

Application Note

AN2289/D
Rev. 0, 6/2002

MPC8xx QMC Usage
Advisory

Heinz Wrobel
Freescale GmbH,
Munich

Part I Introduction

QUICC multichannel controller (QMC), support in some of the MPC8xx family members is quite powerful. However, there are a few issues in setting up a working and stable QMC setup that should be well noted. This document should be used in addition to the *MC68360 QMC Supplement to MC68360 and MPC860 Users Manuals (MC68MH360UM/D)*, hereafter referred to as the QMC manual. It will highlight some issues and add clarifications to the manual. The page numbers and chapters mentioned refer to the QMC manual unless otherwise specified. Note that the information contained in this document also applies to the QUICC MC68MH360, even though references are given based on the MPC8xx family.

1.1 Document Conventions

In the MPC8xx, PowerPC architecture based microprocessor, the MSB is labeled 0, the LSB is labeled 31. The standard for RAMs, ROMs, etc., is that the MSB is labeled as the highest number, i.e., 31 or 15, and the LSB is labeled 0. The notation “x:y” is used to denote bits numbered “x” to “y” in that order. References to the MPC860 User’s Manual refer to the MPC860UM/AD Revision 1. Binary numbers have a “0b” prefix. Hex values have a “0x” prefix.

Part II QMC Manual Clarifications and Corrections

- Page 2-1, Figure 2-1: For the MPC860, the Parameter RAM refers to the four individual 256 byte pages of the four SCCs, resulting in a total of 1KB.
- Page 2-2, Figure 2-2: The figure shows correctly how the channel specific parameters start at the Dual-Port RAM base. Note that on the MPC860, you would normally lose I²C or SPI when using QMC. If you activate the microcode patch from the Freescale web site to relocate I²C/SPI parameters, the channel parameter base will not shift. You will not be able to set up or use channels 0 to 7 or channels 60 to 63.
- Page 2-3, 2.1.2: The note is not fully correct. QMC only uses 0xAC bytes. You can use the microcode patch to relocate I²C/SPI parameters.

Document Conventions

- Page 2-4, 2.1.6: While MCBASE normally points to external RAM, it is permissible to set it up that some or all BDs will actually be placed within free areas of the DPRAM. This may save valuable access time if external memory is slow.
- Page 2-6, Table 2-1: The description for GRFTHR does not recommend a value. You should normally set it to 1 (one) to get an interrupt per frame received.
- Page 2-18, Figure 2-9: Where the interrupt table entry is shown, bit 4 should be included in the description of the channel number field.
- Page 2-22, Table 2-11: Note the description of how to use the POL bit. Be careful to follow it.
- Page 3-1, Figure 3-1: On the MPC8xx, 64 channels are supported. The channel number actually starts at bit 8.
- Page 3-2, 3.2: The STOP RECEIVE command is immediate. It will stop the channel right away and in effect freeze activity on this channel. It will not wait for a frame reception to be finished. As the current buffer descriptor can remain open if a reception was in progress, you can get errors on this buffer once you restart reception again. When restarting, set ZDSTATE first and RSTATE afterwards.
- Page 4-2, Chapter 4: **The NOTE is very important** but not completely correct. After you have handled a table entry you must make sure to clear **everything** but the wrap bit. Do not just clear the interrupt bits or you may confuse your own software with incorrect channel interrupts.
- Page 4-2, 4.1: The fourth paragraph is misleading, stating “the CPM crashes due to an overload of serial data”. The CPM does not crash in the common sense of the word. Rather, it loses track of the mapping of serial data to QMC channels, so it has to cause a global error and halt QMC activity until the user software intervenes. Note that a problem like this typically indicates a very basic performance related design problem. Note well that latency is very important. Overloading of the QMC via the serial line should not occur if you follow the guidelines in chapter 8 about QMC performance.
- Page 4-4, Figure 4-2, Table 4-1. Page 4-5, Figure 4-3: In the MPC8xx family, the bits are not 0 to 7, but 8 to 15. The event information is contained in the lower byte of the SCC event and mask registers.
- Page 4-4, Table 4-1: The IQOV description has a very cryptic last paragraph. This bit should be cleared immediately after reading the SCCE and recognizing this condition by writing a 1 to its location in the SCCE. If this condition occurs, the interrupt last received is lost. It does not overwrite the first entry that is still to be handled by the CPU. The GOV description should refer to an overrun in the SCC’s receiver, as it doesn’t make sense on the transmitter.
- Page 4-5, Figure 4-4: The channel number really starts already at bit 4 for the MPC8xx to reflect the 64 channel capability.
- Page 4-6, Table 4-2: Change the description of bits 4-9 to agree with the changes on the previous page concerning the 64 channel capability.
- Page 5-1, 5: Please note that the CPM “or’s” the various status bits into the BD. You must clear all the status bits generated by the CPM before (re-)enabling the BD or you may confuse your own software with left over old status values.
- Page 5-4, 5.1: Please carefully note this page and the last paragraph to avoid overwritten memory beyond the end of a buffer. The documentation is not quite correct, though. The XTRA information shown in Figure 5-2 is written as 32 bit word. Under special, but quite possible circumstances the XTRA data shown in Figure 5-2 will be written four bytes further than indicated. This means that you **absolutely must** allocate MRBLR+8 bytes for each buffer area for

