

digest of papers

February 28–March 4, 1994


spring

94



comp 

THIRTY-NINTH IEEE COMPUTER SOCIETY INTERNATIONAL CONFERENCE



IEEE Computer Society Press



The Institute of Electrical and Electronics Engineers, Inc.

MICROSOFT CORP.
EXHIBIT 1017

The PowerPC™ 603 Microprocessor: A Low-Power Design for Portable Applications

Sonya Gary[†], Carl Dietz[‡], Jim Eno[†],
Gianfranco Gerosa[†], Sung Park[‡], Hector Sanchez[†]

[†]Motorola, Incorporated, [‡]International Business Machines Corporation
Somerset Design Center, 9737 Great Hills Trail, Austin TX 78759

Abstract

The PowerPC 603™ microprocessor is a low-power implementation of the PowerPC Architecture™. The superscalar organization includes dynamic localized shut-down of execution units to reduce normal-mode power consumption. Three levels of static low-power operation are software programmable for system power management. The 603 PLL (Phase Lock Loop) is capable of generating an internal processor clock at 1X, 2X, 3X or 4X the system clock speed to allow control of system power while maintaining processor performance. Various design features optimize the 603 for both power and performance, creating an ideal microprocessor solution for portable applications.

1: Introduction

The market for a portable computer is dependent on its performance versus its portability and physical features. The role of the microprocessor in a portable system is important for both of these conflicting demands. The microprocessor must provide the performance and flexibility to support a range of applications while promoting the extension of battery life by minimizing power dissipation during both normal operation and standby.

The PowerPC 603 microprocessor was designed with the performance and low-power requirements of portable systems in mind. The 603 combines an efficient 32-bit implementation of the PowerPC RISC architecture and effective static and dynamic power management features. The 3.3V CMOS design can achieve a peak instruction rate of 3 per cycle and keep power dissipation under 3 watts at 80 MHz.

[†]In this document, the terms "PowerPC 603 Microprocessor" and "603" are used to denote the second microprocessor from the PowerPC Architecture family. PowerPC, PowerPC Architecture and PowerPC 603 are trademarks of International Business Machines Corporation. SPEC, SPECint92 and SPECfp92 are trademarks of Standard Performance Evaluation Corporation.

2: PowerPC 603 microarchitecture overview

The architectural organization of the 603 focuses on high performance, low cost and power management. Figure 1 shows a high-level block diagram of the 603.

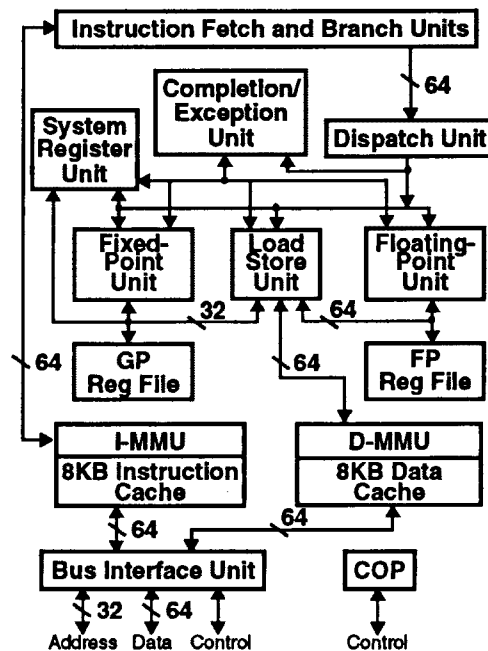


Figure 1: 603 Block Diagram

The instruction fetch unit fetches two instructions per cycle from the instruction cache into a six-entry instruction queue. Branches are folded out of this queue by the branch unit, which can then initiate fetching down either a sequential or target stream. Programmable static branch prediction is used for unresolved branches.

The dispatch unit may issue up to two instructions per cycle to four independent execution units. These units are the fixed-point unit, the floating-point unit, the load/store

unit and the system register unit. With branch folding, a maximum issue rate of three instructions per cycle can be achieved. Execution unit access to general purpose or floating-point registers is handled with rename registers managed by the dispatch unit.

The completion/exception unit tracks all dispatched instructions in a five-entry queue. A maximum of two instructions per cycle are retired in order from this queue. As instructions are retired, architectural register values are committed from the appropriate rename registers. With branch folding, up to three instructions can be completed per cycle.

