



STANDARD  
MICROSYSTEMS  
CORPORATION

# APPLICATION NOTE 8.19

## USB97C102 Programmers Reference Guide

---

*By Hector Carboney*

80 Arkay Drive  
Hauppauge, NY 11788  
(631) 435-6000  
FAX (631) 273-3123

© 2000 STANDARD MICROSYSTEMS CORPORATION (SMSC)



80 Arkay Drive  
Hauppauge, NY 11788  
(631) 435-6000  
FAX (631) 273-3123

Standard Microsystems is a registered trademark of Standard Microsystems Corporation, and SMSC is a trademark of Standard Microsystems Corporation. Product names and company names are the trademarks of their respective holders. Circuit diagrams utilizing SMSC products are included as a means of illustrating typical applications; consequently complete information sufficient for construction purposes is not necessarily given. Although the information has been checked and is believed to be accurate, no responsibility is assumed for inaccuracies. SMSC reserves the right to make changes to specifications and product descriptions at any time without notice. Contact your local SMSC sales office to obtain the latest specifications before placing your product order. The provision of this information does not convey to the purchaser of the semiconductor devices described any licenses under the patent rights of SMSC or others. All sales are expressly conditional on your agreement to the terms and conditions of the most recently dated version of SMSC's standard Terms of Sale Agreement dated before the date of your order (the "Terms of Sale Agreement"). The product may contain design defects or errors known as anomalies which may cause the product's functions to deviate from published specifications. Anomaly sheets are available upon request. SMSC products are not designed, intended, authorized or warranted for use in any life support or other application where product failure could cause or contribute to personal injury or severe property damage. Any and all such uses without prior written approval of an Officer of SMSC and further testing and/or modification will be fully at the risk of the customer. Copies of this document or other SMSC literature, as well as the Terms of Sale Agreement, may be obtained by visiting SMSC's website at <http://www.smsc.com>.

SMSC DISCLAIMS AND EXCLUDES ANY AND ALL WARRANTIES, INCLUDING WITHOUT LIMITATION ANY AND ALL IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND AGAINST INFRINGEMENT, AND ANY AND ALL WARRANTIES ARISING FROM ANY COURSE OF DEALING OR USAGE OF TRADE.

IN NO EVENT SHALL SMSC BE LIABLE FOR ANY DIRECT, INCIDENTAL, INDIRECT, SPECIAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES; OR FOR LOST DATA, PROFITS, SAVINGS OR REVENUES OF ANY KIND; REGARDLESS OF THE FORM OF ACTION, WHETHER BASED ON CONTRACT; TORT; NEGLIGENCE OF SMSC OR OTHERS; STRICT LIABILITY; BREACH OF WARRANTY; OR OTHERWISE; WHETHER OR NOT ANY REMEDY IS HELD TO HAVE FAILED OF ITS ESSENTIAL PURPOSE, AND WHETHER OR NOT SMSC HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

# TABLE OF CONTENTS

REFERENCES.....	5
ACRONYMS .....	5
<b>CHAPTER 1 - INTRODUCTION .....</b>	<b>6</b>
INTENDED AUDIENCE .....	6
<b>CHAPTER 2 – BACKGROUND INFORMATION.....</b>	<b>7</b>
Why C? .....	7
C on the 8051.....	7
MCU Address Spaces .....	9
Interrupt Service Routines (ISRs).....	11
Performance Perspective and Strategy.....	11
Example Program .....	13
Development Equipment.....	13
Executing and/or re-Building the Examples .....	15
The Example Programs .....	16
Coding Style: .....	16
Common Files .....	17
<b>CHAPTER 3 - FIRMWARE BASICS .....</b>	<b>18</b>
Initialization.....	18
CLOCK_SEL Register.....	18
UTIL_CONFIG Register .....	19
GPIOA_DIR Register .....	19
GPIOA_OUT Register .....	19
MEM_BANK Register.....	19
MEM_BANK2 Register.....	20
IOBASE Register .....	20
MEMBASE Register.....	20
<i>Siolnit()</i> .....	20
<i>Platform_Display()</i> .....	21
<i>DBGPRINT()</i> .....	21
<i>printf()</i> .....	21
<i>Com2SendByte()</i> .....	22
<i>kbhit()</i> .....	22
<i>_getkey()</i> .....	22
<i>Delay1uS():</i> .....	22
Code Example: Chpt3.Hex.....	22
<b>CHAPTER 4 - THE MMU.....</b>	<b>23</b>
Allocating and Freeing Packet Memories with the MCU.....	23
USB Reception .....	23
USB Transmission .....	26
<i>Flushing a TxFIFO:</i> .....	28
Memory Management Policy (MMP).....	28
<i>MMPCMD Register:</i> .....	28
<i>IN / OUT Nak Registers:</i> .....	29
<i>PAGS_FREE Register:</i> .....	29
GP_FIFO's:.....	30
<b>CHAPTER 5 - DMA.....</b>	<b>31</b>
Shared Bus Architecture Details: .....	31
DMA Channels .....	32

DMA Transfer Modes .....	32
DMA Transfer Types .....	32
<b>BUS_REQ_INH_TCX:.....</b>	<b>33</b>
<b>DMA_STS AND BUS_STAT_CHXTC: .....</b>	<b>33</b>
<b>DMAC ADDRESS SPACE:.....</b>	<b>33</b>
<b>SAMPLE CODE: CHPT5.HEX.....</b>	<b>34</b>
DMA Direct From MMU to LPT:.....	34
DMA From ISA RAM to LPT: .....	35
DMA From MMU to ISA RAM:.....	36
SCATTER / GATHER DMA .....	36
<i>MMU Memory to ISA Memory Transfers</i> .....	36
<i>ISA Memory to MMU Memory Transfers</i> .....	36
<b>CHAPTER 6 - GETTING ON THE BUS .....</b>	<b>41</b>
Device States: .....	41
USB Suspend/Resume: .....	41
USB Remote Wakeup .....	42
EPCTRL Registers .....	43
<i>Non-ISO OUT EP's</i> .....	43
<i>Non-ISO IN EP's:</i> .....	43
<i>ISO EP's:</i> .....	43
Endpoint Command Register:.....	44
NonControl Endpoint Register: .....	45
RESETS .....	45
<i>POR</i> .....	45
<i>USB Reset</i> .....	46
EP0 Control Transfers .....	46
<i>EP0 Stalls</i> .....	47
<i>EP0 FSM</i> .....	47
<i>CTL-READ Transfers</i> .....	48
<i>No-Data Control Transfers</i> .....	48
<i>Control-Write Transfers</i> .....	48
<i>EP0 Application Callbacks</i> .....	49
APPLICATION POLLING .....	50
DATA TRANSFER.....	50
<i>Data Transfer Overview</i> .....	50
<i>Data Transfer Details</i> .....	50
<i>GPIO Summary</i> .....	51
USB HUB BLOCK.....	51
<i>HUB Initialization</i> .....	52
<i>HUB Control Register1</i> .....	52
<i>USB97C100 Compatibility (HUB BYPASS 2) Mode</i> .....	53

## REFERENCES

Universal Serial Bus Specification Revision 1.1 September 23, 1998 as well as various Device Class Specifications and White Papers are all available at <http://www.usb.org>. This site also has links to many other useful sites.

SMSC USB97C102 and FDC37C67x Data sheets, both available at <http://www.smsc.com>. The 67x is the SIO device on the EVB97C102.

8051 Data Sheet, available from numerous manufacturers, including Intel Corporation, the originator. This is the MCU inside the USB97C102.

8237 Data Sheet, Intel Corporation. This is the DMAC inside the USB97C102.

SMSC LAN91C94/5/6 Data Sheets, available at <http://www.smsc.com>. These devices have a very similar MMU to the USB97C102, with explanation.

The C Programming Language, Brian W. Kernighan and Dennis M. Ritchie, Prentice Hall; can be useful for those having little experience with the language.

## ACRONYMS

BER	--	Bit Error Rate
CFES	--	CLEAR_FEATURE (ENDPOINT_STALL), a USB command
COM	--	For the purposes of this document, an RS-232 serial COMmunications port
DMA	--	Direct Memory Access
DMAC	--	Direct Memory Access Controller
EP	--	EndPoint, as in a USB EP
EP0	--	Endpoint zero, the default pipe for a USB device
GPIO	--	General Purpose Input/Out, as in a device pin
IRQ	--	Interrupt Request
ISO	--	For the purposes of this document, ISOchronous, as in a USB ISO transfer, not International Standards Organization
ISR	--	Interrupt Service Routine
LPT	--	For the purposes of this document, a PC legacy parallel Line Printer port
MCU	--	Microcontroller Unit
MIPs	--	Million Instructions per second
MMP	--	Memory Management Policy
MMU	--	Memory Management Unit
MPU	--	MicroProcessor Unit
POR	--	Power-On Reset
RTL	--	Run Time Library
RxEP	--	Receive EndPoint, as in a USB RxEP
SFES	--	SET_FEATURE (ENDPOINT_STALL), a USB command
SGDMA	--	Scatter-Gather DMA
SIE	--	Serial Interface Engine, as in a USB SIE
SIO	--	Super Input/Output
SOF	--	Start Of Frame, as in a USB frame
TxEP	--	Transmit EndPoint, as in a USB TxEP
USB	--	Universal Serial Bus

## CHAPTER 1 - INTRODUCTION

The basic architectural concept of the USB97C102 device is that all high bandwidth data flow be handled entirely in hardware, with an integrated MCU (MicroController Unit, an 8051 derivative) acting only to manage the flow of data through the various hardware engines. At the center of the device is a multi-ported MMU (Memory Management Unit) that dynamically allocates and frees memory pages grouped into virtual packets both automatically in response to USB traffic, and also under software control by the MCU. On one port of the MMU is the SIE (Serial Interface Engine) that provides a fully hardware-driven interface to the USB, while another port on the MMU is connected to a partial ISA bus interface containing a DMAC (Direct Memory Access Controller, an enhanced 8237 type) that provides a fully hardware-driven interface to external peripheral devices. All three (3) hardware engines (SIE/MMU/DMAC) are capable of operating concurrently with each other and with the MCU, while the DMAC is also capable of interleaving transfers to/from multiple devices concurrently.

As part of its function, the firmware must establish flow control in order to prevent overrun from either the USB or the DMAC in the event that the target of the transfer is not as fast as the data source. Also, while the SIE takes care of all bit and packet level USB protocol issues in hardware, the higher levels of the USB device protocol (e.g., Default Pipe traffic) are the responsibility of the firmware. Note that since this traffic is not bandwidth intensive, the firmware implementation results in absolute flexibility combined with reduced cost, without adversely affecting performance.

This document, along with the companion code which can be downloaded from SMSC's website (Download Code Now), describes detailed register level programming considerations, along with working examples, of the device architecture in general, and USB applications in specific.

### INTENDED AUDIENCE

The intended audience of this document is primarily software engineers, and perhaps their managers, involved with writing firmware for the USB97C102 device. Hardware and systems engineers can also benefit from at least skimming this document in order to better understand the effect of their design decisions on the firmware component of the system.

The reader is assumed to be experienced in register level hardware programming, either in an embedded or Device Driver context. A basic knowledge of the C programming language is required, since all of the example code is written in C, as is an understanding of basic USB operation. Prior knowledge of the 8051 MCU and 8237 DMAC is desirable, but is not a prerequisite.

## CHAPTER 2 – BACKGROUND INFORMATION

This chapter is a potpourri, covering a range of topics including:

- A primer about writing C code for the MCU, including a discussion and example program illustrating the performance impact of coding style and address space usage, as well as a variety of Interrupt handler considerations.
- A description of the development and build environments, including how to set them up and use them to execute and rebuild the example programs.
- A description of the example programs, including coding conventions and common files

An understanding of these topics is assumed in the remainder of this document.

### Why C?

The examples in this document use the C language, as opposed to either ASM or C++ for a variety of reasons. The examples are more clear because they are less cluttered with detail than they would be if they were written in ASM. The quality of C compilers for the MCU has improved to the point that there is now less of a size and performance penalty for using C, rather than ASM, than there once was. When combined with increased time to market pressure, this makes C a reasonable choice for implementing an actual product. As a result of this, most MCU programmers at least have some familiarity with the language, again making it a reasonable choice for a document of this type. For embedded developers with an aversion to C, it might help to think of the compiler as an automated (if sometimes not very smart) ASM generator (just turn on the ASM listing option), which actually is not a bad way to think of the compiler anyway when the target is a small-embedded MCU.

Without getting into a debate over the use of C++ in embedded applications, the language does not seem to have had as much penetration as C at the present time, especially for low-end MCUs. Maybe some day C++ will be the language of choice for a document of this type, but not today.

One non-reason for picking C is portability. The nature of C on the MCU virtually requires the extensive use of non-standard, and hence non-portable, extensions to the ANSI specification of the language. This makes writing C code that is even portable to different compilers for the same MCU an exhaustive exercise in header files. Also, the very purpose of this document is to illustrate coding techniques that are specific to a particular hardware architecture so, even if the source code could be recompiled for a different target MCU, the code still would not do anything useful without the rest of the device hardware being present. Having said this, the example programs are written in such a way as to reduce the porting burden to the extent possible (e.g., non-ANSI typedef's are in header files, no `///  
in-line comments are used, etc.).`

### C on the 8051

This section makes extensive reference to the detailed operation of the 8051; readers unfamiliar with this architecture are referred to the data sheet for a complete description. Developers with substantial experience writing C for embedded MCUs will likely be familiar with this material, and so this section is targeted at that portion of readers who come from a PC device driver background. But do not worry -- later sections will discuss the 8237 DMAC in detail, at which point the tables will be turned. For the developers with a PC background, the closest PC equivalent to programming this MCU is the old MS-DOS COM programs in which all code and data fit in a single 64 KB block (i.e., the TINY memory model) with fixed offsets to the beginning, etc. If you wanted to touch or change any memory or I/O port, you could have your way with it without having to ask any O/S for any approval or assistance. You were writing for maybe a 1 MIPs machine, and not a modern PC with hundred of MIPs and nearly as many MB of RAM, and a big fat OS between you and the hardware. If you are old enough to remember those days, then welcome home. If you are young enough that you cannot believe that people used to really do that for a living, then read on...