QMC. Again, you need to allocate **eight** extra bytes rather than only four as described in the manual to be safe. If you do not allocate the extra eight bytes, you risk that these memory area can be overwritten with XTRA info.

- Page 6-12, 6.2: The QMC examples contain a bad error. The TSA is enabled before the SI-RAM has been set up. Ensure the correct order is used in your code.
- Page 9-2, 9.2: The first paragraph is not quite correct. It is not possible to route a single timeslot to all four SCCs as shown in Figure 9-1. A timeslot can only be routed to one SCC by one TSA. To route a single timeslot to two SCCs at the same time, you have to connect both TSAs to the same TDM line and set up separate mappings for this timeslot. This allows you to do subchanneling for at most two channels by using two SCCs. Note that as of this writing, the MSC microcode is available for sale and documentation can be found on the Freescale webpage. Please contact your Freescale representative for more information.

Part III Notes on QMC Interrupt Handling

The recommended method to handle interrupts by the QMC is to first read in the SCCE register. Write back immediately all the bits that you recognize and handle. This clears the respective interrupt events. It is, e.g., incorrect to first handle all the new interrupt table entries and clear GINT afterwards. To avoid deadlocks in your software, you must clear the recognized interrupt bits in the SCCE before actually handling the interrupt entries. Clear only the bits that you handle. Never clear bits for events in the SCCE that you don't handle. This facilitates debugging.

When you get a new GINT event, you should always handle all the new entries in the queue, not just a single one. After handling an entry, make sure that the entry is completely cleared out with the exception of the Wrap bit. Any entry that does not have the Valid bit set must be completely cleared out, including the channel number. If you don't clear out all the entries on startup or after handling them, these bits will confuse your software when the entry is used again. QMC will only OR in new bits and numbers. It will not overwrite and clear bits that were set previously.

Part IV Notes on QMC Performance and SCC Usage

NOTE

The information in this chapter is for your information only to illustrate why the QMC behaves in a certain way. Freescale can not recommend that you base actual designs on this information as there are various subtle things that can go wrong if the QMC is used in "creative" ways. So be careful what you do with the knowledge outlined below.

On a 50 MHz 860, QMC is rated for a total of 4.1 Mbps of traffic to be distributed over 64 channels. As a general hint, QMC is best used when each 8 bit timeslot handled by an SCC represents 64 Kbps. To further reduce peak load, spreading the QMC load over two SCCs is recommended as described in chapter 2.3 of the manual to make use of SCC FIFOs to avoid bursts of data to be handled by the CPM. Note that, due to its larger FIFO size and implementation as described below, SCC1 alone performs better in borderline conditions than both SCC2 and SCC3 would perform in a spreaded setup. So in general, SCC1 would be first choice for QMC, and spreading load over the other SCCs would be second choice.