The 603 contains separate 8KB, 2-way set associative instruction and data caches, each with a 32-byte line size. Both caches are blocking and allow one access per cycle. The data cache supports writeback operation and snoops, while the instruction cache does not support either. Memory management features include dual 64-entry instruction and data translation lookaside buffers, dual 4-entry instruction and data block address translation registers, and sixteen segment registers.

3: Dynamic power management

Dynamic power management is a software-enabled mode on the 603 which turns on power-saving logic for the execution units, caches and memory management units (MMUs) during normal operation. Once the mode is enabled, no other software intervention is necessary.

Dynamic power management logic automatically manipulates clock regenerators to reduce average power consumption. Power is reduced by eliminating clock switching and inhibiting change of registered values. In a static design, if registered values do not change, the logic those values feed does not switch. Clocks to a particular logic block are disabled or enabled on a cycle-by-cycle basis once it is determined whether that block is needed for instruction execution.

3.1: Dynamic power management implementation using clock regenerators

The 603 clock regenerators produce two clocks, C1 and C2, which feed master and slave latches, L1 and L2, respectively. Each clock regenerator features two 'freeze' inputs which are used for dynamic power management control. The assertion of these inputs force C1 and C2 low. The logic dedicated to support these freeze inputs in the clock regenerators accounts for under 0.3% of the total chip area.

Timing constraints on the assertion of freeze inputs are fairly strict. C1 freeze must be defined by outputs of an L2 latch to ensure freezing the correct C1 pulse and not that

of the previous cycle. It must also be early enough to fully block the next C1 pulse. If the assertion of C1 (C2) freeze is late, a small C1 (C2) pulse may be generated. Functionally this is not a problem since the 603 is a fully static design, and latched data is a don't care during the time C1 and C2 are frozen. However, late freeze assertions do not allow maximum power savings. Power is consumed by the extra clock switching and potentially by spurious activity due to data changes. Therefore, it must be determined early in a cycle if a logic block must be clocked at the end of that cycle. This determination produces the C1 freeze. This freeze is fed through a transparent L1 latch to generate the C2 freeze. This is done to ensure a freeze of the correct C2 pulse and not the one of the previous cycle.

Timing constraints on the negation of freeze signals are based on minimum clock pulse width requirements for L1 and L2 and maintaining sufficient usable cycle time. Negation of both C1 and C2 freezes must be early enough to allow at least the minimum pulse width to latch new values so that the 603 functions correctly. The negation of C2 freeze should also be early enough to allow a full C2 pulse so that usable cycle time is not sacrificed by delaying L2 outputs.

C1 freeze and C2 freeze serve to disable or enable a clock regenerator on a cycle-by-cycle basis so that there is no performance sacrifice to the affected logic. Figure 2 shows the required timing for turning the C1 and C2 clocks off and on using the C1 and C2 freeze signals.




Figure 2: Freeze timing for C1 and C2

3.2: Execution unit dynamic power management

The separation between execution units in the superscalar organization of the 603 allows for independent

dynamic power management in each of the four units. The dynamic power management control in each unit is based on logic already present for instruction dispatch and execution. Only 0.08% of the total chip area of the 603 is attributable to dynamic power management in the execution units.

The 603 execution units employ dynamic power management in a simple fashion. This involves distributed logic for clock freezing based on pipeline stages or the particular instruction dispatched to a unit. The clocks feeding the front-end instruction buffer for each unit are enabled if a valid instruction is being executed in that unit or if a valid instruction assigned to that unit is present in the instruction dispatch buffer. Whether the instruction is actually dispatched is determined too late in the cycle to be part of the C1 freeze equation. Other stages of the execution unit pipelines are frozen separately based on activity each cycle.

The load/store unit is a good example of the efficiency of this method. The last stage of the store pipe is a completed store queue which holds a store until the cache is available or there is a load dependency. Thus a store may remain in the completed store queue for some time. The 603 may freeze the other stages of the store pipeline despite stores remaining in the completed store queue.

603 execution units also freeze clocks based on the particular instruction dispatched to a unit. The system register unit of the 603, which operates on architectural control registers, is a prime candidate for this method. This unit has only one single-cycle execution stage but manages many registers. Clocks to each register are enabled only when that register is being modified, while clocks to the other registers in this unit remain frozen.