When C is compiled for MPUs, arguments are usually passed on the stack, and automatic variables are allocated on the stack. For MPUs with a stack of substantial size, and with the capability to efficiently address the stack memory, this makes a good deal of sense. However, the MCU has neither of these capabilities. As a result, parameter passing is usually done through registers, or in fixed memory locations, or on a "simulated" stack if the arguments cannot fit in registers. Use of a simulated stack is very slow, so it is suggested that any functions for which speed is a consideration limit their arguments to the types and numbers that can be passed in registers. Even then every extra argument passed increases the overhead of the call, so "less is more" when it comes to performance. One way of satisfying this suggestion would be to pass a single pointer to a structure, but that is not a good solution for the MCU because it does not have the flexible addressing modes (e.g., based indexed

displaced addressing) that CISC MPUs usually have. The MCU is also not very efficient at doing the address calculation necessary to access a complex data structure. Although it is generally frowned upon (and for good reasons), the use of global variables is the best solution from the perspective of performance.

In addition, short functions are good candidates for implementation as macros; as macros, there is no call and return overhead, which can be a substantial portion of the total execution time when functions are short. Also, the final code size might not even be larger because all of the code related to parameter passing and register saving/restoring is eliminated. Using macros instead of small functions is a good thing for the MCU. Another benefit of macro functions is that they avoid the issue of reentrancy, which is discussed next.

Automatic variables are usually not placed on the MCU stack both because the MCU would not be able to efficiently access them if they were, and because the stack tends to be small. Instead, automatic variables are usually allocated in one of the data address spaces of the MCU, which are described in more detail in a later section. In order to avoid wasting valuable data memory for automatic variables that are not presently being used, a good Linker/Locator will use data overlays within the procedures of a given thread. Since different threads (e.g., foreground and ISR) can execute concurrently, their data cannot be overlaid. This brings up the issue of recursive and reentrant code. While it is possible to compile code for the MCU so that it is reentrant (i.e., executed by more than one thread at a time), such code cannot have its data overlaid with the data from any other thread, so it consumes more memory. Also, its data must be relocatable because it might be necessary to have multiple instances at the same time. Due to a lack of efficient addressing mechanisms for this type of allocation in the MCU, such code will execute much more slowly than standard functions. The situation is essentially the same for recursive functions, which call themselves within the same thread, because of the need for multiple simultaneous data instances. In general, both are undesirable from a performance point of view. Avoiding recursion mostly involves algorithm design, so there is not much that can be presented here with respect to universal techniques. However, there are two easy ways to avoid reentrancy: one is to use macros, as was described above, and the other is to cut-and-paste the function and to give each copy a slightly different name. This can work very well if the reentrancy is between two different threads, for example the foreground thread and a single ISR. Having duplicate functions may increase code size somewhat, but each copy is smaller and executes much faster than a reentrant version of the same function, so the increase in code size, if any, is often worth it.

In order for data overlays to work properly, the Linker/Locator must be able to unambiguously determine the execution context for every code section, which can be difficult if the code calls through function pointers. The easiest solution to this is not to use function pointers, but if they must be used, then the compiler should be set to produce an ASM listing, and this file and the Linker output files should be carefully inspected to be sure that the tools are correctly implementing the software design as intended.

Note that the Linker/Locator will often mistake uncalled functions for separate threads and will not overlay their data with any of the actual threads. This behavior can be somewhat annoying during initial coding, in which case it is common to write a number of functions before actually using any of them. One solution, albeit not a pretty one, is to create a dummy function that calls each of the uncalled functions with dummy arguments conditional on an argument to the function being TRUE; this dummy function is then called from e.g., main() with a binary flag that is FALSE, so the code never actually executes, but it gives the Linker enough information to know how to overlay the data; see Chapter 3 and Chapter 6 for an example of this. Some compilers also permit defining the thread each function is supposed to execute in, which achieves the same result. The only problem with this approach is that each function's source code needs to be changed in order to move it from one thread to another.

Considering the combination of the above issues, it is usually a good idea to develop a function hierarchy that is relatively flat (i.e., not too deeply nested) and orderly, without a lot of cross calling. This combination reduces the impact of call/return/parameter passing overhead and gives the linker good opportunity for achieving RAM savings through the data overlay mechanism.

The MCU does not contain a barrel shifter; as a result of this, shifts are implemented one bit at a time in a loop, so they are best avoided where possible. However, a good compiler will recognize shifts by a multiple of 8 as being a change in address, and this can be handled quite efficiently, especially if the argument is in DATA space. Examples of this would include macro's for HIBYTE(a), LOBYTE(a) and MAKEWORD(a,b) (see Type.H for the definitions of these macro's).

Another efficient sequence involves a conditional jump based on a single bit in a byte being either set or cleared -- the MCU has JB (Jump if bit set) and JNB (jump if bit clear) instructions, and a good compiler will use these whenever it gets a chance. For example,



**if (myVar & 0x08) { }.**

Most compilers offer the option of producing an ASM listing in addition to the C listing, and it is often useful to enable this feature and at least glance at what the compiler is doing, especially when working in a new environment. Sometimes small changes in the source code can make big changes in the resulting object code. The linker output file should also be inspected, especially to make sure that the linker is correctly understanding the memory map (e.g., 256 bytes of internal RAM, as opposed to some other size), that the stack is of sufficient size, and that any external data RAM or firmware ROM is properly located. It is also the author's personal preference to set all warning levels to maximum and to take any warnings seriously, but it is recognized that this is a matter of individual taste.

Speaking of ASM, if it happens that there is some function for which execution speed is absolutely critical, it is always possible to write some of the code in ASM and call these ASM functions from C. Since the quality of the code produced by compilers has improved much in recent years, this should not be necessary, and the performance gain will likely be small if it is done, but it is always an option.

### **MCU Address Spaces**

The MCU contains 256 bytes of internal RAM. The low 128 bytes of this address space can be accessed using either direct or register indirect (e.g., @R0, @R1) addressing, while the high 128 bytes can only be accessed using indirect addressing. Since direct addressing is somewhat faster than indirect addressing, it is desirable to locate variables whose access speed is important in the low 128 bytes. A portion (16 bytes) of the low 128 bytes can also be addressed as individual bits and, unlike most MPUs, the MCU contains a special Boolean Processor and so is quite efficient at manipulating these. The MCU also has an external data memory address space of 64 KB. However, access to this address space is substantially slower than either direct or indirect addressing modes because it involves the use of the 16-bit DPTR register.

[For completeness, the upper 128 bytes of internal memory, when used with direct addressing, accesses yet another address space: the Special Function Register, SFR, and there is also an external address space mechanism that only uses an 8-bit address with paging. However, neither of these is salient to the discussion, which follows.]

C compilers for the MCU are quite flexible in their ability to enable the programmer to define where variables are stored and how they are accessed based on how they are declared in the source code. Since ANSI C has no provision for this, compilers implement extensions to the ANSI language, which limits portability between different compilers, even for the same MCU. For the Keil compiler, variables located in the low 128 bytes of internal RAM are referred to as type "data" and direct addressing is used to access them. Variables located anywhere in the 256 byte internal RAM are referred to as type "idata" and indirect addressing is required to access them, even if they are ultimately located in the low 128 bytes, since their actual location is not known at compile time. Variables located in the external data address space are referred to as type "xdata", and the DPTR register is always used to access them. Bit addressable variables are referred to as type "bit" for obvious reasons.

If all of this seems a little bit bizarre, well it is. If all of this seems a little confusing, do not worry about it. The C language makes all of this easy to hide in header files (for the most part), which are included in the example programs.

The on-chip peripherals in the USCB97C102 (e.g., the SIE and MMU) are mapped into the xdata address space. In addition, there are 4 addressing windows, all of which are mapped into the xdata address space of the MCU, that permit access to the address spaces of the external buses. There is a window for the ISA I/O address space, one for the ISA Memory address space and two for the "Flash" Bus address space. [As an aside, note that the name "Flash" Bus is something of a misnomer, in that the bus can be used to interface to any combination of memory devices (ROM/EPROM/EEPROM/FLASH/RAM, etc.) and/or memory-mapped peripheral devices; a more descriptive name would be the MCU Bus, since it is a generic bus that is owned full time by the MCU.] Each memory window has its own bank select register associated with it, which is also mapped into xdata space that permits moving the window anywhere in the entire address space of the corresponding bus. The following table summarizes the xdata address windows:

ADDRESS SPACE (BUS & TYPE)	TOTAL SIZE	WINDOW SIZE	WINDOW LOCATION (XDATA)	BANK SELECT REGISTER (XDATA)
ISA I/O	64 KB	256 bytes	0x4000-0x40ff	IOBASE [0x7F71]
ISA Memory	1 MB	4 KB	0x5000-0x5fff	MEMBASE [0x7F72]
Flash Memory	1 MB	16 KB	0xc000-0xffff	MEM_BANK [0x7F29]
		16 KB	0x8000-0xBfff	MEM_BANK2 [0x7F28]

The MEM\_BANK and MEM\_BANK2 registers also control paging in the code address space of the MCU, and this must be considered if the entire firmware is larger than 16 KB.

In particular, the bottom 16KB of code address space (0x0000-0x3FFF) always maps straight through to the Flash Bus (i.e., 0x00000-0x03FFF). The next 16 KB of code space (0x4000-0x7FFF) is a movable 16 KB window whose 6 MSB's are controlled by the MEM\_BANK register. Out of the upper 32 KB of both the code and xdata spaces, the bottom 16KB (0x8000-0xBFFF) is a movable 16KB window whose 6 MSB's are controlled by MEM\_BANK2 register and the upper 16KB (0xC000-0xFFFF) is a duplicate (i.e., an alias) of the lower -upper 16KB (0x4000-0x7FFF) of code space, which is controlled by MEM\_BANK register. The following table summarizes Flash Bus mapping into the MCU address spaces:

ADDRESS (SIZE)	CODE	XDATA
0xC000-0xFFFF (16 KB)	Set by MEM_BANK	Set by MEM_BANK
0x8000-0xBFFF (16 KB)	Set by MEM_BANK2	Set by MEM_BANK2
0x4000-0x7FFF (16 KB)	Set by MEM_BANK	(On-chip peripherals & ISA)
0x0000-0x3FFF (16 KB)	0x0000-0x3FFF	(On-chip SFR's, etc)

Unlike the other on-chip peripherals, the on-chip DMAC is mapped into the ISA I/O address space, so its registers are accessed in the MCU xdata address space, just like the SIE and MMU, once the IOBASE register is set.

From the discussion above, it might seem undesirable that the on-chip peripherals are all mapped into the xdata address space, and it is when considering the performance penalty involved. However, the other address spaces are very limited in size, and it would be even worse to lose a substantial portion of one of those address spaces instead (especially data or idata). It is important for the device programmer to recognize this situation because it has an impact on firmware performance. Firmware should be written so as to access device registers as little as possible. For example, if the value contained in a register is needed multiple times, then it should be read once and saved in a variable, and then this variable can be read multiple times with far greater speed.

The existence of all of these different address spaces raises an interesting question: what happens with pointers? The answer is: it depends. Each address space (data, idata, xdata and code) needs either a 1 or 2 byte pointer to span it. As a result, if the particular address space that a pointer references is known or implied by usage (e.g., perhaps it is explicitly typed when it is defined), then a straight pointer is all that is needed. However, if a pointer is to be "generic" in that the same pointer can be used to reference any address space, then an extra byte must be added to it in order to identify the address space to which the pointer is currently assigned -- hence, the 3-byte pointer. Some compilers provide support for 3-byte pointers, some only support 3-byte pointers, and some do not support 3-byte pointers at all, comprising yet another portability issue. It is the author's opinion that generic pointers are a mixed blessing: they are terrible from a performance point of view, but they do permit writing code with fewer non-ANSI directives in it, and they make it trivial to move data items from one address space to another because none of the pointers to the data need to be changed when the data is moved (by re-typedef'ing it).

## Interrupt Service Routines (ISRs)

One very nice feature of the MCU is that it contains multiple Register banks. The current register bank can be changed quickly by setting 2 bits in the PSW. This can be used to substantially reduce the interrupt latency time by avoiding having to save the entire register bank on entry to the ISR (and restoring it on RETI). Compilers for the MCU support this feature by (surprise) using extensions to ANSI C. For the Keil compiler, the "using" function attribute causes the compiler to insert code in the function that will save the existing register bank and switch at entry, and will restore the previous register bank at return.

In the USB97C102 architecture, the most important ISR is for IRQ 0, which is the one that handles USB traffic. It is desirable to keep the latency for servicing this IRQ as short as possible, and the use of Register Bank switching is strongly encouraged. Note that each additional register bank used is located in the "data" address space, so some RAM is lost this way, but the alternative would be to push the registers on the Stack (usually in "data" space), so the total RAM consumed is the same, and the bank switch is much faster.

One nice feature about the USB97C102 when it comes to the handling of the **ISR** registers, which contain the Interrupt Status bit(s), is that -- **these registers are cleared by writing a "1" to the corresponding bit(s)!** This feature eliminates the problem of having IRQs cleared by accident (bits cleared on read), which existed in the older version of this part. In the past, it was not possible to read the registers without clearing the IRQs. The ISR bit(s) were automatically cleared each time the associated register was read. As a result, each time the register was read, all pending interrupts had to be serviced before continuing normal operation. Adding this feature to the USB97C102 allows the designer to individually clear the associated bit(s) in the associated "Interrupt Source" registers as the corresponding interrupts are handled.