Document Conventions

With the current implementation of the 860 it is possible, but **definitely not** recommended, to exceed the 4.1Mbps line rate with certain restrictions if, on average, the data rate from the TSA to the SCC does not exceed the upper performance limit for QMC. If you plan to do this anyway, you should definitely contact your Freescale support before proceeding. This can be achieved by having “burst” data to the SCCs or by spreading out single TSA timeslot routing to have silent times for the SCC to get the required average rate. The SCC FIFOs will take up the data bursts and leave the QMC microcode in the CPM time to actually process the data. Note however, that you **can not** simply take the FIFO size and assume that data bursts of that size will be handled correctly. There are two FIFOs per SCC to take into account. The RX FIFO could theoretically be filled completely with a burst of data if enough silence follows to have the microcode process the data in the FIFO. The TX FIFO, however, must be treated very carefully.

The microcode must always be able to fill data, flags, or idles into the TX FIFO so that there is no FIFO underrun when the FIFO contents are put on the line via the TSA. You can consider the maximum QMC bit rate performance to be the possible fill rate of the TX FIFO by the microcode. One additional issue is that there can be a certain delay of >8 CPU clocks from the start of the microcode until the first data can be handled.

If the line rate exceeds the fill rate, you need to be aware of a feature in the QMC called “load balancing”. Certain operations in the QMC, like opening BDs, may take a significant amount of time as can be seen in chapter 8.3 of the manual. These operations will only be started by the microcode if the TX FIFO is filled beyond a certain threshold. If it is not, the microcode will fill idles or flags instead into the FIFO as configured to keep the channel on the line alive. For the 32 byte FIFO in SCC1, this threshold is set at 16 bytes. For the other SCCs with 16 byte FIFO, the threshold is set at 12 bytes. If these thresholds are not exceeded, the QMC will not open TX BDs and you will only see idles/flags on the line.

The QMC works with 32 bit granularity when accessing the FIFO, i.e, it always reads or writes four bytes at a time internally. This granularity is important when thinking about the thresholds. This means that for SCC1 you could theoretically have a burst of 16 bytes length and for the other SCCs a burst of 3 bytes going out to the TSA. Another few bytes would already empty the TX FIFO to the threshold. On the other hand, the microcode will already start to process data after the initial startup while the FIFO data goes. With the microcode start up delay taken into account, this means that, e.g., for a 25MHz CPU and 8Mbps line rate, anything beyond 3 bytes burst on SCC2-4 would not work as expected. On SCC1, there is enough time for the microcode to start data processing in the background for larger bursts. The internal CPM priority handling and execution times mean that you should leave a healthy margin above the thresholds to make sure that the TX side will work and open BDs as you expect it to. Always remember that handling data bursts is not specified in the QMC manual. Freescale does not support this way of operating QMC.

Part V Document Revision History

Table 1. Document Revision History

Rev. No.	Substantive Change(s)
0	Initial release.



THIS PAGE INTENTIONALLY LEFT BLANK



THIS PAGE INTENTIONALLY LEFT BLANK

How to Reach Us:

Home Page:

www.freescale.com

E-mail:

support@freescale.com

USA/Europe or Locations Not Listed:

Freescale Semiconductor
 Technical Information Center, CH370
 1300 N. Alma School Road
 Chandler, Arizona 85224
 +1-800-521-6274 or +1-480-768-2130
support@freescale.com

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
 Technical Information Center
 Schatzbogen 7
 81829 Muenchen, Germany
 +44 1296 380 456 (English)
 +46 8 52200080 (English)
 +49 89 92103 559 (German)
 +33 1 69 35 48 48 (French)
support@freescale.com

Japan:

Freescale Semiconductor Japan Ltd.
 Headquarters
 ARCO Tower 15F
 1-8-1, Shimo-Meguro, Meguro-ku,
 Tokyo 153-0064
 Japan
 0120 191014 or +81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.
 Technical Information Center
 2 Dai King Street
 Tai Po Industrial Estate
 Tai Po, N.T., Hong Kong
 +800 2666 8080
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor Literature Distribution Center
 P.O. Box 5405
 Denver, Colorado 80217
 1-800-441-2447 or 303-675-2140
 Fax: 303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document. Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

