

A 800 MHz IOP PowerPC SOC with PCI-X DDR266 and DDRII SDRAM 667

By Gerard Boudon, Alan Wall*, Joe Foster**, Barry Wolford*

IBM Microelectronics 91105 Corbeil-Essonnes, France

*IBM Microelectronics Austin Tx USA

** HP Storage division Houston Tx USA

Abstract — A PowerPC system-on-a-chip processor which integrates high speed state of the art with a rich mix of conventional peripherals is described. The PowerPC core and caches achieve frequencies as high as 800 MHz at a supply of 1.5 V and active power consumption as low as 6W. The system with on chip L2 cache executes up to 1600 DMIPS and can be used as IO Processor for RAID Disk array application. The SOC occupies 90 mm² in a 0.13 um, 1.5 V nominal-supply, bulk CMOS process. The high performance memory controller includes several DLL's allowing fine tuning of key Read and Write SDRAM signals.

Index Terms—IOP processor - SDRAM DDRII - PCI-X DDR

I-INTRODUCTION

This PowerPC system-on-a-chip (SOC) design platform is intended to address the high-performance RAID market segment. The SOC uses IBM's Core-Connect technology [1] to integrate a rich set of memory and I/O interfaces including SDRAM DDR2 controller, PCI-X DDR, bit stream XOR, I2O messaging, DMA controller, Ethernet 1Gb, parallel Bus, UART and IIC bus support.

Technology

- CMOS 0.13 um Copper
- 7 levels of Metal
- 11.757 million of gates
- Gate area = 3x12 channels of 0.4um

Packaging

- 29mm FC-PBGA (Flip chip Plastic Ball Grid Array)
- 1mm pitch
- 528 Signal I/Os
- 783 Pads

II- SYSTEM OVERVIEW

This SOC design consists of a high performance 32-bit processor core, which is fully compliant with the PowerPC specification. The processor core for this design was based upon an existing, fixed voltage PowerPC 440 core [2]. The core includes a hardware multiply accumulate unit, static branch prediction support and a 64-entry, fully-associative translation look aside buffer. The CPU pipeline is five stages deep. Single cycle access, two-way set associative, 32-kbyte SRAM instruction and data caches are connected to the processor core.

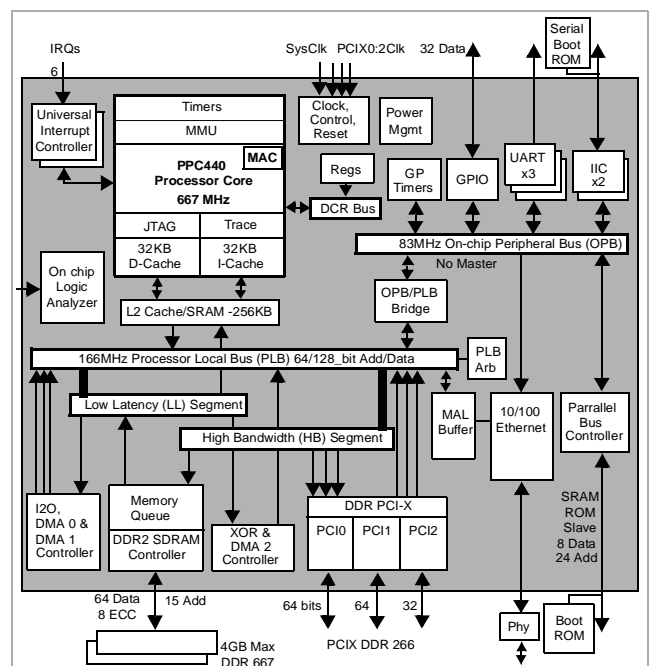


Figure 1 SOC IOP processor block diagram

A second level cache of 256KKB is also integrated improving processor performance by increasing percentage of cache hits. This memory can be used also as an on chip SRAM memory. Included are redundant bit for parity and spares that can be connected after test and configure with on chip fuses

III ARCHITECTURE: CROSSBAR PLB BUS

The key element of this SOC for high speed data transfer rate is the central crossbar PLB (Processor Local Bus) [1]. Two out of Eleven masters can access simultaneously to one of the two PLB slave bus, one specialized in High bandwidth(HB) data transfer and the second one with Low latency (LL). The same physical memory bit in the SDRAM can be accessed either on the HB or the LL slave bus through two aliased addresses. The Crossbar architecture separates read and write data busses allowing simultaneously operations with two independent masters. It also has separated Address and Data busses allowing one master to practically do Reads and Writes simultaneously. Translation of address is done at each crossing of the PLB bus. The 11 masters can access 4 slaves on each slave bus and through the On Chip Peripheral Bus (OPB) 9 more slaves for which performance is less critical.