As in any interrupt-driven system, care must be exercised if any hardware or software resources are shared between the foreground and background. The classic situation to avoid is the one in which, while the foreground code is in the middle of a read-modify-write, the ISR executes and changes the value, and then the foreground over-writes the value from the ISR. An example of such a case in this device is the **GPIOA\_OUT** register in the situation when both the foreground and ISR threads are manipulating GPIO pins, but the same class of situation can result as a matter of sharing software resources (e.g., RAM variables) rather than hardware registers. One way to avoid this problem is for the foreground thread to disable IRQs while accessing the shared resource, but other mechanisms are possible as well. If IRQs are disabled by the foreground thread, then it should be for the shortest amount of time possible (ideally just a couple of microseconds) in order to avoid a significant negative impact on the IRQ latency time.

In other situations, a portion of the hardware is shared between the foreground and ISR threads, but there is no read-modify-write issue; a common example of this is a numeric coprocessor in an MPU system. In these cases, the ISR can simply save/restore the register set so that the foreground thread does not even know the ISR had used the hardware. An example of such a situation in the USB97C102 is the MMU, with the **PNR**, **PRL** and **PRH** registers.

Also as usual, any variables used by the ISR need to be global in one way or another and, in order to avoid problems with reentrancy, do not call any of the ISR code from the foreground. Of course, the ISR itself can never be called from any foreground thread because it uses a RETI, rather than a RET.

## Performance Perspective and Strategy

Although much of the previous discussion was concerned with the impact of specific coding techniques on performance, it is useful to realize that much of the code in any given application for this component is not related to performance in any way. For example, when a USB device is first attached to the Bus, it is enumerated, reset, configured, etc. Required execution times for these operations (per the USB Specification) is measured in units of milliseconds, so performance is not a concern here. There are also occasional Control Transfers to the Default Pipe (EP0), but these are also not performance sensitive since they do not occur with high frequency. When writing code of this type, it is permissible to use any and all of the techniques (e.g., complex data structures, etc.) that might have an adverse effect on performance if using those techniques is appropriate. For example, the USB descriptors are essentially a complex data structure, so writing them that way is a natural expression of the coding solution.

As far as the remaining code is concerned, it is helpful to quantify just what performance level is required or desired. For example, Isochronous applications are essentially hard real time -- it is an absolute requirement that the software be fast enough to handle the stream, but any further speed improvement serves no useful purpose. In the case of Bulk data applications, once the software is fast enough to saturate either the USB or the peripheral

device, the same situation results. From the software perspective, performance in this component architecture is really an issue of how many packets are handled in each USB frame. Note that this is different from the situation in many other USB components in which the actual data bytes must flow through the MCU. **Based on this definition, a USB camera that delivers 1,000 byte Isochronous packets is a low-performance application with respect to software since only a single packet needs to be handled each USB frame.**

Once specific performance numbers are established, the data flow needs to be planned:

Generally, USB OUT packets will arrive in an `Isr0()` function, which is the ISR that handles interrupts from USB traffic flow (any RX, any TX, etc.). At arrival, the Rx packets need to be validated and saved for consumption by the foreground device handlers, either in software or hardware queues (the component has both). This is also the time to update the current packet count for each BULK EP and, if an EP has reached its limit, to make that EP "busy" so that further OUT packets to that EP will be NAKd. In the foreground, each device handler checks the state of its peripheral device. If the peripheral device has just completed transferring a packet, then the handler must free the packet in the MMU so that the packet memory can be used for receiving additional packets. If the EP is Bulk, then the handler must update the packet count and, if it is low enough, make the EP "not busy" so that future OUT packets sent to that EP will be received and ACKd. Finally, if the peripheral device is ready for another packet and one is available in its packet queue, the handler must start sending the next packet to the peripheral device. For high performance peripheral devices, a DMA hardware interface should be used (rather than PIO), so starting the next packet involves setting up a new DMA session.

The data flow for Tx is similar to Rx, but backwards. If the peripheral device is below its limit on the number of packets it is permitted, and the peripheral device is ready to fill a new packet, then the handler must allocate a packet and start the peripheral device filling it. If the peripheral device has finished filling a packet and is below its limit on how many it is allowed to queue for transmission on the USB, then the handler must queue the full packet. In both cases, the count of packets owned and packets queued must be adjusted. After the Host reads the packet on the USB, the packet will appear in the Tx Completion queue. It is a matter of choice by the programmer whether to handle this in `Isr0()` or in the foreground but, regardless of wherever it is done, the packet must be removed from the completion fifo, and the count of packets owned and packets queued must be updated.

[To be truthful, the word "must" in the paragraphs above is a little bit strong, since simplifications are often possible, but the description represents the most general case.]

From the above, it can be seen that the MMU acts in a managerial role supervising the traffic flow and ownership of shared finite resources such as packet memories and Tx queues. There are at least two obvious ways that code like this can be implemented: either centralized or distributed. In the centralized case, there is a set of functions that encapsulate data structure(s) that maintain the present state of the system, and decide when to busy/un-busy RX EPs and when to grant or refuse requests for packet allocations and tx packet queuing. In the distributed case, each individual EP handler function maintains its own state and decides when to allocate, free and queue packets, etc. As usual with software, still other approaches are possible, and any approach that yields a correct solution is equally valid.

Since all of the above code executes for every packet that is transferred on the USB, the execution speed of this code is critical if high performance, defined in terms of the number of packets per USB frame handled, is to be achieved. From the previous discussion, a number of techniques can be applied in order to obtain best performance:

- locate all variables in data or `idata` address space
- implement any short functions as macros
- bank switch the registers in the ISR
- keep the function hierarchy shallow
- consider global variables -- they are necessary for the ISR anyway
- if any arguments are passed to functions, pass them in registers
- avoid complex data structures -- use simple arrays or scalars
- avoid 3-byte pointers if possible

Using all of these techniques in combination will make a big difference in the execution speed of the code, and can still result in very legible and maintainable code if thoughtfully applied. If at that point the performance is still less than desired, it becomes important to understand clearly where the MCU is spending its time. An easy way to determine this is by pulsing GPIO pins at the entry/exit of the major functions. Pulsing GPIO pins in this fashion

does add a couple of microseconds to the execution time, but that is relatively small, and the information it provides is critical to understanding where the MCU is spending its time. The functions that are consuming the most time can then be inspected in the ASM listing of the compiler in order to understand why they are taking so long, and to see what, if anything, can be done to reduce their execution times. Sometimes recoding a function in ASM can help, but modern compilers generate reasonably efficient code, so the improvement is usually not much. Any substantial improvement usually comes from changes in algorithms, data structures, or the address space in which variables are located, which is the reason why it is so important to think all of this through carefully before writing the code in the first place.

### Example Program

The example program **Chpt2.Hex** illustrates the effects of address spaces and coding style on performance. Unlike the other examples in this document, this program is \*NOT\* meant to be executed. Instead, the ASM listing should be inspected.

The program consists of a set of 4 functions that each pushes an entry onto the head of a software queue. The functions differ in the address space and organization of the queue. By inspecting the ASM listing, quantitative performance differences can be determined. The following analysis is for the typical case in which the queue is neither full, nor does the head pointer wrap around.

**PushDqueHead()** uses simple scalars and an array in data space. The function executes 15 instructions that consume 20 processor cycles.

**PushIqueHead()** is the same, except that the variables are located in idata space. The function executes 17 instructions in 22 processor cycles. By inspecting the ASM code, it can be seen that the extra instructions and cycles are a result of the head pointer being located in idata space, not the array. If the head pointer were in data space, then the instruction and cycle count would be the same as the previous example. From this it can be concluded that:

1. idata access is only a little slower than data access.
2. idata access for array elements is identical in speed to data access, so arrays should usually be located in idata space.

**PushIQueHead()** uses a data structure located in idata space. The function executes 24 instructions in 31 processor cycles. From this it can be concluded that there is a major performance impact involved in the use of data structures (in this case 50%), even fairly simple ones. Inspection of the ASM code reveals that the extra time is spent doing the address arithmetic to access the structure elements.

**PushXqueHead()** uses the same data structure as above, but it is located in xdata space. The function executes 39 instructions in 64 processor cycles, making it about twice as slow as the same example in idata space, and 3 times slower than the first 2 examples. From this it can be concluded that there is a major performance penalty associated with xdata access. Inspection of all of the DPTR manipulation in the ASM listing shows why this is the case.

### Development Equipment

It is not the purpose of this document to specifically recommend or endorse any particular product(s) of any particular manufacturer(s). For each of the items described below, a variety of manufacturers offers a range of products that appear to be suitable. However, in order to provide concrete examples with explicit instructions, it is necessary to do so in the context of a particular hardware and software environment. Following is a list of the environment used to develop the example programs in this document:

#### 1. USB Host System

A standard PC running Windows 98SE and/or Win2000. The extra Host software consists of RW2.Exe and UsbSmsc.Sys, both supplied on the companion code disk. Since this system will often be used for the purpose of Driver and/or Application software testing, it is the author's preference to treat this as a test machine -- it is assumed that it could crash at a moment's notice, losing everything on its hard disk and disrupting any LAN that it might be connected to. As a result, no development work should be done on it, no important data should be stored on it, and it should not be connected to any network. In addition, nothing but Host test software should be executed on it in order to avoid corrupting any test results due to interactions with other hardware and/or software, except as an explicit part of the testing.

## 2. MCU Development System

A standard PC running Windows 98SE, Win2000, or any other OS capable of hosting all of the development tools used. This system hosts the MCU Compiler, USB Protocol Analyzer, and ROM Emulator, each of which is described below. Ideally, everything on this system should be full production quality, with no Beta or pre-Release anything. Connection to a network for backup and/or printing services is encouraged. Any Host software development (either drivers and/or applications) can be performed on this machine, provided that all tools used are of suitable production status. Of course, any such Host software should never be executed on this [development] system.

Because one of the particular Analyzers and the Emulator used each requires a parallel port, a board containing a second parallel port (set as ECP with Legacy LPT2 assignments) was added to the system. The Emulator can also make use of a serial port for various purposes, and it was connected to COM1, but this connection is not used in the context of this document.

Since a number of the programs make use of the COM2 port on the 67x SIO device on the EVB97C102 for console I/O, a null modem was used to connect the COM2 port on the EVB to the COM2 port on the MCU Development System. HyperTerminal was used to establish the communications link, with settings of 115.2 Kbaud, 8 data, 1 stop, no parity, no flow control, TTY emulation. Because no hardware flow control is used, the null modem can be of the simple 3-wire type, with the TxD and RxD crossed, and the grounds connected together. For this to function properly, the jumpers on the EVB must be set to connect the COM2 transceivers to the SIO, rather than the USB97C102, device. For Assy 6126 Rev. B, JP9 and JP10 must have jumpers between pins 1 and 2.

## 3. MCU Development Tools

The Keil 8051 C Compiler V5.0 ([www.keil.com](http://www.keil.com)). For installation, just use the default settings.

## 4. USB Protocol Analyzer

When executing the programs in this document, access to an Analyzer is not required, since the software has already been developed and tested. However, when developing new software, at least part time access to an Analyzer is a practical necessity, at least in the author's opinion. An external trigger input is desirable in order to permit triggering the Analyzer from software-controlled pulses on either the Host or the Target system. A trigger output is desirable to permit triggering an oscilloscope or Logic Analyzer.

Representative manufacturers of Analyzers include Genoa Technology [[www.gentech.com](http://www.gentech.com)] and CATC [[www.catc.com](http://www.catc.com)]. When using the CATC USB Inspector, the companion software searches the LPT ports for the Analyzer, so it is best to install the device on LPT1 so that the software does not touch the Emulator while searching for the Analyzer.

## 5. ROM Emulator

The TechTools UniROM UR08-1M-90 ROM Emulator ([www.tech-tools.com](http://www.tech-tools.com)). This device was connected to LPT2 during the code development for this document.

Connecting UniROM to the EVB97C102 consists of plugging the 32-pin target cable into the Flash Dip Socket (U14) on the EVB97C102. Plug the 34-pin header end of the cable into the UniROM. This connector is polarized, making it nearly impossible to plug in backwards.

If the ROM Emulator has a suitable Reset output with which to drive the Target, it is most convenient to connect it. Otherwise, it will be necessary to hold the Target reset (using the manual push button on the EVB97C102) during firmware downloads. For the Emulator used here, the Reset output is fully programmable, with a setting of LOW TRISTATE being correct for interfacing with the EVB97C102 reset circuitry. The EVB97C102 connection is at TP7, and the Emulator connection is at the Feature Connector pin 5; micro-hook cables suitable for making this connection are supplied with the Emulator.

A workable alternative to a ROM Emulator is a Monitor ROM, but that is not the approach used in this document for a variety of reasons: Monitor ROMs are extremely compiler-specific, require that an external RAM device be available, and often consume hardware resources beyond just their memory footprint (e.g., a Bank Select register when executing the application from RAM, a COM port for download, etc.). As a result of these characteristics, it is often not possible to use a Monitor ROM on final target hardware because the necessary resources might not be present. In addition, the target hardware might not contain even a single COM port for a debugging console, much less an additional one for code download. For situations like this, the ROM Emulator used here contains a COM port that the firmware can use as a debugging console even in situations where the target hardware has no COM

port of its own. Also, since ROM Emulators usually interface to the development system using a parallel port, rather than a COM port, code downloads are much faster than with a ROM Monitor; this is especially true if the MCU's serial port is used instead of a 550A type.

## 6. Oscilloscope or Logic Analyzer

An HP54645D Mixed Signal Oscilloscope (<http://www.hp.com>).

When executing the programs in this document, access to a scope is not required, since the software has already been developed and tested. However, when developing new software, at least part time access to a scope is a practical necessity, at least in the author's opinion.

