


# PowerPC Embedded Processors Application Note



## System Design Considerations

March 21, 1996

IBM Microelectronics  
Dept D95/Bldg 060  
3039 Cornwallis Road  
Research Triangle Park, NC 27709  
Version: 1.0

*Abstract - This document addresses topics and questions frequently asked of PowerPC Embedded Processor Applications engineers. Much of this information can be found in the PPC403Gx User's Manuals and the associated data sheets.*

### Clock Symmetry

Readily available oscillators are generally specified as providing 60/40 or 55/45 percent duty cycle symmetry. The PPC403Gx DRAM controller uses the negative edge of SYSClk for generating /CAS and, in some modes, /RAS. In contrast, all other signals are relative to the rising edge of SYSClk. Designers should account for the effect of duty cycle variation when timing their DRAM interface. One option which reduces symmetry concerns is using an oscillator at twice the desired frequency and dividing its output by two.

### Bit Ordering and Programmable Static Bus Sizing

The PowerPC architecture specifies that bits are numbered left to right, starting with bit 0. Therefore, a 32-bit bus is designated D[0:31] with D[0] being the most significant and D[31] the least significant bit.

The PPC403Gx processors provide programmable static bus sizing on a per-bank basis. This greatly simplifies interfacing external devices of less than 32-bits because no byte steering logic is required. To make use of programmable static bus sizing, connect external devices as follows:

Word	D[0:31]
Halfword	D[0:15]
Byte	D[0:7]

Static bus sizing is advantageous in that it allows direct connect of 8- and 16-bit devices to the PPC403Gx. From the viewpoint of a programmer there is no difference between a load word instruction directed to an 8-bit device or to a 32-bit device. In the 32-bit case, fetching the requested data requires one access whereas four bus operations are used for the 8-bit case. The difference becomes significant in systems where bus latency is of concern to the designer, especially in regards to cache line fill and line flush operations.

The PPC403Gx processors typically execute out of cache over 80% of the time. If there is no external bus master or DMA activity during this time, the address and data buses are unused. To save power, the



data bus is tristated when not in use. In addition, during byte and halfword write operations, the unused data bus byte lanes are high-Z. Treatment of the address bus between accesses is different depending on the processor selected. The PPC403GA always tristates the address bus between accesses, the PPC403GB continues to drive the last address, and the PPC403GC allows the user to select between these two options. The ability to hold the last address has been added to prevent noise from causing undriven lines to approach the switching threshold of receivers on the line. If this happens and external multiplexers or drivers on the address lines do not have sufficient input hysteresis, their outputs may switch needlessly, creating additional noise and increasing the total system current draw.

## Cache Implications

To improve system performance, PPC403GA, PPC403GB and PPC403GC processors have 2 kB instruction and 1 kB data caches. Both caches are two-way set associative and organized with a 16-byte line size. Following a reset the caches are disabled and software must explicitly invalidate them before enabling them.

When the processor requests data from a cacheable region and that data is not in the cache, a cache miss occurs. If the external bus is not performing high priority DMA or in use by a bus master, the cache line fill operation is on the external bus two cycles after the miss.

The processor performs a line fill either sequentially from the cache line starting address or target word first, depending on the bank configuration register related to the given memory address. Target word first begins the memory access at the requested word, continues to the end of the cache line and then fetches the beginning of the cache line. With sequential line fills the access starts at the beginning of the cache line and hence the processor must wait until the requested data word is accessed. As such, sequential line fill mode should only be used when interfacing to devices, such as burst mode ROMs, that do not support target word first accesses. The requested data item becomes available to the processor one cycle after it is read from memory. When a cache miss occurs on the data side and the new data displaces a dirty cache line, a cache line flush immediately follows the fill operation. In this case the bus is busy for twice the line access time. Note that line flushes are always sequential from the starting address of the cache line.

Because cache line fill and flush operations are atomic bus operations with respect to themselves, no other operations, such as DMA or an external hold request, can interrupt them. In systems with hard real-time requirements, the time required for fill and flush operations must be considered when calculating bus latency. With 32-bit wide DRAM operating at 2-1-1-1 a line operation consumes six cycles. The bus then remains busy for an additional cycle during the DRAM precharge cycle. In contrast, an 8-bit memory operating at 2 wait states per transfer requires 48 bus cycles to provide the same amount of data to the processor.

Since only line operations exploit the advantage of fast page mode DRAM, the best technique for initializing external memory regions is via cache line flushes. Because the current data in external memory is overwritten with the new data, cache line fills are not required. Software can accomplish the initialization of a cache line without causing the old data to first be filled into the line, by using the `dcbz` instruction. Normal stores can then be used to set the desired data into the line, and then `dcbst` can be used to efficiently flush the line to memory.

