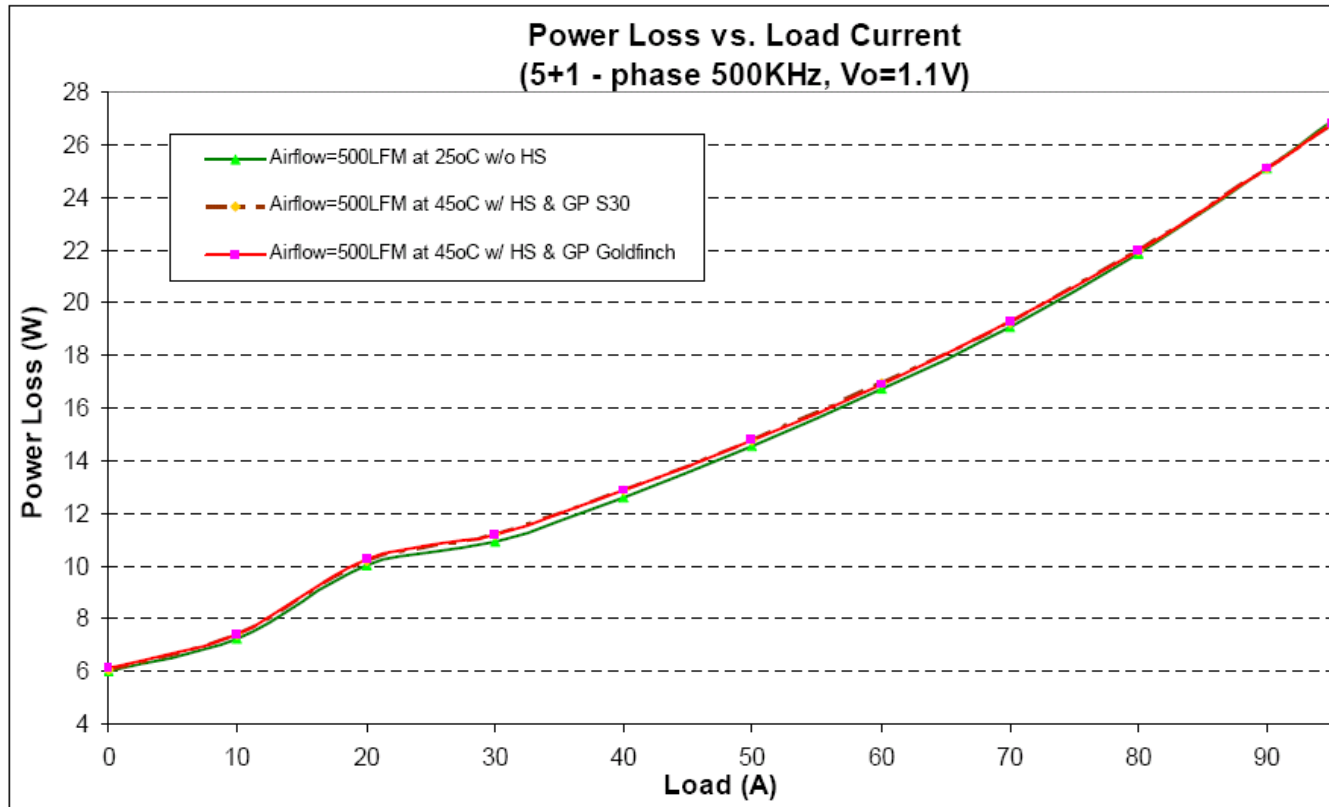
Design Approach

- Spacious form factor allows 500kHz switching frequency, 5 + 1 phases, and “value” MOSFETs without heat sink
- Conservative design – large amount of bulk capacitance results in low control loop BW ($F_{sw}/7$)
- VDD Output
 - 5 Phase
 - 120nH 7 x 10mm ferrite bead inductors
 - 9 x 470uF SP + 10 x 22uF Output Caps
- VDDNB Output
 - 1 Phase
 - 220nH 7 x 10mm ferrite bead inductor
 - 9 x 470uF SP + 10 x 22uF Output Caps

VDD Efficiency




VDD Power Loss



- 2.9Watts / in² – No cooling from motherboard

Thermal Image



VID=1.1V, 25°C, 500LFM, IVDD=95A, IVDDNB=20A, No Heat Sink

Summary

- AMD's "Barcelona" brings new power requirements
 - Separate supplies for CPU Core and memory
 - Serial VID Control of VR Output Voltage
- "Barcelona" compatible components have been developed to provide Computing OEMs with a range of solutions for the "Power Design Triangle"