Key features to look for are digital storage, numerous channels (8 or more is desirable), and deep memory (at least 100 K points/channel is desirable). Assuming that the scope is only being used for software development, it is not necessary to have analog channels or to sample at high speeds (10 Msps is about enough). It is also not necessary to have the elaborate trigger capabilities that are standard in Logic Analyzers these days, since it is a simple matter to have the software pulse a GPIO pin and provide a direct trigger when the desired event occurs, etc.

## Executing and/or re-Building the Examples

The following procedures are based on the use of the hardware and software just described. If different equipment is used, then other means will need to be used to achieve a port, which is beyond the scope of this document.

The Code Disk contains project files, listings, and final HEX files (in Intel Hex Format) for all of the example programs for the EVB97C102 Assy 6126. As a result, building the programs is not necessary in order to download and execute them.

In order to download the HEX files to the ROM Emulator, a set of **CHPTxx.BAT** files is provided. Each of these BAT files in turn executes the **Download.Bat** file, which produces a **Download.Cfg** file, which is finally passed to the **UrLoad.Exe** program to do the actual download to the ROM Emulator. The reason for all of this indirection is to ease the task of porting to different Host/Target/Emulator environments.

The **DownLoad.Bat** file should be modified in order to change:

1. The Host drive/path/name of the download program (e.g., c:\unirom\UrLoad.Exe)
2. The Host port to which the Emulator is attached (e.g., LPT2)
3. The Target memory type, size and base address being emulated (e.g., 128K Flash @ 0)
4. The Target reset circuit type (e.g., LOW TRISTATE)

Of course, if a different ROM Emulator, or a Monitor ROM, is used then other means will need to be used to achieve a port, which is beyond the scope of this document.

Rebuilding the example programs can be done from the Keil compiler by opening the desired project (e.g., **ChptXX.PRJ**, from the Project--Open Project dialog), and building it (from the Project -- Make: Build Project dialog). Of course, if a different compiler, etc. is used, then other means will need to be used to achieve a port, which is beyond the scope of this document.

If it is desired to make new projects using the Keil compiler, then it is important to adjust several settings:

1. Set the compiler to produce an assembly listing (if desired) (using Options -- C51 Compiler... -- Listing -- Include Assembly Code checkbox).
2. Define the symbols: **\_KEILC\_DBG** (using the Options -- C51 Compiler... -- Misc. -- Symbols for DEFINE command).
3. Set the internal RAM size to 256 bytes (using Options -- BL51 Code Banking Linker ... -- Size/Location).
4. Set the xdata RAM location (if used) to 0xC000 (using the same tabbed dialog sheet as above).

It is also possible to compile the example programs in Microsoft Visual C/C++, although the resulting binary cannot be executed. Version 4.0 was used in the development of the code examples. Even though the binary cannot be executed, it can still be useful to author code in MSVC as a check of the portability of the code, and in situations where the developer is more comfortable with that environment. The companion code disk contains a single MSVC project file **ProgRef.MDP** that contains subprojects for each of the example programs in this document. Since MSVC embeds absolute paths in the project file, it is best to copy everything to C:\USB\PROGREFV1\_1 and build from there. Under this directory should be additional subdirectories called CHx\Release, where "x" is 2, 3, 5, and 6; this is where the binaries for each subproject are placed.

If it is desired to make a new project in MSVC, do the following:

1. Create a new project using File -- New -- Project Workspace -- Console Application, and give it a name (e.g., ChptXX).
2. Use Insert -- Files into Project... to add the desired source (\*.C) files.
3. Define **\_MSVC\_DBG** using Build -- Settings -- C/C++ -- Category General -- Preprocessor Definitions. Optionally, set the Warning Level to 4 (maximum) on the same sheet.
4. Optionally, on the C/C++ Customize sheet, select the Disable Language Extensions checkbox in order to obtain the best assurance that the source code is generic ANSI C.

Note that because of the way the Keil compiler handles SFRs and SBITs, it was necessary to add some special code to **USB97C102.H** for non-Keil compilers like MSVC. This code only defines the subset of the SFRs and SBITs used by the example programs, so it might be necessary to add more if any new code makes use of additional SFRs and/or SBITs.

## The Example Programs

### **Chpt2.Hex:**

Described earlier, provides an example of different ways of implementing a software queue. Unlike the other examples, this code is not meant to be executed, but the ASM listing should be inspected to see and understand the impact of address spaces and coding style on performance.

### **Chpt3.Hex:**

This program illustrates basic programming techniques for the USB97C102 device and EVB. Topics covered are initialization, RAM access and testing (to qualify the hardware setup), DBGPRINT(), console I/O, use of GPIOs, software generated time delays, etc.

### **Chpt5.Hex:**

This program illustrates a variety of DMA programming techniques using the LPT port as an example.

### **Chpt6.Hex:**

This program illustrates an actual (albeit simple) USB application that is fully Chapter 9 compliant. While the primary emphasis is on coding techniques, the code is intentionally structured in a way that completely separates the core USB software from the actual application, with the intent that the application code can be easily replaced while leaving the core USB code intact.

### **Coding Style:**

Significant effort was expended to make the code examples as fully ANSI C as possible. All non-ANSI typedef's, etc. are declared in Types.H (with the exception of SFRs and SBITs, as was already discussed). There are a few places where some pragma's were needed in the code, but these are only used where needed. Obviously, no "//" in-line comments are used.

All variable names begin with a lower-case letter, while all function names begin with an upper case letter.

All function names begin with an upper case letter and contain no embedded underscores (e.g., MyFunction(a,b)). Macro functions that are called in the same way as an actual function use the same naming convention. However, macro's that use a calling convention that is different from what an equivalent function would use (e.g., if a variable



would be passed by reference to an actual function, but the variable name is passed to the macro) are named with all upper case letters and underscores separating words. In order to avoid confusion with register definitions, parentheses () are ALWAYS used with macro's, even when no arguments are passed (e.g., MY\_MACRO(a,b), YOUR\_MACRO()).

All variable names begin with a lower case letter, and type BYTE is assumed, unless one of the following prefixes is encountered:

- wFoo -- WORD
- bFOO -- bit
- cFOO -- code const
- szFOO -- ascii-Z string
- wzFOO -- unicode-Z string
- pFOO -- pointer

Also, structures have no prefix, since their identity as structures, rather than BYTEs, can easily be discerned by context. No attempt was made to prefix variables more fully (e.g., with their memory space, etc.) for fear that having too long a prefix on every variable would do more harm than good.

Register names follow the same convention as macro's, using all upper case letters and underscores; they can be differentiated from macro's because no parentheses are used (e.g., BUS\_REQ). Bit fields within registers are identified by a trailing underscore and a name that is indicative of the register to which it corresponds (e.g., BUS\_REQ\_HREQ\_ is the HREQ bit field in the BUS\_REQ register).

Manifest constants (i.e., #define's) use all upper case letters and underscores as well, but are prefaced with a lower case c (e.g., cUSB\_EP0\_MAX\_SIZE) so they can be differentiated from register names.

### Common Files

A number of files are common to many of the example programs, and will likely be useful in other applications:

**Types.H** contains definitions for all non-ANSI data types.

**USB97102.H** contains register and bit definitions for the device.

**USB97102.C** contains some useful misc. device-specific functions (e.g., HwReset).

**SIO.C/H** configures the SIO device and provides low-level console I/O functions.

**Debug.C/H** provides DBGPRINT and high-level console I/O redirected to SIO.C/H.

**Usb.H** provides definitions for USB related data types, etc.

**UsbCore.C/H** provides core code for handling connection and transfers on the USB.

## CHAPTER 3 - FIRMWARE BASICS

All of the examples in this document make use of the **Types.H** file, which encapsulates all of the non-ANSI data type declarations, as well as defining a number of useful macro's, etc. All of the executable examples (i.e., everything except **Chpt2.Hex**) also make use of the **USB97102.H** file, which contains defines for every register and bit field in the **USB97C102** device; all references to data types, registers and bits make use of the nomenclature in these two files. Experienced C programmers will recognize the use of the "volatile" keyword, which must be used with all memory-mapped I/O devices in order to prevent the compiler's optimizer from removing [what it thinks are] redundant memory accesses.

Several additional files are used in every example in this document, and every function contained in them is described in this section:

**Debug.C/H** - provides debug console I/O

**SIO.C/H** - provides code to initialize the SIO device and operate the COM port

**USB97102.C** - provides initialization code, etc.

### Initialization

At the end of POR (Power-On Reset), each core in the **USB97C102** device is initialized, every register (except for the DMAC) is filled with defined values, and the MCU begins executing code at address zero using the Ring Oscillator.

Each of the example programs calls the **USB97C102HwInit()** function in **USB97102.C** early in the execution of **main()** in order to complete the reset process. Since many of the values used by this function are application dependent, the code makes use of several global variables in code space, whose values are set by the application, in order to determine which values to place in the various device registers.

Even though this function configures the SIE and resets the MMU, it is important that it does **\*NOT\*** enable any USB Endpoints, even EP0, since this must wait until a USB Reset is received. This, and many other USB related issues, is discussed in detail in a later chapter.

Since the DMAC registers do not have defined values after POR, the function **DmacReset()** in **USB97102.C**, which is called by **USB97C102HwInit()**, places valid entries in every DMAC register.

### CLOCK\_SEL Register

One of the first things to do is to select the desired clocks for the MCU and the DMAC and to disable the Ring Oscillator. Note that it is important to select the new MCU clock source **BEFORE** disabling the Ring Oscillator; doing otherwise can shut off the clock to the MCU, which would halt execution.

Usually it is desired to use the highest speed clocks for the MCU and the DMAC in order to obtain best performance, but there are situations in which this is not the case. One example is any application in which power consumption is more important than performance. Since power consumption increases approximately linearly with clock frequency for CMOS devices like the **USB97C102**, lower clock frequencies translate directly into lower power consumption. Another consideration is the speed capabilities of the various hardware devices connected to the **USB97C102**, some of which might not be capable of operation with the fastest clocks. For the MCU, it is possible to slow the clock frequency before accessing a slow external peripheral device and then to restore the higher speed clock after the access. Since the DMAC operates asynchronously with respect to the software, the same technique cannot be used, and the DMAC should usually be set at a clock frequency, based on the capabilities of the slowest attached device that remains constant.

### **UTIL\_CONFIG Register**

The four lowest GPIO pins are multi-function, and this register provides control of the function of each of the individual pins. Although this register can be manipulated dynamically at run-time, it is most common to "set it and forget it" at initialization time.

In applications that use the MCU's internal Serial Port, this register is used to enable the Tx and Rx signals onto the GPIO0 and GPIO1 pins. This Serial Port is not nearly as capable as the 550A UARTs found in contemporary SIO devices, but it can be useful in some applications.

This register also permits input trigger selection for the MCU timers either from the USB SOF or from a GPIO pin. The ability to hardware trigger a timer from the SOF has a variety of uses, including (i) detection of a USB suspend, (ii) detection and reconstruction of missing SOFs, and (iii) intra-frame time measurement for Isochronous rate feedback. The ability to trigger a timer from a GPIO pin, combined with the fact that the timers can create interrupts, can be used to provide additional GPIRQs. However, the standard GPIRQs are more than sufficient for most intended applications of the device.

Since the desired setting of this register is application dependent, the sample code again makes use of a global BYTE in code space whose value is defined in the application code.

### **GPIOA\_DIR Register**

All GPIO pins are bi-directional, and this register controls the direction of each individual pin. Note that a pin's direction must be explicitly set in this register, regardless of the pin's function having been set in the **UTIL\_CONFIG** register. For example, setting the GPIO1/TXD pin function to TXD in the **UTIL\_CONFIG** register does *\*NOT\** automatically make the pin direction an output pin; the corresponding bit in **GPIOA\_DIR** register must be set in order to do this. Note that at POR, all pins are defined as inputs, so it is necessary to use a resistor (or other mechanism) in order to pull each signal high or low if a defined logic level is needed prior to the time when the software can set the appropriate registers.

### **GPIOA\_OUT Register**

This register defines the logic level applied to GPIO pins whose direction is set as output. Note that if this register is manipulated in an ISR(s) in addition to the foreground, then the foreground code should disable interrupts before access to this register and then restore the IRQ enable when finished; see **Chapt3.C** for an example of this.

### **MEM\_BANK Register**

This register controls a movable 16 KB window that appears identically in both the CODE and XDATA address spaces (i.e., Von Neuman model) of the MCU at address 0xC000. The memory accesses are to the FLASH Bus, and this register controls the high 6 bits of the 20 bit physical address. Among other things, this register can be used for bank selecting code pages, which is described later in this section.

For applications that need to store more data than what will fit in the internal RAM, this address space can be used to access external RAM. ROM Monitors often require this (Von Neuman) style of memory as well. On the EVB97C102, there is a 128 KB RAM located at physical address 0x40000 (i.e., 256K above address zero); to map the start of this RAM, fill the **MEM\_BANK** register with 0x10 (the high 6 bits of the address, right-aligned in this register), which is what **Chpt3.Hex** does.

There are a number of considerations if this address space is used for RAM that contains variables declared as BYTE\_X, etc. First of all, the compiler's start up code will execute before the application has an opportunity to set the **MEM\_BANK** register contents. This means that either any variables allocated there must not be initialized by the startup code, or that the start up code must be modified to set this register before doing the initialization. Of course, the Linker also needs to be informed of the location of the 0xC000 window in order to generate the proper code.

Do not be misled by the name "FLASH Bus" since nothing about this bus restricts its usage to only Flash Memory devices. With its 8-bit Data, 20-bit Address and ISA-style nRD and nWR signals, it can be used for any combination of RAM, ROM, EPROM, EEPROM, FLASH, and peripheral I/O devices. The two primary differences between this bus and the ISA bus are (1) the ISA bus has DMA, while this bus does not, and (2) the ISA bus must be shared with the DMAC while this bus is owned full time by the MCU. As a result of this last fact, it is preferable to interface all non-DMA peripherals to this bus.