The Crossbar is made with large Muxes with 11 inputs each time with 64-bit Addresses, and 128-bits Data Write, plus twice 11 Muxes with 64-bit Address, and 128-bits Data Read from the PLB slave busses: This represents very large amount of wiring in the center of the device.

The PowerPC440 CPU is a 32 bit processor that can address up to 4 GB of physical address. The 64 entries TLB transform this address to a real address of 36 bits for a 64GB address space.

These 36 bits are decoded on the 64 bit address PLB bus to access one of the slave. It can be noticed that only some masters such as the PCI-X interface can access to a 64-bit address space.

IV SDRAM DDR 2 at 667MHz

Running DDR SDRAM at 667MHz is a big challenge for the design of the memory controller.

New techniques must be added with several DLL lines and automatic timing adjustment to compensate external wires. The following table shows the differences between the DDR1 and the DDR2

	DDR2 SDRAM	DDR1 SDRAM
Transfer Rate	400/533/667Mbps	200/266/333/400Mbps
Clock Freq	200/266/333MHz	100/133/166/200MHz
Prefetch Size	4-bit	2-bit
Burst Length	4/8	2/4/8
Data Strobe	Differential	Single Ended
Supply Voltage	1.8V	2.5V
I/O Interface	SSTL_18	SSTL_2
Power (400Mbps)	IDD1:247mW (MAX.)	IDD1: 527mW (MAX.)
Package	FBGA	TSOP (II)
DIMMS	240 pins	184 pins
Command Set	Same as DDR	
Basic Timing	Same as DDR	
New Functions	- ODT (On Die Termination) - OCD (Off-Chip Driver) calibration - Posted CAS - AL (Additive Latency)	

The signals between the SDRAM and the Memory controller must follow the rules;

- Commands entered on each rising CLK edge
- DQS edge-aligned with data for READs
- DQS center-aligned with data for WRITEs
- SDRAM DLL to align DQ and DQS transitions with CLK.

Automatic fine phase advance/delay by steps near 12ps of the Clock/Data/DQS and DM signal is used to adjust timings and thus compensate skew due to loads and length of external wires, as drifts due to voltage and Temperature.

The following Read and Write waveforms has been capture on the SDRAM interface with a DDR2 400.

One of the biggest difficulty is to capture data on a Read operation. The figure 5 shows an external signal named feedback which can help to adjust the store in the Memory buffers with the first DQS rising transition.

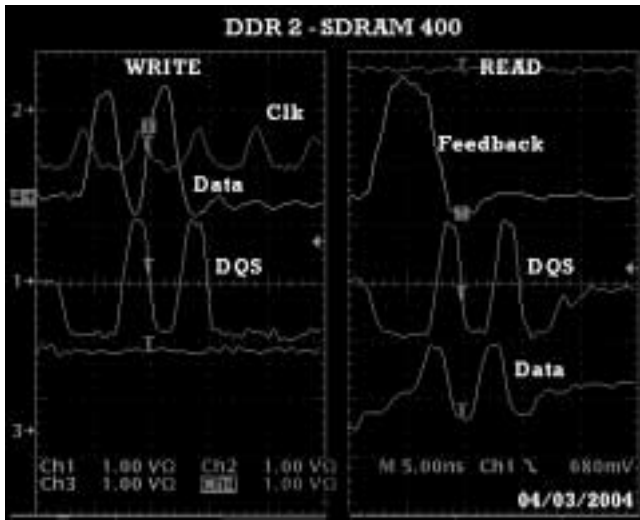


Figure 2 Typical waveform DDRII 400MHz

V - PCI-X DDR 267

Traditional PCI is a multidrop type of bus which limit its performance. The PCI-X DDR mode 2 [5] is an evolution to point to point bus such as the PCI Express while maintaining compliancy with legacy PCI. The PCI-X mode 2 introduces DDR scheme to double the performance to up to 17 Gbits/sec for a 64 bit bus. By comparison, a PCI Express 4 channel as a peak throughput of 10 Gbits/sec.