3.2.1: Estimating effectiveness of dynamic power management in the execution units

The effectiveness of using dynamic power management in the execution units varies with the type of code run. For example, the floating-point unit clocks will be frozen continuously if integer-only code is run. If code is scheduled such that all of the units remain busy continuously, few clocks will be frozen.

In evaluating the power saving potential of dynamic power management in the execution units, the percentage idle time of each execution unit was estimated while running different types of code. This was done using an architectural modeling tool, the Basic RISC Architecture Timer (BRAT) [1], which collected statistics while running SPEC92 benchmark traces. Cycles when the dispatch buffer contained an instruction for a unit were not included as idle time for that unit. Though idle time figures indicate the time an entire execution unit may be frozen, they do not account for the time during instruction

execution when some portions of that unit may remain frozen.

The BRAT results, shown in Figure 3, illustrate that for some applications, each functional unit may be idle during a large percentage of the run time. Since dynamic power management freezes execution unit clocks during this idle time, it can be effective in reducing the average power consumption of the execution units.




Figure 3: Execution Unit Idle Time for SPEC92

3.3: Cache and MMU dynamic power management

The 603 has blocking caches. This ensures that when a cache miss occurs, all other accesses are held off until miss data is returned from memory. During this time the caches and MMUs are idle and can have their clocks disabled by dynamic power management logic while waiting for data. Depending on the memory latency and the processor to bus clock ratio, this idle time could be many cycles. Blocking caches allowed a straightforward implementation of dynamic power management in the caches and MMUs with little impact on full chip area. The entire dynamic power management logic for the caches and MMUs is approximately 0.20% of the total chip area.

3.3.1: Data cache dynamic power management

Dynamic power management logic freezes all clock regenerators to the data cache while waiting for miss data. This is true for all single beat or burst reads from memory. If a miss in the cache requires a writeback to memory, or castout, the clocks will be frozen after the castout is complete. If a miss in the cache does not require a castout, the clocks are frozen immediately after the miss address is sent to the bus interface unit. The clocks remain frozen until data is returned from memory. Clocks are automati-

cally enabled for each beat of data of a burst read. If there are multiple processor clocks between data beats, the clocks will be frozen between each beat. Figure 4 shows the data cache power management state machine.

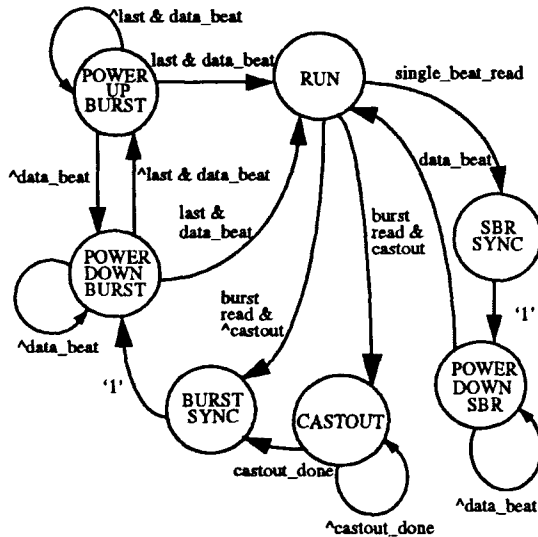


Figure 4: Data Cache Power Management State Machine

3.3.2: Support for bus snooping

The 603 supports bus snooping to maintain memory coherency. If the data cache clocks are frozen due to a miss in the cache, the dynamic power management logic is required to unfreeze the data cache clocks to service snoop requests. If a snoop hits to a modified line in the cache, the data cache clocks will remain enabled for the snoop castout. If the snoop misses or hits to an unmodified line, the data cache clocks can be frozen until it receives a data beat from memory or another snoop request. Figure 5 shows the state machine used to control the unfreezing of clocks for snoop lookups and snoop castouts.