In addition to the movable window described above, there is also a single fixed 16 KB window mapped to address zero both on the Flash Bus and in the MCU code address space. For applications that can fit in 16 KB of code

space, this window is sufficient to contain the program code and the movable window can be used for external data access as was previously described. For applications that require more code space, the movable window is also mapped into 0x4000-0x7FFF in the MCU code space, and can be used for program storage. If the MEM\_BANK register is set to 0x01 (the POR default value), then this second window will map straight through to the Flash Bus, which provides a contiguous 32 KB code space starting at address zero on the Flash Bus; this is what Chpt6.Hex does (V1.1 and later). Treating the low 16 KB window as the “root” segment, placing the overlaid code in the upper 16 KB movable window, and using the MEM\_BANK register to select the desired overlay can support code overlays. For applications that require both large code space (with or without overlays) as well as external RAM or I/O access, a simple decoder (e.g., a PAL, CPLD, FPGA, etc.) can be placed on the Flash Bus that maps the top (e.g., 1 KB) portion of each 16 KB page (except probably for the bottom page) to the RAM or I/O device(s), with the lower portion (e.g., 15 KB) of each page selecting the desired ROM page.

### MEM\_BANK2 Register

This register controls a movable 16 KB window that appears identically in both the CODE and XDATA address spaces (i.e., Von Neuman model) of the MCU at address 0x8000-BFFF. The memory accesses are to the FLASH Bus. MEM\_BANK2 register controls the most significant 6 bits of the 20-bit Flash address bus. These 6 bits are used for bank-selecting one of the 64 16K byte Banks/Pages in the external flash bus.

**Note:** MEM\_BANK2 register resets to a default value 0x00, mapping directly to BANK/PAGE 0.

### IOBASE Register

This register is used to access the 64k bytes of ISA I/O that is available on the ISA bus. The ISA bus is accessed through a 256 byte movable window that appears at XBYTE [0x4000]; the most significant 8 bits of the 16-bit ISA I/O address are contained in the **IOBASE** register. Since it is common to have peripheral devices spread throughout the ISA address space, it is common for this register to be re-written extensively at run-time.

It is common for foreground code to perform this type of manipulation, so if this register is manipulated in an ISR as well, it is important that it be saved and restored in the ISR in order to avoid disrupting the foreground code.

In order to address an I/O port located (e.g. at 0x0577h) on the ISA bus, load IOBASE register with 0x05h and access the I/O port through 0x4077h in the Xdata space. If wanting to access the registers of the 8237A DMA controller, which is internal to the USB97C102, load IOBASE register with 0x00h and R/W to 0x4000 in the Xdata space. IOBASE register is loaded with 0x00h, since the registers of the DMA controller are located in the ISA I/O space at base address 0x0000h. Only the DMA controller can access the MMU's packet buffers via the ISA bus.

### MEMBASE Register

This register performs a similar function to the **IOBASE** register, except that it controls a 4 KB movable window into the ISA memory address space that appears at XBYTE [0x5000]; the high 8 bits of the 20-bit ISA memory address are contained in the **MEMBASE** register. All of the comments about the **IOBASE** register above apply to this register as well.

#### Shared Bus Architecture Basics:

Note that the ISA bus is shared between the MCU and the DMAC, so the MCU must acquire ownership of the ISA bus before it can access it. A variety of issues related to this is discussed in detail in a later chapter. For the moment, it is sufficient to mention that the functions **IsaAcquire()** and **IsaRelease()** in **USB97102.H** provide the mechanism by which the MCU can acquire ISA bus ownership and release it back to the DMAC respectively. Since ownership of the ISA bus by the MCU suspends all DMA traffic, it is usually desirable to disable interrupts for the duration of the ownership in order to prevent DMA suspension for an extended period of time while the ISR executes; the functions **InterruptDisable()** and **InterruptEnable()**, also in **USB97102.H**, provide the means for doing this. Note that, like any 8051 derivative, the **EA** bit provides a global interrupt enable (when TRUE), and it is usually desirable to save/restore this bit when manipulating it.

### SioInit()

This function, which is contained in **Sio.C**, initializes the FDC37C672 SIO device on the EVB. Since the examples in this document only make use of a single COM port and the LPT port, these are the only functions that are initialized. For other applications, the device contains an additional COM port (with a full set of handshake lines on the EVB), an FDC (Floppy Disk Controller), P/S Keyboard/Mouse controller, and the COM2 port used in these

applications can be reconfigured as an IrCC (InfraRed Communications Controller) that supports IrDA FIR (Fast InfraRed), as well as a variety of Consumer IR protocols. For detailed information about the SIO device, the interested reader is referred to the device Data Sheet.

Briefly, the SIO device is intended for use on a PC motherboard controlled by a PnP/APM and/or ACPI BIOS. As a result of this, all hardware resource assignments are fully configurable and each device can be disabled as well. The **SioInit()** function simply configures the COM2 and LPT devices to use the resource assignments defined in **SIO.H**. **SIO.H** then defines various registers and bit fields based on these values. It is instructive to look carefully at the manner in which **SIO.H** performs the register definition -- each register is defined in such a way that the final xdata address is calculated at compile time, and not at run time, which results in optimum performance; a comment to this effect appears in the source code.

Note that this function saves and restores the state of the **EA** bit and disables interrupts before acquiring the ISA bus in order to access the SIO. On return, the **BUS\_REQ** and **IOBASE** registers are restored to their previous settings, along with the **EA** bit. Although this register save/restore is not necessary for a function that executes only at initialization time, the technique is useful in many other places, and this basic code sequence will be seen often in the examples.

Since the examples use the COM2 port for debug console I/O, **SioInit()** calls the function **Com2Init()** to initialize the COM port before it returns. The COM port is set for 115.2 Kbps, 8 data, 1 stop, and no parity. Note that the COM2 port on the EVB has no handshake lines, so there can be no hardware flow control.

### **Platform\_Display()**

One EVB (ASSY 6126 Rev. B) contains a HP 8-Character LED display that is connected to ispLSI 1032 I/O lines. This function, which is contained in **Debug.C**, will display the first 8 characters from the string to the display. For the examples in this document, the corresponding string "ChapterX" of the corresponding program will be displayed. In practice, this function is used as a debugging tool if the software appears to get hung at some unknown location; A "failed" string will be displayed. Also, by placing calls to this function at various points in the code, it becomes possible to determine what section of code executed last before the system hung. Execution time is relatively short, on the order of about 10 microseconds, so even relatively frequent calls will not have a significant impact on execution timing.

### **DBGPRINT()**

This is a macro function defined in **Debug.H** that is conditionally compiled based on a "**DBG**" command-line switch to the compiler. When this switch is defined, the function simply passes its argument to **printf()** in the RTL (Run-Time Library); when this switch is not defined, **DBGPRINT()** compiles to nothing. This permits keeping extensive debugging information in the source code, while being able to make a compact production release by simply undefining **DBG**. It also permits disabling **DBGPRINT** in various sections of the source code by using preprocessor directives (**#ifndef DBG / #define DBG, #ifdef DBG/ #undef DBG**).

There is also a **DBGTRACE** macro function that behaves in a similar way, except that it is also gated by **TRACE\_ON**. Changing the value of this symbol in the source code permits disabling TRACE messages while leaving PRINT messages enabled. One use for this capability is to use TRACE messages to show the code execution sequence for "normal" conditions, and to use PRINT messages to display any error conditions. Once the code seems to be functional, the TRACE messages can be disabled, resulting in smaller code size and MUCH better performance (since printing through a COM port is really S--L--O--W), but any errors are still displayed. If the cause of an error is not obvious, then the TRACE messages can be enabled again simply by uncommenting a single line in **Debug.H**. The author found all of this very useful when developing the example programs.

### **printf()**

This RTL function eventually calls the **putchar()** function, which appears in **Debug.C**. In the example code, this function just redirects to **Com2SendByte()** from **Sio.C**. For other target hardware platforms this code could instead make use of the MCU Serial Port, a COM port contained in a ROM Emulator, or any other output device that happens to be available.

Unfortunately, this function is usually not reentrant, even in RTL's intended for embedded applications, so none of this can be used in an ISR. However, the execution time of a **printf()**, especially through a COM port, is so long that you would not want to do it in an ISR anyway. If it is necessary to display information from an ISR, it is suggested that a circular buffer be used in which the ISR places very short messages, perhaps a single DWORD or so, and the foreground polling loop can then send them to the terminal. By keeping ISR's as short and as simple as

possible, it should not usually be necessary to do this, so none of the sample code illustrates the technique; as the text books say, "It is left as an exercise for the interested reader."

### **Com2SendByte()**

This function, from **Sio.C**, does the usual dance with **EA**, **BUS\_REQ** and **IOBASE**. The hardware resource assignments, register access, etc. are all from **Sio.H**.

The one interesting twist is that, while it is waiting for the UART to be ready for Tx, it restores the registers and the **EA** bit in order to permit DMA and IRQs to proceed. This is critical if **DBGPRINT**'s are to be used successfully in an actual USB application, since not doing so would suspend IRQs and DMA for up to 100 uS (@ 115.2 Kbps, and much longer at slower speeds) every time this function was called, which would not be a good thing.

### **kbhit()**

This function, from **Debug.C**, simply redirects to **Com2IsRxRdy()** in **Sio.C**, similar to the arrangement with **putchar()** above. Once again, redirection could be to any other available input device for other target hardware platforms. **Com2IsRxRdy()**, as usual, gets its hardware assignments from **Sio.H**, does the usual **EA**, etc. dance as before, and simply returns the appropriate bit from the UART's **LSR**. Note that this function is fairly fast since it never has to wait for anything, unlike **Com2SendByte()**.

### **\_getkey()**

This function, also in **Debug.C**, simply redirects to a function in **Sio.C**, this time **Com2GetByte()**. The same methods are used here as for the other functions previously described.

**HINT:** Note that this function will wait for a keystroke, so it can be useful to check **kbhit()** first before calling this function in order to avoid waiting for an indefinite period of time.

### **Delay1uS():**

From **USB97102.C**, this function implements an accurate software time delay based on an MCU clock of 24 MHZ. Note that the function does not change the IRQ enable, so it is up to the caller to do so if that is what is really desired; it is not often desired to disable IRQ's while waiting in a time delay function, but in very special circumstances (e.g., doing precise timing interfaces to a hardware device) it has its uses.

### **Code Example: Chpt3.Hex**

This program initializes the hardware, prints a message on the debugging terminal, waits for a keystroke, performs a RAM test on both the ISA and FLASH busses, and pulses a GPIO pin in an infinite loop.

In order to execute this program, the COM2 jumpers on the EVB must be set to select the SIO, as opposed to the MCU COM port. On Assy 6126 Rev. B, JP9 and JP10 must have jumpers between pins 1 and 2. See Chapter 2 for a general description of the equipment setup.

The use of "goto" statements in C code is something of a religious issue. Many C purists feel that it was a mistake to even include it in the language, but many of these same people have no problem using **setjmp/longjmp**; go figure... It is the author's style to use goto's in lieu of C++ or Win32 Structured Exception Handling (e.g., **try/throw/catch** or **try/leave/finally**), especially in situations in which there is some non-trivial exit processing to do. Consistent with this philosophy, there are a total of 3 places in the entire sample code for this document in which goto's are used; readers who are offended by this style are welcome to modify the code to eliminate them.

The **RamTest()** function performs a test of the RAM devices on each bus; it is based on a simple traveling-zero's/traveling-one's algorithm. Note that it does the dance with **EA** and **BUS\_REQ** even though it is not needed in this particular application since no IRQs or DMA are enabled; the same can be said for disabling/enabling IRQs around the GPIO pulsing in the final loop. However, that exact sequence is required in real applications in which IRQs are being processed and the ISR will often be accessing the same **GPIOA\_OUT** register. It was the author's judgment that the reader would be better served by illustrating these techniques right from the beginning, rather than showing coding techniques at this point that would need to be unlearned in a later chapter in order to create code that could be used in a real application.

The reader is **strongly encouraged** to execute this program as part of validating the equipment setup. In particular, failures on either of the RAM tests can be caused by an improper interface to the ROM Emulator, as was described in Chapter 2.

## CHAPTER 4 - THE MMU

The MMU (Memory Management Unit) is responsible for managing the 4 KB Data Buffer RAM that is used for all USB communications. The memory is organized as 32 pages of 128 bytes each. Since each USB packet has an 8-byte packet header associated with it, this means that a maximum size (64 bytes) USB BULK packet can fit in a single MMU page. Larger packets are handled by the MMU concatenating multiple (up to 10) pages together. Even when packets are physically comprised of multiple pages, the MMU creates the illusion of a single large virtually contiguous packet for all 3 ports: the SIE (interface to USB), the MCU (for PIO), and the DMAC.

The MMU in this device is extremely similar to the MMU in the SMSC LAN91C94/5/6 devices, and the interested reader is referred to those data sheets for a more complete discussion of the architecture; the focus of this chapter is on the register-level programming instead.

### Allocating and Freeing Packet Memories with the MCU

The MCU can instruct the MMU to allocate a packet of a desired number of pages, or to free a specific packet, by using the **MMUCR** (MMU Command) register. The MCU can tell when an allocation has completed by checking the **ARR** (Allocation Result) register. Assuming that the MMU has enough free memory pages to satisfy the request, the allocation is very fast (a couple of microseconds or so), so the firmware should just wait for the allocation to complete (as opposed to leaving the procedure and doing other work).