Among the new features of the PCI-X mode 2 there are DDR, ECC, and OCD control. The circuit of figure 3 that control impedance of the Off chip driver is based on the comparison of a group of NFETs in parallel with an external calibration resistance in a first step. The Adjustment is done by turning on/off several NFETs. In a second step when NFETs are calibrated, the PFET impedance is compared to the NFET which should be equal to the external Resistance. Then the result of On/off PFET/NFET is applied to the final FET's of the Off chip driver.

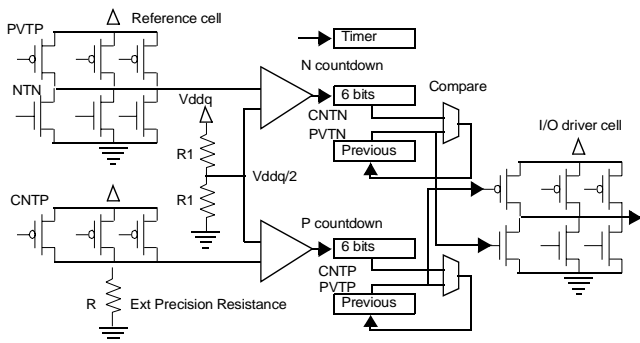


Figure 3: OCD Off-Chip Driver control circuit for PCI-X DDR mode 2

VI - CLOCK DOMAINS

The merge on a single SOC of various cores such as a 800MHz CPU, three PCI-X DDR266 and a SDRAM DDRII 667 leads to the implementation of five (5) PLLs.

Memory and CPU are synchronous and clocking built from the same external low frequency system clock. Due to its high speed operation, two PLL are in cascade. One the CPU-PLL for the CPU and most of the peripherals, and the second the DDR-PLL exclusively used for the Memory controller generating the 2x clock signals. The boundary between these 2 synchronized clocks is done at the PLB bus.

To achieve a minimum skew of 200ps at the PLB bus, the two PLL's are in series with the CPU-PLL signal that input the DDR PLL adjusted in timing with the end of the CPU clock tree feeding CPU latches. The feedback of the DDR-PLL is taken after the clock tree such that zero delay is added from the entry of the main clock. Thus this scheme can support the 2 PLL's in each corner of the die inside their respective cores.

Each of the PCI-X have his own clock and are not synchronized. We have two clock domains in the PCI interface, and special care was done to re synchronized data, between PLB bus and External PCI agents.

VII - POWER DISSIPATION

The PowerPC architecture is well reputed for its low power dissipation coupled with high performance.

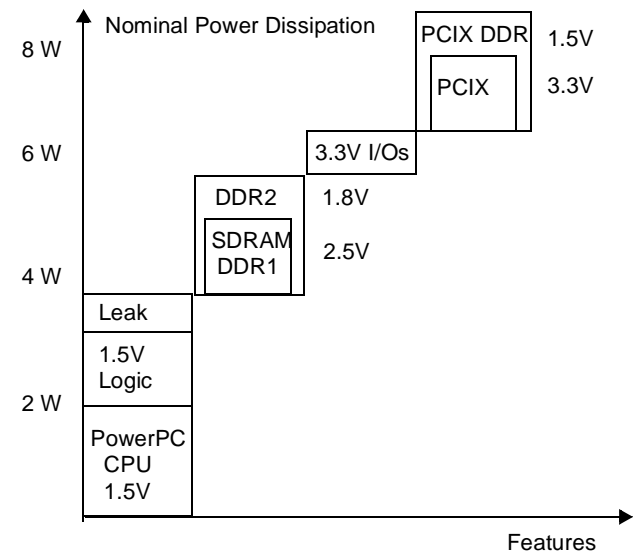


Figure 4 Power dissipation breakdown

The power breakdown of the various cores on the chip highlights the growing importance of the Memory controller and PCI-X busses at high frequency.