The data cache power management state machine and the snoop power up state machine do not interact. They are totally free-running and operate independently, but together they define the C1 freeze signal for the data cache:

$C1_FREEZE = ((POWER_DOWN_SBR \& \wedge data_beat) \vee (POWER_DOWN_BURST \& \wedge data_beat)) \& SNOOP_IDLE \& \wedge snoop_request$

This freeze equation allows the data cache to be powered down due to a read miss, then to power up for a snoop request and snoop castout, and then power back down if read data has not been returned from the bus. This approach was found to be much simpler than using one state machine to handle the coordination of cache misses with snoop accesses.




Figure 5: Snoop Power Up State Machine

3.3.3: Instruction cache dynamic power management

The dynamic power management logic for the instruction cache will freeze the clocks on all burst or single beat read misses. The state machine which controls this is identical to the data cache state machine except that the castout section of the state machine is omitted because the instruction cache does not include a modified state. Also, the instruction cache is not snooped so there is no need to power up the instruction cache while waiting for data from the bus.

3.3.4: MMU dynamic power management

Both instruction and data memory management unit clocks are frozen for any burst or single beat read. The MMUs only need to be operational for initial lookups and not during snooping or for data beats from memory, so they are frozen for the entire miss. However, system register unit accesses to the MMUs may occur while the clocks are frozen. If this situation is detected, the dynamic power management logic will unfreeze the clocks until the next power down condition is met. Figure 6 shows the state machine which controls the dynamic power management of the MMUs.

3.3.5: Estimating effectiveness of dynamic power management in the caches and MMUs

To determine the potential power savings of the dynamic power management logic in the caches and MMUs, the percentage of time that the cache would be idle waiting for data from memory was estimated. It was found that the cache would be idle 60% of the time waiting for data, assuming the following: a 6-1-2-1 memory system, a 2:1 processor to bus clock ratio, an 80% cache hit rate, a cache access every cycle, and no bus snooping. Therefore, dynamic power management logic could freeze the clocks

to the caches approximately 60% of the time. The MMUs can be frozen for a greater percentage of the time because they are only used for initial lookups and their clocks do not need to be enabled for each data beat from memory.




Figure 6: MMU Power Management State Machine

3.4: Dynamic power management results

Table 1 shows total power measured (instantaneous current measured manually) for various applications with 603 silicon using dynamic power management. Power figures include both internal and external power consumed; internal power is measured for internal Vdd supplied to the 603 only, while external power is measured for external Vdd supplied to 603 I/Os plus I/O external pull-ups and terminators. Table 2 shows the percentage decrease in internal power dissipation for the same applications if dynamic power management is used. All of the dynamic power management logic accounts for approximately 0.6% of the total 603 chip area.

Table 1: Dynamic Power Management Results

Freq (MHz)	Power (W)*				
	C linpack	dhrystone	hanoi	heapsort	nsieve
25	0.96	0.85	0.86	0.83	0.90
33	1.16	1.06	1.07	1.02	1.13
50	1.59	1.49	1.48	1.45	1.59
66	1.97	1.86	1.84	1.80	1.98
80	2.21	2.20	2.17	2.12	2.33

*3.3V nominal, room temperature, 80MHz in 2:1 bus mode

Table 2: Percentage Decrease in Internal Power Dissipation Using Dynamic Power Management

C linpack	dhrystone	hanoi	heapsort	nsieve
8.5%	14.0%	13.8%	14.2%	16.4%

4: Static power management

The 603 provides three static power management modes, Doze, Nap and Sleep, which are programmable through a hardware implementation register. Static power management allows power management software or an operating system to reduce average power consumption when the 603 is idle for any extended period of time. The names Doze, Nap and Sleep are indicative of the progressive increase in power savings obtainable. Once any one of the static power management modes is enabled, the 603 completes execution of all outstanding instructions and achieves a quiescent state. Once the quiescent state is reached, the 603 disables all clocks to units that are not required to be functional during a particular power management mode. Once clocks are stopped and the processor is in one of the power-saving modes, an external event asserting one of the wake-up signals, such as the external interrupt input pin, will bring the 603 out of that mode. The 603 will resume instruction execution by jumping to the address of the appropriate interrupt vector. This operation is common to all three modes. A more detailed description of the individual modes appears in the following sections.

4.1: Doze mode

Doze mode allows the 603 to maintain cache coherency while in a power-saving mode. The snoop logic and the data cache are kept active to service snoops as they occur on the bus. If a snoop hit occurs, the necessary cache update and a snoop copyback bus operation will occur while the 603 remains in Doze mode. However, if an address parity error occurs during a snoop, or a bus error during a snoop copyback causes a machine check condition, then the 603 will exit the Doze mode and take a machine check interrupt.