The following block of code will allocate a packet:

```
MMUCR = (MMUCR_ALLOCATE_ | (numPages-1));
while (ARR & ARR_FAILED_); /* wait for MMU to finish */
pkt = ARR & PN_MASK_; /* save result */
```

The following block of code will free a packet:

```
PNR = pkt;
MMUCR = MMUCR_RELEASE_;
```

Note that the packet number, saved in the variable "pkt" in the above code examples, is analogous to a handle -- it tells the MMU which packet is being referred to, but has no particular relationship to any physical memory address, etc. The detail of what to do with these packet numbers is discussed in the following sections.

The **MakePkt()** function in **Chpt5.c** shows an example of allocating a packet, and also of filling it with data, which is discussed in the USB Transmission section below. The **FreePkt()** function in **Chpt5.C** shows how to free a packet using these methods.

### USB Reception

As each Rx packet arrives from the USB and the MMU allocates a packet memory for it, it pushes the packet number onto a queue that the MCU can access by reading the **RXFIFO** register. The MCU can then inspect the packet header (the first 8 bytes of each packet) in order to decide what to do with it, based on things like the Data Toggle PID, Endpoint Address, etc. The MCU then has a choice of dropping the packet and freeing the memory, or of just removing the packet from the **RXFIFO** queue but leaving the memory allocated.

Normally, handling USB Rx packets is done in **Isr0()**, which is the ISR that handles IRQ0 interrupts; an example of unmasking this interrupt is shown in **main()** in **Chpt6.C**. To do this, the corresponding bit in the **IMR\_0** register must be cleared, and IRQ0 must be enabled in the MCU's **EX0** SBIT:

```
EX0 = 1
IMR_0 = ~INT0_RX_PKT_;
InterruptEnable();
```

Having done this, the ISR for IRQ0 will be executed each time a USB Rx packet arrives. In the ISR, the MCU can tell if the IRQ is from a USB Rx by inspecting the **RXFIFO\_EMPTY\_** bit in the **RXFIFO** register. Assuming that there is one or more packets in the **RXFIFO**, the MCU can access the packet at the head of the **RXFIFO** by doing the following:

```

PRL = 0; /* packet header begins at offset zero */
PRH = PRH_RCV_ | PRH_READ_ | PRH_AUTO_INCR_;
pktHdr0 = MMU_DATA;

/* TODO: read any other bytes of interest */

```

Setting the **PRH\_RCV\_** field in the **PRH** register tells the MCU to access the packet at the head of the **RXFIFO**, rather than the one in the **PNR**, for subsequent transfers through the **MMU\_DATA** register. The other fields in the **PRH** setting tell the MMU to read from the packet, as opposed to writing to the packet, and to increment addresses after each access. If it is desired to read from a packet that is not at the head of the **RXFIFO**, this can be done by placing the packet number in the **PNR**, and then setting the **PRH** for **PRH\_PNR\_**. For example:

```

PNR = pkt;
bOldEA = EA; /* see next section in Reference Guide text */
InterruptDisable();
PRL = 6; /* byte count is at offset 6 */
PRH = PRH_PNR_ | PRH_READ_ | PRH_AUTO_INCR_;
EA = bOldEA;
cntLo = MMU_DATA;
cntHi = MMU_DATA;

```

Note that, in the code samples above, the **PRL** register is ALWAYS written before the **PRH** register; this is a device requirement and is not optional. As a result of this fact, **it is critical** that interrupts be disabled in any foreground thread that uses these registers if they are also shared by an ISR. Failing to do so can result in the ISR writing the **PRL** register after the foreground has done so, but before the foreground has written the **PRH** register, which violates the device requirement for accessing these registers. The code sample above illustrates this technique, which is used often in the example programs (see Chpt5.C and UsbCore.C).

A careful reading of the USB97C102 Data Sheet reveals that the MMU can take up to 1.218 uS after the **PNR** is written to present valid READ data. One might think that it is necessary to insert extra delay after setting **PRH** before reading **MMU\_DATA**, but this is not the case. The reason for this is the way the MMU is interfaced to the MCU in its XDATA address space. Recall that XDATA access must use the **DPTR** register, and that the **PRH** and **MMU\_DATA** are at different addresses in the XDATA address space. As a result of this, the most efficient code possible for reading the **MMU\_DATA** register after writing the **PRH** register is as follows:

```

MOV DPTR, #6000H
MOV A, @DPTR

```

Each of these is a 2-cycle instruction, with the actual XDATA access happening in the 2nd cycle of the **MOV A, @DPTR** instruction. As a result, an absolute minimum of 3 instruction cycles is guaranteed, which is more than the time that the MMU needs.

The Data Sheet says that the MMU sequential access time is 588 nS. Since access to the **MMU\_DATA** register requires a **MOV A, @DPTR** instruction, which takes 2 instruction cycles, this access time is satisfied, even when the MCU operates at its maximum clock frequency of 24 MHz, so no additional delay is required.

Getting back to the received packet, Byte0 in the packet header is the most interesting when handling a USB Rx packet. First of all, there is the **PKT\_HDR\_BAD\_CRC\_** bit that will be set if the packet arrived with a bad CRC; such packets should always be dropped:

```

if (pktHdr0 & PKT_HDR_BAD_CRC_ ) {
    MMUCR = MMUCR_REMOVE_RELEASE_;
    goto _next_pkt;
}

```



Note that the **REMOVE\_RELEASE\_** command applies to the packet at the head of the **RXFIFO**, and not to the packet number in the **PNR**; the command both removes the packet from the **RXFIFO** and releases the memory that had been backing the packet number.

Since in practice the USB has a very low BER, it should not be often that a packet will arrive with a bad CRC, but properly written software will check for it nonetheless. The fact that the occasional bad CRC packet is received is useful for ISO applications in which it is necessary to know in which specific USB frame the packet arrived, even though the data is unusable. As an aside, the packet header also contains the USB frame number in which the packet arrived, which is also useful in ISO applications.

The **pktHdr0** byte also contains the EP address in the low nibble:

```
ep = pktHdr0 & EP_MASK_;
```

The target EP is significant in order to determine how to handle the reception. For example, the Data Toggle PID should be ignored for ISO, but it must be checked for BULK packets, and Control Write packets on EP0, in order to detect duplicate packets, which should be dropped. The **pktHdr0** byte contains 2 bits for this purpose: **PKT\_HDR\_LAST\_TOG\_** contains the Data Toggle value of the previous packet received on the same RxEP, while the **PKT\_HDR\_SAME\_TOG\_** bit indicates if the current packet's Data Toggle is the same as the previous packet. One might assume that simply checking the **PKT\_HDR\_SAME\_TOG\_** bit would be sufficient for BULK packets, but it is not because a Control Transfer (e.g., CFES, etc.) might have executed between the receptions. The solution is to maintain a bit variable in software for each RxEP that keeps track of the expected Data Toggle. As each packet arrives, the Data Toggle for that packet is calculated and compared with the expected value; packets with the incorrect Data Toggle value are discarded.

```
switch (ep) {
case 1:
    if (pktHdr0 & PKT_HDR_LAST_TOG_) /* expect opposite toggle from last time */
        bThisTog = 0;
    else
        bThisTog = 1;
    if (pktHdr0 & PKT_HDR_SAME_TOG_) /* unless it turns out to be the same */
        bThisTog = ~bThisTog;

    if (bThisTog != bEp1RxToggle) { /* mismatch, so drop pkt as a duplicate */
        MMUCR = MMUCR_REMOVE_RELEASE_;
        goto _next_pkt;
    }

    bEp1RxToggle = ~bEp1RxToggle; /* invert toggle for next time */

    /* TODO: whatever you do with packets for this EP */
    MMUCR = MMUCR_REMOVE_; /* remove from RXFIFO, but keep memory */

    break;

case xxx:
```

The situation for duplicate packets is analogous to that for bad CRC -- it should not be expected to happen often, but it can happen, and properly written software will handle it as described.

As can be seen from the above code samples, it is necessary to use at least the **PRH** and **PRL** registers, and sometimes also the **PNR**, in order to handle a USB reception. However, these registers are also used by the foreground thread whenever it has to access a packet memory. Assuming that the receptions are handled in an ISR, it is necessary for the firmware to save these registers on entry and to restore them on return. It is also necessary for the foreground thread to disable interrupts while accessing the PRL/PRH register pair, as was previously described.

```
void UsbRxIsr(void)
{
```

```

if (!(RXFIFO & RXFIFO_EMPTY_)) {
    /* Save MMU registers */
    oldPNR = PNR;
    oldPRL = PRL;
    oldPRH = PRH;

    do {
        /* TODO: Handle all receive packets */
        while { (!(RXFIFO & RXFIFO_EMPTY_));

        /* Restore MMU Registers */
        PNR = oldPNR;
        PRL = oldPRL;
        PRH = oldPRH;
    }
}

```

A careful reading of the USB97C102 Data Sheet reveals that the **PNR**, **PRH**, and **PRL** registers should not be modified for at least 2.5 uS after the previous write to **MMU\_DATA** when writing to a packet, which the foreground thread might have been doing when the IRQ happened. However, just saving the MMU registers takes much more than 2.5 uS, so the required time is more than satisfied without adding any extra delay.

An alternative strategy for ISO applications is to use the **INT1\_SOF\_** IRQ instead of the **INT0\_RX\_PKT\_** IRQ; doing so will cause the ISR to execute exactly once for each USB frame, which is a desirable characteristic in some ISO applications. In addition, if the code is carefully structured, it might be possible to never disable IRQs in the foreground. If this can be achieved, and if the ISR code is written with minimal branching and branch balancing, then the execution of every part of the ISR will exhibit low jitter from one USB frame to the next, which is a desirable characteristic in most ISO applications.

Another strategy issue is where to check for bad CRC and Data Toggle -- in the ISR or in the foreground. It is the author's preference to do it in the ISR so that the corresponding packet memories can be freed as quickly as possible, and it also keeps all of this code in a single place, instead of distributing it among the foreground device handlers. In addition, it avoids any consideration of synchronizing with EP0 transfers like CFES. However, either approach is valid, provided that the checks are performed someplace in the code.

In applications that use 64-byte BULK packets, a fast Host can deliver a packet roughly every 50 uS. In most cases, it is desirable to have an **RXFIFO** loop that is at least as fast as the packet arrival time in order to avoid an overrun situation. This suggests that the minimum necessary processing should occur in the **RXFIFO** loop, and the code should be written for the best possible execution speed. This in turn affects the style with which the code should be written, as was previously discussed in Chapter 2. It is this concern for speed that makes it reasonable to even consider deferring the CRC and Data Toggle validation.

Regardless of how the receive handling code is designed, it is essential that it be implemented in such a way that the **RXFIFO** never overflows and the MMU never runs out of free pages. The reason for this is that it is a requirement that every USB device must always be capable of receiving a SETUP packet on EP0; if either of the above conditions occurs, the device will not be capable of receiving anything, including a Setup packet, rendering it out of compliance with the USB Specification. The details of both Rx and Tx Memory Management Policy (MMP) code are discussed in later sections of this chapter.

## USB Transmission

In order for the MCU to request that the MMU queue a packet for transmission on the USB, the first thing the MCU needs is a packet. This can either be a packet that the MCU allocated, or it can be one that was allocated by the MMU as a result of a USB reception. In fact, taking a packet number that arrived in the **RXFIFO** and retransmitting it is the most basic form of Loopback test with this device.

Once a packet has been obtained, the next thing to do is to write the desired byte count to the header portion of the packet; this is how the SIE knows how many bytes to send on the USB. Finally, the data portion of the packet can be filled, starting at a byte offset of 8, since the first 8 bytes are considered to be the header by the SIE.

The byte count is usually written to the packet by the MCU using PIO (Programmed I/O) through the **MMU\_DATA** register. The data portion of the packet is usually filled by the DMAC (which is discussed in a later chapter) for high

performance devices, but can also be filled by the MCU for slower non-DMA devices, and is almost always filled by the MCU when handling EP0 traffic. When writing to a packet using the MCU, the sequence is as follows:

```

PNR = pkt;
bOldEA = EA;
InterruptDisable();
PRL = 6;          /* offset of count low byte in packet header */
PRH = PRH_PNR_ | PRH_WRITE_ | PRH_AUTO_INCR_;
EA = bOldEA;
MMU_DATA = LOBYTE(wSize);
MMU_DATA = HIBYTE(wSize);

while(bytesToSend--)
    MMU_DATA = *pBuf++;

/*
 *   TODO: be sure to wait at least 2.52 uS before changing PRH or PNR,
 *         which is not hard to do on this MCU
 */

```

Setting the **PRH\_PNR\_** field (actually, it clears a bit) in the **PRH** register causes the MMU to use the packet number in the **PNR** register, rather than the head of the **RXFIFO**, for the subsequent access through the **MMU\_DATA** register. For an example using this technique, refer again to the **MakePkt()** function in **Chpt5.c**.

Each of the 16 USB TxEPs has a 5-deep TxFIFO associated with it. Each endpoint has a three bit up/dn counter, **TX\_FIFOx**, which will maintain the number of packets queued for transmit at that endpoint. The counter is incremented when there is a push on the TxFIFO of the corresponding endpoint, and it is decremented when there is a pop on the Tx FIFO of the corresponding endpoint.