VIII-IMPLEMENTATION

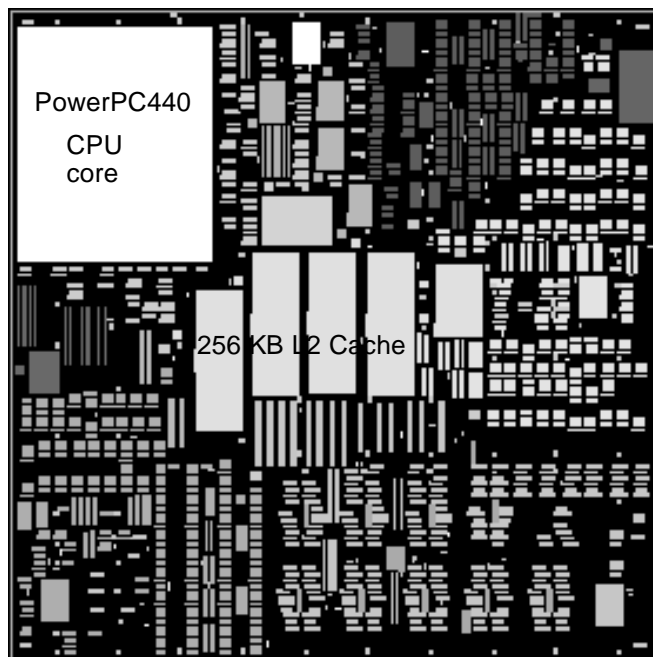


Figure 5: 9.6 x 9.6 mm Chip layout showing I/O circuits - PLL and DLL's, all SRAMs

Due to the large number of I/O (783) needed to integrate all the peripherals, the I/Os are placed all over the places in the die. A peripheral approach for IO implementation is possible with staggered structure, but it would have resulted in a larger die size, and a more noise sensitive part because of large simultaneous switching.

The device is based on an ASIC with integration of software based core - also named IP's - at the exception of the PowerPC CPU core which is a precharacterized hardware with optimized timing analysis and tuned clock distribution to achieve 800MHz. By comparison same CPU core runs only at 600MHz if implemented as a soft core with the best optimization tools.

Logic is described in Verilog and logic synthesis done with Synopsys synthesis tool. The physical design including floorplanning, placement and wiring has been done with IBM proprietary Chip Bench. Special care was taken in physical implementation for minimization of noise induce by coupling and simultaneous switching on top of the conventional signal integrity verification.

Extensive simulation of each core as simulation after complete integration has been done, resulting in a first pass good product.

IX- TEST RESULTS

A special board with modular approach for PCI-X and DIMMS peripherals attachments has been developed. It permits to debug the SOC device with DDR1 and DDR2 SDRAM as PCI, PCI-X and PCI-X DDR connectors. Debug was done with the Riscwatch debugger through the JTAG serial link I/O.




Figure 6: Board used for debug with DDR2 DIMMS close to the IOP processor and PCI-X bus analyzer

CONCLUSION

A SOC integrating a PowerPC CPU core with large number of state of the art and conventional peripheral has been designed and tested good on its first pass of silicon. The CPU has been tested at 800MHz and SRDRAM DD2 at 533 MHz. At the time of this publication 667MHz SDRAM as PCI-X DDR 266 mode 2 were not tested because DDRII 667 DIMMS and bus analyzer was not available.

REFERENCES

- [1] IBM Corp. (1999) Coreconnect Bus Architecture, Hopewell Junction, NY. [Online]. Available: http://www.chips.ibm.com/techlib/techlib.nsf/product-families/CoreConnect_Bus_Architecture
- [2] IBM Corp.. (2000) PowerPC Embedded Cores, Hopewell Junction, NY. [Online]. Available: http://www.chips.ibm.com/techlib/techlib.nsf/products/PowerPC_440_Embedded_Cores
- [3] JEDEC STANDARD DDR2 SDRAM SPECIFICATION (Revision of JESD79-2) January 2004
- [4] W.Lau "Overcoming DDR-2 interface challenges" EDN, Jan 22, 2004
- [5] *PCI-X Addendum to the PCI Local Bus Specification, Version 2.0*