Along with the snoop logic, the time base/decrementer logic is kept active in Doze mode. This provides uninterrupted timer functionality during Doze. A decrementer interrupt will cause the 603 to exit Doze mode. Asserting any of the pins INT_, SMI_, MCP_, SRESET_ or HRESET_ will also cause the 603 to exit Doze mode in ten system clocks (SYSCLKs) or less. Using a 1:1 processor to bus clock ratio, this is the worst case logic delay from the external pin assertion to all clocks being enabled.

4.2: Nap mode

In Nap mode, all logic is disabled except the time base/decrementer logic. Since cache coherency cannot be maintained in Nap mode, the system program must flush the data cache before entering Nap mode if system memory is expected to be altered while in Nap mode. Otherwise, a data cache flush is not necessary since the 603 keeps the cache contents unaltered while in Nap mode. A pair of handshake signals, QREQ_o and QACK_o, are provided for the system to allow the 603 to go into Nap only when cache coherency will no longer be a problem. When the 603 is ready to enter the Nap mode, QREQ_o is asserted. The system logic responds with an active QACK_o to allow the 603 to proceed with entering Nap mode.

Compared to Doze mode, further power savings are achieved with the data cache and bus snooping logic disabled. Furthermore, the receivers of most input and bidirectional pins are disabled for added power savings. Outputs maintain their normal idle state. A decrementer interrupt or assertion of any one of the pins mentioned previously with regard to Doze mode will wake the 603 from Nap mode. As in the Doze mode, the wake-up latency is ten SYSCLKs or less.

4.3: Sleep mode

Sleep mode allows a maximum power savings by disabling the clocks to all units. The same QREQ_o/QACK_o handshake protocol exists for Sleep mode as for Nap. Unlike the Doze and Nap modes, more power savings can be achieved in Sleep mode by disabling the PLL and SYSCLK. In Doze and Nap modes, the SYSCLK and PLL configurations must remain in the same state as they were prior to the commencement of the power-saving mode. Sleep mode provides the ability to dynamically manage system power by allowing system logic to disable the PLL or SYSCLK, change the SYSCLK frequency, or change the processor to bus clock ratio through the PLL configuration pins.

Since Sleep mode does not automatically disable the PLL, the system logic can implement several different levels of power savings depending on the wake-up response time requirements. For example, if a quick wake-up from sleep is required, then by leaving the SYSCLK input and the PLL configuration unaltered, the 603 will wake from Sleep within ten SYSCLK cycles or less. However, if maximum power savings is a requirement, then the PLL and SYSCLK may be disabled completely, reducing power dissipation to leakage levels. However, when coming out of Sleep mode, a maximum of 200 usec is required for the PLL to relock to the new SYSCLK frequency. After this relock time, any of the external pins mentioned

in the Doze section may be applied to wake-up from Sleep.

4.4: Static power management results

Table 3 shows total power measured with 603 silicon using the static power management modes. Sleep figures are listed for configurations with the PLL and SYSCLK enabled, with the PLL disabled and SYSCLK enabled, and with both the PLL and SYSCLK disabled.

Table 3: Static Power Management Results

Freq. (MHz)	Power (mW)*				
	Doze	Nap	Sleep		
			PLL on CLK on	PLL off CLK on	PLL off CLK off
25	133.2	49.4	38.8	12.8	4.7
33	168.0	62.0	47.8	13.5	4.7
50	241.7	88.8	66.2	15.1	4.7
66	307.1	113.0	88.5	17.7	4.7
80	366.1	135.1	105.5	19.3	4.7

*3.3V nominal, room temperature, 80MHz in 2:1 bus mode

5: Other low-power design features

Several other areas of the 603 were targeted for low-power design. These include the caches and memory management units, the phase lock loop, clock distribution and clock regenerators, standard cell and datapath libraries and the bus interface unit.

5.1: Cache and memory management design

The data portion of each cache is divided into eight 1KB subarrays, which were specifically designed for low-power operation. Each subarray holds one word of the cache line for both sets of data. This organization was chosen to minimize power: since each cache employs double-word datapaths, there is never a need to access more than two subarrays during any given cycle. If only one word of data needs to be read or written, only one subarray is accessed.