The empty/full status of each TxFIFO is available to the MCU in the **TXSTAT\_A - TXSTAT\_D** registers. When the MCU wants to transmit a packet on a given TxEP, it must first check to make sure the corresponding TxFIFO is not already full, and then the MCU must push the packet onto the TxFIFO. Once the packet is filled with data, its packet size is set (the order is not important, just as long as they both get done), and it is known that there is room in the desired TxFIFO, then it is time to queue the packet for transmission. To do this, the packet number must be written to the **PNR**, the desired USB Endpoint Address must be written to the **TX\_SEL** register, and the **MMUCR** must be written with the command: **MMUCR = MMUCR\_ENQUEUE\_**. The following code sequence illustrates these techniques by queuing a packet for transmission on txEP2:

```

if (!(TXSTAT_A & TXSTAT_A_    EP2TX_FULL_)) {
    PNR = pkt;
    TX_SEL = 2;
    MMUCR = MMUCR_ENQUEUE_;
}

```

When IN tokens arrive from the Host, the SIE will transmit the packets in the order they were pushed on the corresponding TxFIFO. As each packet transmission is completed, the packet number is pushed onto the Tx Completion Queue, which the MCU can access in the **TX\_MGT2** register. Note that the MCU is not required to inspect this queue, it is simply available if the firmware author wishes to use it; there is no problem with letting the queue overflow. One use for this queue is the situation in which the firmware is implementing Memory Management Policy code, and as a part of that code, it is keeping track of how many packets are presently owned by each TxEP. By inspecting the Tx Completion Queue, the code can correctly decrement the count for the corresponding TxEP as each packet is actually sent. For example:

```

while (!(temp=TX_MGMT2) & CTX_EMPTY_) {
    pkt = temp & PN_MASK_;
    /* TODO: whatever you want to do with the packets */
}

```

When a packet is transmitted, the default case for the MMU is to free the corresponding packet memory automatically. However, this feature can be disabled using the **TX\_MGT\_MEM\_DALL\_** bit in the **TX\_MGT** register.

One possible reason for not having the packets freed automatically is to permit the packet memories to be recycled after transmission without having to go through a new allocation. However, the MMU can allocate packets quickly, so there is usually no obvious benefit to this approach.

#### Flushing a TxFIFO:

There are situations in which it is necessary to flush the packets that have already been queued for transmission on a particular EP; an example of this is when handling a CFES command. There is a **RESET\_TX\_** command in the **MMUCR**, but this will **ONLY** reset the specified TxFIFO -- it will **NOT** free the associated packet memories. In order to release the packets, the MCU must remove the packets from the TxFIFO and free them one at a time; the **POP\_TX** register provides the mechanism for doing this. The following code sequence illustrates the correct way to flush a TxFIFO:

```
MMUTX_SEL = 2;
while (! (TXSTAT_A & TXSTAT_A_EP2TX_EMPTY)) {
    PNR = POP_TX & PN_MASK_;
    MMUCR = MMUCR_RELEASE_;
}
MMUCR = MMUCR_RESET_TX_;
```

Although the methods used to control the state of each EP are described in a later chapter, it is important to note here that the EP corresponding to a TxFIFO that is being flushed must **\*NOT\*** be enabled during the flushing operation due to the possibility of the Host issuing an IN token during the flush. Note that when handling a CFES command, the EP is already STALL'd, so this requirement is automatically satisfied.

#### Memory Management Policy (MMP)

The Memory Management Policy feature permits limiting the number of received packets in memory per endpoint. It allows the USB97C102 to dynamically utilize the memory buffer; supporting 32 endpoints with finite buffer memory.

#### MMPCMD Register:

This register allows the MCU to access and control the up/down counters for each endpoint. A five-bit up/down counter will be implemented for each endpoint. Each counter will be incremented by the MCU to initialize the limit, then decremented by the hardware as packets arrive at its corresponding endpoint, and incremented by the MCU after it releases the packet. If the count reaches 0, and the MMP feature is enabled, then the hardware will not receive the packet and will NAK non-isochronous OUT tokens. If the count is zero, it will not decrement further; if the count is 31 it will not increment further. The MCU can enable or disable this feature independently for each endpoint. The default condition is disabled.

If MMP feature is disabled before counter reaches zero, the endpoint counters will still count, but there will be no MMP action taken when the counter reaches zero.

Setting the MMP method support:

```
    :
    :
{
    #if 0

    endpoint_rx_busy(0);

/* Resetting the counter to zero and disabling the MMP */

    MMPCMD = MMP_RST_DSB_ | 0;

/* Increment the count */

    MMPCMD = MMP_INCREMENT_ | 0;

/*Enabling the MMP */

    MMPCMD = MMP_ENABLE_ | 0;
```

*/\* This will cause the count and the enable / disable state to be latched into the MMPSTAT reg \*/*

```
MMPCMD = MMP_GETSTATE_ | 0;

#endif
:
:
}
```

#### **IN / OUT Nak Registers:**

During a well-implemented token decode, a function upon receiving a data packet, may return any one of the three handshake types. If data is corrupted, the function returns no handshake. If the data packet was received error-free and the function's receiving endpoint is halted, the function returns STALL. If the transaction is maintaining sequence bit synchronization and a mismatch is detected then the function returns ACK and discards the data. If the function can accept the data and has received the data error-free, it returns ACK. If the function cannot accept the data packet due to flow control reasons, it returns NAK.

NAK can only be return by functions in the data phase of IN transactions or handshake phase of OUT Transactions. The host can never issue NAK. NAK is used mostly for flow control purpose to indicate that a function is temporarily unable to transmit or receive data, but will eventually be able to do so without need of host intervention.

Knowing this, the IN\_NAK register's bit(s) are set every time the SIE responds with a NAK to IN tokens on the corresponding endpoint. It is reset when the MCU writes a one to the corresponding bit. OUT\_NAK register works the same way as IN\_NAK register except that the corresponding bit(s) are set when SIE responds with a NAK to OUT tokens. For example:

*/\* host is not retrieving the data fast enough \*/*

```
if (!recovery_shunt_remainder)
{
    recovery_shunt_remainder = 1;
}
```

*/\* if tx is in progress we can't purge, wait until we see an innak, so we know that no xmit is in progress. \*/*

```
INNAK_LO = 0xFF;
endpoint_tx_busy(2);
while (!(INNAK_LO & 0x04));
```

```
txfifo_purge(2);
```

*/\* value of last pkt sent out \*/*

```
recovery_start_pkt = npkts_this_cyl;
```

```
endpoint_tx_enable(2);
```

#### **PAGS\_FREE Register:**

The MCU can determine the current number of free memory pages by inspecting the **PAGS\_FREE** register, but this is of limited use in a USB application while pages are constantly being allocating and freed. There is a **PAGS\_FREE\_NAK\_ALLRX** bit in the **PAGS\_FREE** register that can be used to reduce the traffic, but this should not be used often, if at all, due to performance considerations. Potentially this combination could be used to implement a stochastic MMP algorithm.

It is the author's preference to use a deterministic MMP approach in which the peak memory usage of each EP is planned in advance, and the run-time code consists of making sure that no EP ever goes over its limit. An example of such an approach appears in a later chapter.

### **GP\_FIFO's:**

In the Data Sheet, these are lumped in with the ISA Bus Control Registers. In this document, they appear here in with the MMU Registers. In the actual device, they do not have anything to do with either of these, but they are so similar to the FIFO's that are part of the MMU, that this seems like the best place to put them.

Each of these FIFO's is byte-wide and 8 deep, with its own status register that indicates the empty/full status, just like the TxFIFO's. At POR, they are cleared empty. Software should never read from them when empty, or write to them when full, because the result is not defined. Use of these FIFO's is as follows:

```
/* push a packet onto GP_FIFO1 */  
if (!(GPFIFO1_STS & GPFIFO_FULL_))  
    GP_FIFO1 = pkt;  
  
/* pop a packet from GP_FIFO1 */  
if (!(GPFIFO1_STS & GPFIFO_EMPTY_))  
    pkt = GP_FIFO1;
```

There is an example of using these FIFO's in **Ep0RxIsr()** in **UsbCore.C**, which is discussed in a later chapter.

## CHAPTER 5 - DMA

### Shared Bus Architecture Details:

The USB97C102 contains an 8237 DMAC, which is the same device used in the PC. However, the interface to the MCU in this device is different than the MPU interface in the PC because, unlike an MPU, the MCU does not have any HRQ (Hold ReQuest) or HLDA (HoLD Acknowledge) signals with which to accomplish the traditional interface. Although this section summarizes the DMAC operation, the interested reader is referred to the 8237 Data Sheet for a complete description of the device.

In order to understand the operation of this interface, it is useful to first understand how the DMAC normally interfaces to an MPU. In this situation, when the DMAC wants to perform a transfer on the bus, it issues an HRQ to the MPU. The MPU will complete whatever instruction it is currently executing, it will then release the bus, and it will indicate this release to the DMAC by activating its HLDA signal. Once the DMAC sees the HLDA, it proceeds to drive the bus, and signals this by activating its AEN signal. Devices attached to the bus can then differentiate between MPU and DMAC transfers by inspection of the AEN signal. When the DMAC has finished with its transfer, it will release the bus and the AEN and HRQ signals. When the MPU sees the HRQ signal release, it will release its HLDA and will then drive the bus again. Because of this hardware handshaking between the MPU and the DMAC, the DMA transfers can occur in between the MPU instruction executions; this technique is sometimes called "cycle stealing" because the DMAC is effectively stealing bus cycles from the MPU.

Since the MCU has no HRQ or HLDA signals, the handshaking described above needs to be accomplished in software on this device, which is why it is important for the programmer to fully understand the operation of the traditional hardware approach. The key to the software approach is the **BUS\_REQ** register. Setting the **BUS\_REQ\_HLDA** bit in this register enables a hardware gate that issues an HLDA signal to the DMAC whenever the DMAC issues an HRQ; in this state, the DMAC will obtain ownership of the bus "immediately" whenever it asks for it. If the MCU clears the **BUS\_REQ\_HLDA** bit, then the DMAC will complete its current transfer, after which it will release the bus and deactivate AEN. Since the MCU can read the state of AEN in the **BUS\_REQ\_AEN** bit, it can use this bit to tell when the DMAC has actually released the bus and the MCU then owns it.

There are **IsaAcquire()** and **IsaRelease()** macro functions defined in **Usb97102.H** that illustrate this process; fortunately the macro's themselves are much shorter than the explanation of how they work:

```
#define IsaRelease()    BUS_REQ = BUS_REQ_HLDA_

#define IsaAcquire()   BUS_REQ = 0x00; while (BUS_REQ & BUS_REQ_AEN_)
```

[Aside: note that the trailing semicolons are intentionally omitted from the #defines so that these macros are used exactly like functions in the source code that calls them. Also note that, since they are macro's, there is no issue of reentrancy should it be desired to use them in multiple threads e.g., foreground and an ISR.]

Note that the **IsaRelease()** macro always executes "immediately" because all it has to do is set the HLDA bit. However, the **IsaAcquire()** macro is another story. In the same way that an MPU cannot release a bus in the middle of an instruction, a DMAC cannot release a bus in the middle of a transfer. As a result, the **IsaAcquire()** macro can take an appreciable amount of time to execute, depending on what the DMAC is doing at the time; this is described in detail in the following sections.

It is important to realize that all DMA transfers are suspended for the duration of the time that the MCU owns the ISA bus. This is unlike the cycle-stealing approach in the PC in which, while the MPU is setting up one DMA channel for a future session, the other channels can continue to transfer concurrently. As a result of this, it is important for the MCU to acquire the bus as infrequently as possible, and to retain ownership for the shortest amount of time possible. As part of this, it is usually a good idea to disable IRQ's before acquiring ownership and to not reenale them until the bus has been released; failure to do so can cause the MCU to retain ownership for an extended period of time if substantial time is spent executing ISR code during the ownership interval. For example:

```
oldEA = EA;
InterruptDisable();
IsaAcquire();

/* set up the next DMA session as quickly as possible */
```

```
IsaRelease();  
EA = oldEA;
```

This type of code sequence was seen extensively in Chapter 3, and this is the reason for it.

## DMA Channels

The DMAC contains 4 independent DMA channels. Each DMA channel can be individually programmed for Mode and transfer Type, as described in the following sections, and can be individually enabled. The exception to this is for Memory-To-Memory DMA, which always uses both DMA channels 0 and 1, as described below.

### DMA Transfer Modes

The DMAC supports three (3) different transfer modes: Single, Demand, and Block. As part of discussing these, it is important to differentiate between a DMA cycle, a DMA transfer and a DMA session: a "session" is the movement of the entire block of data programmed on a given DMA channel, and is composed of transfer(s); a transfer is an indivisible unit of data movement within which the DMAC cannot release the HRQ signal (analogous to an MPU not being able to release the bus within a single instruction) and is composed of DMA cycle(s); a DMA cycle is an individual bus cycle, which is also indivisible.

In **Single Transfer Mode**, a single data byte is transferred for each DRQ issued by the attached device. The DMAC is able to release HRQ after each individual byte transfer in response to the MCU deactivating HLDA. As a result, **IsaAcquire()** executes rapidly if all of the attached devices use this mode. However, this mode offers low performance because of the hardware overhead involved, so most high bandwidth peripherals do not use it.

In **Demand Transfer Mode**, bytes are transferred for the duration of the time that the attached device holds DRQ active. Most devices that use Demand Transfer Mode (e.g., LPT in ECP mode, audio codec, etc.) contain a counter that limits the maximum size of each burst in order to prevent the device from holding the bus so long that other devices are starved for data. Note that the DMAC cannot release the bus in the middle of a transfer, only in between transfers, so the execution time of **IsaAcquire()** becomes bounded by the largest burst size of the attached devices; this is usually on the order of 10 uS or so, which is not too bad. Examples of this type of transfer are included in the example code for this chapter.

In **Block Transfer Mode** the entire DMA session is executed in response to a single DRQ from the attached device (i.e., the entire session is treated as a single transfer), and the DMAC cannot release the bus for the duration of the transfer. In this mode, the only limit on the amount of time that the bus will be owned is the size of the session programmed by the MCU. For the USB97C102 device, the primary application for Block Transfers is when performing Memory-To-Memory DMA (discussed below) to move a packet between the MMU and external ISA RAM; since this is most common with BULK packets, whose size is limited to 64 bytes maximum, the transfer time is 64 uS at 8 MHz. It is desirable that **IsaAcquire()** not be called with IRQ's disabled while such a transfer in progress, especially if the session size is large, and techniques for avoiding this are discussed in a later section of this chapter.