The PPC403Gx family of processors double maps memory in that A0 does not participate in address decoding. Therefore, a given block of external memory can be accessed with bit A0=0 or A0=1. By appropriately programming the data and instruction cache controllability registers (ICCR and DCCR) one of these two addresses can be cacheable, whereas the other is not. For some tasks the ability to concurrently reference the same area of memory both cached and non-cached may greatly improve

system performance. For example, clearing a region of DRAM memory is fastest using the cache because fast page mode (burst) memory access cycles are used during the line flush process. In contrast, scatter/gather type operations generally require less bus activity when performed in non-cached mode.

## Cache Coherency Issues

The PPC403Gx processors do not provide hardware support for maintaining cache coherency during DMA and external bus master operations or in a multi-processor configuration. For example, DMA operations read and write memory without regard to the contents of the on-chip cache. To ensure that the cache and external memory remain coherent during these operations software must be written such that DMA and bus mastering activities only reference areas of memory that are not present in the cache.

## Cache Flushing

As stated above, explicitly flushing the cache is required in situations where external devices can read or write memory that might be present in the cache. In these cases, software must ensure that prior to allowing external devices access to the bus that the cache and external memory are coherent. The PPC403Gx processors do not provide an instruction that flushes the cache. There are two avenues which a program may pursue. If the addresses for which coherency must be enforced are known, then data cache block flush (dcbf) instructions should be issued for these addresses. This technique is well suited for cache-aware code written to exploit cache performance advantages.

When the contents of the cache are unknown it will be necessary to flush the entire cache to ensure coherency. Because it is not possible to flush by congruence class, another technique is required to provide reasonable performance. The recommended algorithm is as follows:

- Execute a series of 64 dcbz instructions to a 1 kB block of scratchpad memory. Since dcbz establishes the address in the cache without actually making a BIU request, it is possible to use an address that is not even valid in the address space. However, the addresses used must be designated cachable, as dcbz to a non-cachable address will cause an alignment error. This series of dcbzs fills the cache with the scratchpad addresses, thereby displacing all previous cache lines, causing the flushing of the dirty data.
- Issue a series of 32 data cache congruence class invalidate (dccc) instructions, one for each congruence class. This eliminates the dirty “garbage” scratchpad addresses established by the dcbzs.

This sequence of operations flushes out to memory all previously dirty contents of the data cache, without causing subsequent line fills to the dcbz addresses since dcbz establishes a line in the cache without actually loading that line from memory. In establishing these lines, dirty cache lines are flushed from the cache. The series of dccc instructions then removes the dcbz addresses and data from the cache, again without actually flushing the data to memory. It is necessary that interrupts be disabled during this cache flushing process.

Care must be taken when working with areas of memory using both cached and non-cached accesses. Before accessing memory in a non-cachable manner, software must ensure that the location is not present in the cache. A non-cachable access to an external memory location that is also present in the cache is considered a programming error and in addition is a violation of the PowerPC architecture.

## IVR Pin



The IVR pin connects to a diode stack on each of the 403Gx I/O pins and provides a current path for preventing transient overvoltages from damaging the processor. In systems where all the devices that drive into the PPC403Gx operate at +3.3V, IVR should be connected to +3.3V. Otherwise IVR must be connected to +5V. Note that IVR does not determine the high output level; the PPC403Gx always drives +3.3V for logic one conditions. As a result, to minimize static current draw, any pull-up resistors should be connected to a +3.3V source.

## **TestA, TestB**

These pins are used for factory test of the product. For normal use, TestA must be connected to ground and TestB to +3.3V or +5V.

## **Single Cycle Peripherals**

Single cycle peripherals require external logic to gate /OE and the desired /WBEx signals with a secondary clock to generate read and write pulses. Care should also be taken to ensure that the peripheral places its drivers in the high-Z state quickly enough to avoid data bus contention during back-to-back read/write operations. This can be a concern because unlike some processors, those in the PPC403Gx family do not insert turn-around or dead cycles between transactions.

## **Timing Specifications**

The PPC403Gx data sheet provides timing information that includes minimum and maximum specifications for items such as output delays and setup times. Designers should note that for a given device these parameters track each other in a roughly linear fashion. A fast chip has outputs that become valid with the published minimum output delay, inputs that actually require less setup time than specified and input hold times as given in the data sheet. Conversely, a slow device requires more input setup time, less hold time and drives its outputs later in the cycle.

## **DRAM Interface**

The PPC403Gx embedded processors allow direct connect of DRAM devices via an integrated DRAM controller. The only exception is the required addition of an external address multiplexer with systems that make use of external bus masters. This is needed because the processor normally multiplexes the DRAM row and column address onto the address bus during internal operations. Since a bus master provides its own address, the multiplexer allows the external master to drive the full 32-bit address, and then the DRAM controller to gate the row and column address to the DRAM at the appropriate time.