Each subarray is held in a constant state of precharge if the array is not being accessed. That is, each bitline and output driver of each subarray is constantly being precharged when the subarray is idle. To save power, the pre-

charge signals are not clocked and only turn off while the subarray is being accessed. The precharge of the output drivers is disabled only if the access is a read access.

Special attention was given to the timing of the precharge signals and the word lines in each subarray to guarantee that no bitcell is ever enabled while the bitline precharge devices are enabled. To further save power, each subarray employs pulsed word lines to eliminate unnecessary bitline discharge once the sense amplifiers have been strobed. The sense amplifiers within each subarray function as data latches once they have been strobed, eliminating the need for clocked output data latches. For speed reasons, the sense amplifiers for both sets of data within each subarray are strobed for load hits. For castouts, however, only the sense amplifiers for the data set being cast out are strobed. Each subarray contains a way-select mux, allowing both sets of data to share the same output drivers.

The tag portion of each cache consists of a single array. To save power, each tag array is precharged when idle. Due to the speed requirements of supporting a read-modify-write operation each cycle, the precharge signals within each tag array are clocked. Like the data subarrays, each tag array employs pulsed wordlines and strobed sense amplifiers. Unlike the data subarrays, however, each tag array's wordline pulsing logic is designed to limit the bitline discharge to approximately one volt, regardless of frequency. Limiting bitline discharge to one volt saves power at lower frequencies, where one of each pair of bitlines would discharge completely.

The 603's separate instruction and data MMUs each consist of three arrays: a sixteen-entry Segment Register array, a 4-entry Block-Address-Translation (BAT) array, and a 64-entry, 2-way set associative Translation Lookaside Buffer (TLB) array. Like the cache tag arrays, each of these arrays is precharged when idle, but the precharge signal is clocked. Each of these arrays is so small and fast that sense amplifiers are not required and thus pulsing the word lines of these arrays is unnecessary. This provides a further savings in DC power. The TLB array contains a way-select mux, allowing both sets of TLB data to share the same output drivers.

5.2: Phase Lock Loop

The 603 integrates a synthesizing analog PLL which maintains proper edge alignment between the 603's internal clock and SYSCLK. In addition, the 603 PLL can be configured during power up for a processor clock frequency that is 1X, 2X, 3X, or 4X the bus clock frequency. This allows reduced system power for a given level of processor performance. The PLL is designed to lock to a wide range of SYSCLK frequencies from 16 MHz up to

66 MHz. If lower frequencies are desired, the PLL can be bypassed altogether. During Sleep mode, all analog circuits used in the PLL design are completely shut down in order to attain milliwatt performance.

5.3: Clock distribution

A passive clock distribution network was chosen over a set of distributed clock buffers in order to minimize active power. An HTREE style distribution network (see Figure 7) using metal3 and metal4 produced the best rise/fall time performance while maintaining a low total routing capacitance. In addition, skew was minimized among all branches. Grid-like layouts were discounted due to their high routing capacitance. Global clock skew was kept under ± 100 pS with this scheme; rise and fall times were under 1.2 nS while average power due to the HTREE driver and the HTREE itself was kept under 100 mW at 80 MHz.