When the processor is initialized, all of the DRAM RASs are disabled until the corresponding bank registers are programmed. During this time, and until the bank is configured for a DRAM device, no refreshes are generated. As such, for most DRAM devices it is necessary to wait until eight refresh cycles have occurred before proper DRAM operation is guaranteed. During this start up time, code must not attempt to use the just-enabled DRAM.

Access time for fast page mode DRAM devices is determined by the relationship of several signals. Assuming that address setup times are met for RAS and CAS becoming active, data is available a deterministic time after the RAS and CAS signals become active. To improve DRAM system performance, RAS should ideally become active immediately after the address becomes valid. In real systems this is problematic and the easier solution is to drive RAS active one-half cycle after the address is driven. Depending on the system clock, this delay in driving RAS may require running the memory subsystem at 3-2-2-2 rather than 2-1-1-1 timing.

The PPC403Gx processors provide an option for the timing of the DRAM RAS signal. Normally, RAS is activated one half cycle after the address is driven by the processor. Using Early-RAS Mode, an internal delay line provides a RAS signal that becomes active a quarter clock cycle after the address is driven. Because an internal delay line is used to generate this timing, when Early RAS Mode is employed in a system that stops the system clock a delay of 700 ns is required before DRAM can again be accessed after the clock is restarted.

When DRAM Read on CAS is selected in the IOCR, DRAM read data is latched on the rising edge of CAS instead of on the rise of SYSCLK. This mode is provided to allow more access time and to compensate for some amount of board capacitance by latching data on the CAS line deactivation. In other words, the PPC403Gx latches data when CAS actually goes high, not when it begins to drive CAS high. As stated in the datasheet, capacitive loading must not delay the CAS low-to-high transition such that the time from the CAS rising edge to the next SYSCLK rising edge is less than 15.5 ns.

## **Chip Sockets**

AMP manufactures a socket assembly which can be used with the 160 pin PQFP PPC403GA and PPC403GC processors. The socket assembly consists of a socket (P/N 822114-1) and a cover (P/N 822115-1). Additional information including dimension drawings and a recommended PC board footprint is available via AMP Fax at 1-800-522-6752.

## **Ready and Bus Timeout**

The PPC403Gx processors provide a READY line for device paced transfers. Typically, after receiving a valid chip select a peripheral drives READY low until it is ready to proceed with the access. On the PPC403GA and PPC403GB processors, READY is a synchronous input and hence the designer must ensure that setup and hold time requirements are met. When driving READY from an asynchronous source the peripheral's READY signal must be synchronized to the system clock by passing it through a flip-flop triggered off the system clock. During reads, data is latched into the processor one cycle after READY is sampled high. In contrast, write operations complete either in the cycle READY becomes active or one cycle later depending on the setting of the WEoff bit in the appropriate bank register.

The PPC403GC provides identical functionality with the exception that setting bit 24 in the IOCR configures READY as an asynchronous input. In this mode READY synchronization takes place within the processor and as a result two cycles are added to the timings previously discussed.

A 128 cycle bus timeout is included in the PPC403Gx processors to prevent an errant device from causing system lockup. If an access, including not ready time, extends beyond 128 cycles an imprecise exception occurs. This timeout cannot be disabled on the PPC403GA and PPC403GB. Furthermore, because the exception is imprecise, the processor likely cannot be restarted from the exception point. Designs with peripherals that require more than 128 cycle access times should use the PPC403GC. This processor allows the device-paced timeout to be disabled by setting bit 21 in the IOCR.

## **Reset**

Systems must include a power on reset function to activate /RESET. /RESET should be driven by an open-collector or open-drain driver in combination with an external pullup resistor that can be overdriven by the processor's pull-down driver. Because the /RESET input is not Schmidt triggered, designers should avoid simple RC circuit power on logic.

Upon recognizing a system reset, initiated by either software or the falling edge of /RESET, the processor drives /RESET active for 3 (GA & GB) or 2048 (GC) clocks. After the given number of cycles have



occurred the processor stops driving /RESET and waits until /RESET rises enough to be recognized at an inactive level. The processor latches the value of the boot width pin (BOOTW) from two clocks before the /RESET line is sensed as negated and compares it to the BOOTW pin value after /RESET is inactive. If BOOTW follows /RESET the boot memory is configured for 16-bit operation, whereas BOOTW tied low indicates an 8-bit boot memory and BOOTW high 32-bit memory. Once /RESET is negated the processor begins operation by fetching the first instruction from address 0xFFFFF7FC. Note that the 403GC waits 32 clock cycles after /RESET is negated before fetching the first instruction.

IBM will continue to enhance products and services as new technologies emerge. Therefore, IBM reserves the right to make changes to its products, other product information, and this publication without prior notice. Please contact your local IBM Microelectronics representative on specific standard configurations and options.

IBM assumes no responsibility or liability for any use of the information contained herein. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of IBM or third parties. NO WARRANTIES OF ANY KIND, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE ARE OFFERED IN THIS DOCUMENT.