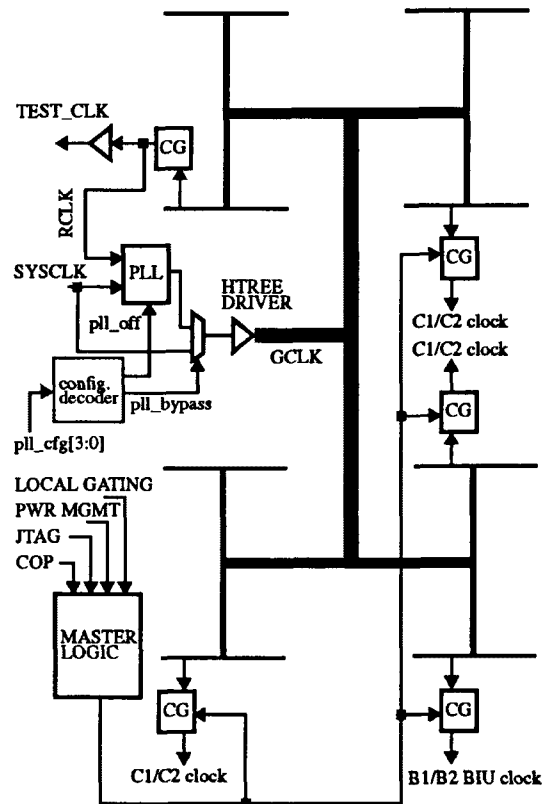


Figure 7: HTREE Clock Distribution Network

5.4: Clock regenerator design

The output drivers of all 603 clock regenerators are designed with enough granularity to cover a wide capacitive load range from 0.2 - 2.2 pF. As a result, lower average power is obtained as compared to other schemes which use 'dummy' loads in order to maintain constant clock latencies.

5.5: Standard cell and datapath libraries

Static logic design is used throughout the standard cell and datapath element libraries where the (speed X power) product is optimized. Special care was taken during logic synthesis in repowering standard cells for greater drive capability. Likewise, datapath elements were only repowered in order to meet certain timing constraints.

5.6: Bus Interface Unit

In normal operation, the bus interface unit (BIU) reduces power consumption by turning off the 64 bit data bus receivers when the 603 is not reading from the bus. In this way power is saved whenever the data bus is used by other masters for operations such as DMA transfers or when the 603 is performing write operations. Dynamic power management logic is not implemented in the BIU due to bus snooping requirements. Static power management modes do affect the BIU. In Nap and Sleep modes, clocks to the BIU are disabled along with any power-consuming circuitry, such as a voltage reference generator. Also, when waking up from Nap or Sleep, the BIU resumes normal operation only after the bus clock and processor clock have been synchronized.

6: CMOS Technology

The 603 is fabricated with a 3.3 volt 0.5 micron CMOS technology using four layers of metal. Split instruction and data caches account for over half of the 1.6 million transistors used within a 85 mm² die. The design is fully static and LSSD compliant. 603 physical characteristics are summarized in Table 4.

7: 603 performance

The 603 design is optimized for both power and performance. Table 5 provides estimated SPEC92 performance for an 80MHz 603-based system using a 2:1 processor to bus clock ratio and a second level external cache [2].

8: Conclusion

The PowerPC 603 microprocessor is a low-power implementation of the PowerPC Architecture with power management capability and competitive performance. The area-conscious design also makes the 603 a cost-effective alternative. The 603 combines all of the necessary features to make it the ideal microprocessor solution for portable applications.

Table 4: 603 Physical Characteristics

Technology:	0.5 um Nwell CMOS, four-layer metal
Die Size:	7.39 mm X 11.47 mm
Transistors:	1.6 million
Memory:	split 8KB caches
Bus:	split 32b address, 64b data, 16 - 66 MHz
Voltage:	3.3 V nominal
Power:	3.0 Watts Max @ 80 MHz
Package:	240 pin wire-bond CQFP, C4 CQFP
Frequency:	16 - 80 MHz
Processor to bus clock ratios:	1:1, 2:1, 3:1, 4:1
I/Os:	165 signal, CMOS/TTL, 5V tolerant

Table 5: 603 Performance Estimate

Frequency	SPECint92	SPECfp92
80MHz	75*	85*
*estimated based on simulated results		

Acknowledgments

The authors gratefully acknowledge the 603 design team for their support in the design, implementation and verification of the power-related features of this part. In particular, Peter Ippolito provided the clock regenerator design which made the implementation of the power management features possible. Also, the support and promotion of the 603 as a power-conscious microprocessor by the 603 project managers, Arturo Arizpe and Jim Kahle, was invaluable.

References

- [1] Poursepanj, A. et al., "The PowerPC™ 603 Microprocessor: Performance Analysis and Design Trade-Offs," *Proceedings of COMPCON 1994*, February 1994.
- [2] Burgess, B., et al., "The PowerPC™ 603 Microprocessor: A High Performance, Low Power, Superscalar RISC Microprocessor," *Proceedings of COMPCON 1994*, February 1994.
- [3] Arizpe, A., and Kahle, J., "The Future Direction of PowerPC™," The Sixth Annual Microprocessor Forum, October 18, 1993.