### DMA Transfer Types

Although the DMAC offers others, the three (3) transfer types of greatest relevance are Memory Read, Memory Write and Memory-To-Memory.

In Memory Read and Memory Write, the transfer is between the memory and the attached device. For these transfers, the attached device in hardware generates the DRQ, and usually either Single Transfer or Demand Transfer mode is employed. As a result of this, **IsaAcquire()** executes fairly rapidly when all enabled channels are used in this way. Examples of this type of transfer are included in the example code for this chapter. If the target memory is the ISA RAM, then only channels 2 or 3 can be used, while all 4 channels are capable of device DMA with the MMU.

For Memory-To-Memory transfers, channels 0 and 1 are both used. In addition, a bit in the **DMA\_CMD** register must be set in order to establish this mode of operation. Unlike device DMA, Memory-To-Memory transfers must always use Block Transfer Mode, which means that **IsaAcquire()** can take a long time to execute. In addition, the channels should remain masked to hardware DRQ's, and a software DRQ is used instead. Examples of this type of transfer are included in the example code for this chapter.



It should be noted that Memory-To-Memory transfers have an adverse effect on performance since each individual data movement involves two back-to-back DMA cycles: the first cycle reads from the source memory into a temporary register inside the DMAC, and the second cycle writes to the destination memory from the temporary register. In addition, there is usually a transfer either to or from a Device as the ultimate source or sink of the data, so there is a total of 3 DMA bus cycles involved in each byte movement. As a result of this, Memory-To-Memory DMA as a method of performing scatter/gather should not be done for very high bandwidth devices. However, it can be extremely useful for certain types of devices, a prime example of which is an FDC, in which it is desired to be able to read or write an entire track of the media in order to obtain the best performance from the physical (i.e., mechanical) device; since the sustained throughput is fairly low, the relative inefficiency of the Memory-To-Memory DMA is not an issue, and the overall performance is increased due to the improved utilization of the mechanical device.

### **BUS\_REQ\_INH\_TCx:**

Normally, the DMAC activates the TC (Terminal Count) signal during the final transfer of a DMA session; this informs the attached device that this is the final transfer. In the USB97C102 device, firmware can control the gating of the TC signal for each individual DMA channel in order to prevent a device from seeing the TC signal. This provides a mechanism for doing Scatter-Gather DMA in software. Since disabling the TC to a device will normally prevent that device from issuing an IRQ, the DMAC has the ability to issue an IRQ as a result of a channel TC and/or DRQ, which is globally enabled and disabled using the **INTO\_ISADMA\_** bit in the **IMR\_0** register; individual DMA channel DRQ and/or TC IRQ's are enabled and disabled using the bits in the **BUS\_MASK** register. Of course, TC can be polled, as previously described, instead of interrupt driven, if desired.

As an example, suppose that a number of small individual packets (e.g., 64 byte BULK packets) arrives on the USB, and that it is desired to create the illusion of a single large contiguous DMA buffer and session when transferring the packet contents to a device. Due to the way the MMU allocates packet numbers, the organization of the DMAC address space for the packets, and the 8-byte packet headers prepended by the SIE, it is never possible for the payload data of the packets to be physically contiguous. However, by setting up multiple DMA sessions, one for each packet, and masking the TC to the device for all but the final session, from the perspective of the device, the multiple small DMA sessions appear to be one large DMA session.

### **DMA\_STS and BUS\_STAT\_CHxTC:**

The DMAC contains a **DMA\_STS** register that, among other things, contains a bit indicating the TC status of each channel. These bits can be useful in order to determine whether a given channel has completed the programmed DMA session or not. It is important to note that these bits **CLEAR ON READ**, so if this technique is used for multiple channels, then it will be necessary to shadow these bits in software. A suggested technique is to make a function e.g., **BYTE DmaGetSts()** which reads from the physical **DMA\_STS** register and OR's the contents into a shadow byte, and returns the final shadow byte. A companion function e.g., **DmaClearTC()** can do the same, but will also clear the corresponding bit in the shadow copy; this same functionality also needs to occur each time a DMA session is setup on a channel. Of course, every experienced firmware author has developed a favorite set of techniques for dealing with such hardware, and any solution is valid that provides a properly working result.

However, since the DMAC is mapped in the ISA address space, accessing the **DMA\_STS** register requires that the MCU must first acquire ISA bus ownership, which defeats one of its best possible uses -- to determine if a DMA session is done BEFORE acquiring the ISA bus, not afterwards. However, the **DMA\_STS** register is shadowed in the **BUS\_STAT** register, which is mapped directly in the XDATA space, so the MCU can read this shadow register without acquiring the ISA bus. In addition, since this is a shadow register, and not the physical **DMA\_STS** register, reads of this register are non-destructive i.e., reading the shadow copy clears no bits. This permits firmware to rapidly determine if the DMA session on any combination of channels has completed or not without having to acquire the ISA bus in order to do it. It is important to note that the physical **DMA\_STS** register must be read during session setup in order for the shadow register to be useful and that, since the shadow register is backing the physical **DMA\_STS** register with the clear-on-read characteristic, the **BUS\_STAT** register must be shadowed in software as well if the TC from multiple DMA channels is to be checked.

### **DMAC Address Space:**

The DMAC is capable of addressing a 64 KB address space. As is shown in the USB97C102 Data Sheet, the low 32 KB of this address space maps straight-through to the bottom 32 KB of the ISA Memory address space. This permits the DMAC to access an external RAM on the ISA bus, assuming that one is placed in this region. This can be useful if it is desired to transfer more data to or from a peripheral than will fit in the MMU buffers (e.g., reading/writing an entire track on an FDC, handling multiple max. size IrDA frames, etc.).

The DMAC can directly address each of the 32 packet memories in the high 32 KB of its address space. Each 1 KB block of this region corresponds directly to each of the 32 packets. Note that this tacitly places a limit of 1016 bytes on the payload data, since each packet has an 8-byte header.

## Sample Code: Chpt5.Hex

This program sends a short test page to an HP-PCL printer (e.g., HP DeskJet, HP LaserJet, or compatible) attached to the EVB using 2 different DMA techniques. If a physical printer is not available, it is possible to create a "NULL Printer" by jumping pins 11 (BUSY) and 24 (GND) on a male DB25 connector, or making the equivalent connection on the EVB. The program makes use of DMA channel 3 for the LPT DMA, so the DMA3 jumpers for the SIO device must be installed on the EVB; for Model 6126 Rev. B, these are at JP23 and JP24. The program displays its progress on the COM2 port, with the same arrangement as **Chpt3.Hex**.

There is a set of 3 BYTE\_C strings that contain a Prolog to be sent to the printer before the actual message text, an actual message, and an Epilog to be sent after the message to tell the printer to render the page. These are contained in **cProlog[]**, **cMsg[]** and **cEpilog[]** respectively. The **MakePkt()** function allocates 3 packet memories and copies these strings into them, using techniques discussed in Chapter 4. Although it is not necessary for this example, the packets are treated as having an 8-byte header in the same way as they would if they were to be transmitted or received on the USB. The companion function **FreePkt()** is used to free the packet memories at the end of the program.

Although the details of ECP operation are beyond the scope of this document, suffice it to say that there is a mode called "Parallel Port Fifo Mode" (mode 010) in which data is sent using DMA with standard Centronics handshakes; flow control uses the BUSY signal only. **SioInIt()** in **Sio.C** configures the LPT device for ECP operation. Setting the ECR for 0x40 sets mode 010 with DMA disabled, and setting the ECR to 0x48 sets the same mode with DMA enabled. The interested reader is referred to the SIO Data Sheet for a more complete discussion.

The purpose of this code example is to focus on the issues related to DMA in general, rather than ECP in particular. In order to reduce ECP related clutter, things like manipulations of the ECR occur in functions that are separated from the DMA related functions, which would hardly be considered best practice for an actual ECP application. There are numerous comments in the code suggesting changes to be made in order to improve performance for an actual ECP application.

### DMA Direct From MMU to LPT:

In the first DMA method, the function **LptSendPkt()** is used to send each packet directly from the MMU to the LPT. The function begins by calling **LptStopDma()** to disable DMA in the ECP device while a new DMA session is being set up. This function sets the ECR to 0x40, with the usual IRQ disable/restore, **IsaAcquire()**, etc. **DmaPktWithDev()** (discussed below) is then called to setup and start the DMA session. **LptStartDma()** is then called to enable DMA in the ECP device by setting the ECR to 0x48, and the function then waits for the TC in the **BUS\_STAT** register before returning, which indicates that the DMA session has completed. Note that the ECR must be re-written to 0x48 each time it sees a TC, which happens on every DMA session in this example. If the TC were masked in the **BUS\_REQ** register, then the device would never see any TC's, so the code to write the ECR each session could be removed, which would reduce the execution time of the software. In this example, the TC mask for **BUS\_STAT** is the manifest constant **cLPT\_TC\_MASK**, which is derived from the **LPT\_DMA** constant in **SIO.H**. The use of constants in this way is good for execution speed, since everything is calculated at compile time, rather than at run time. The MCU is also quite efficient at testing for a single bit set using a JB instruction, as was discussed in Chapter 2.

**DmaPktWithDev()** is a general purpose function that will perform a DMA transfer between any packet memory on any DMA channel using any Mode and Transfer Type in either direction (i.e., memory read or write). As such, it should prove to be a useful point of departure for any particular DMA application.

The function begins by reading the packet size from the packet header; it then validates the size by making sure that there is at least 1 payload data byte; it does not check for a maximum of 1016 bytes. Note that in many applications the packet size is limited to 64 payload data bytes (e.g., USB BULK packets), and in many cases the packet size is already known to be valid before being passed to a DMA setup function, so this code could be made more efficient by doing the size check on a BYTE basis, or skipping it altogether. The size is next adjusted for the DMAC by subtracting 9, which accounts for the 8 byte header and the fact that the DMAC requires that the session size be the number of bytes **MINUS 1**.

The key to the flexibility about the arbitrary channel, mode, transfer type etc. is the **dmaChMode** argument that is passed -- this is the value that will ultimately be placed in the **DMA\_MODE** register, which contains all of this information. In order to be able to determine which address and count registers to use, the **dmaCh** portion of the byte is masked off, and a (volatile) pointer (**pDmaReg**) is initialized to point to the address register of the specified channel. Looking at **Usb97102.H**, it can be seen that the **DMA\_ADDRx** and **DMA\_CNTx** registers are arranged in consecutive order starting at 0x4000, which is why **pDmaReg** is initialized the way that it is. In the example code, the **dmaChMode** value is obtained from **cLPT\_DMA\_CH\_MODE**, which is based on manifest constants in **Usb97102.H** and **SIO.H**.

The IRQ state is then saved and disabled, and the ISA bus acquired as usual. In order to be able to access any of the DMAC registers, the **IOBASE** register is set to **DMA\_IOBASE**. Since DMA channels 0 and 1 can be used for both Memory-To-Memory and device DMA in the same application, a check is made to see if the current channel is either of these and, if it is, then the **DMA\_CMD** register is written to clear Memory-To-Memory mode. Note that this could be skipped in an application if this is not relevant.

In general, it is a good idea to mask any channel while programming it, and then to unmask it when finished. This is accomplished by writing the **DMA\_MASK** register, which is done next. Next, the **DmaClearByteFF()** macro function is called in order to clear the byte pointer flip-flop. This flip-flop controls the reading/writing of WORD registers in the DMAC -- the first access after clearing the flip-flop is the low byte, followed by the high byte. For this reason, if DMA is ever done in an ISR as well as in the foreground, the foreground thread **MUST** disable IRQ's, at least while programming any WORD registers in order to avoid having the ISR upset the state of this flip-flop. The byte flip-flop toggles after each byte write so, if consecutive word registers are programmed, it is not necessary to re-clear the flip-flop before each word.

The starting address of the session is then written to the appropriate **DMA\_ADDRx** through **pDmaReg**. Of course, this pointer could have been dereferenced as an array, but that results in slower code execution, so the pointer notation was used. Note that the address value is 8 in the low byte to skip past the packet header. The high byte of the address has the MSB set to 1 in order to select the upper 32 KB of the DMAC address space, which is where the packet memories are mapped, and the packet number is shifted right by 2 bits in order to take into account the 1 KB packet locations. The **pDmaReg** pointer is post-incremented after writing the high address byte so that it points to the appropriate **DMA\_CNTx** register, and the 2 bytes of count are written.

Next, the **DMA\_MODE** register is written to set the transfer mode, direction, etc. Next, the **DMA\_STS** register is read in order to clear any old TC's that might be left over from a previous session, since this code will detect the end of the session by polling for the TC in the **BUS\_STAT** shadow register. Finally, the DMA channel is unmasked to enable it, the ISA bus is released to the DMAC, and the IRQ enable state is restored.

In USB DMA applications that use BULK packets, a DMA session must be set up for each packet, so the efficiency of the code that does this is usually of great concern. As a result, it is suggested that **every reasonable effort be made to streamline this code as much as possible**. For example, do not waste time checking the packet size if it is already known to be valid; do not use a WORD if a BYTE will do. Do not waste time setting the **DMA\_MODE** each session if it is the same every time; do it at initialization instead. Do not pass the **dmaChMode** at all, but have a separate dedicated function for each DMA device; doing so also eliminates the need to use a pointer for register access, which also improves performance. Consider masking TC's in order to avoid having to set e.g., the **ECR** each time a session is set up. If it is necessary to set a register like **ECR**, do it right in the dedicated DMA function just before releasing the ISA bus, rather than in a separate function where the entire **EA/BUS\_REQ** dance has to be done again.

#### **DMA From ISA RAM to LPT:**

Although this is out of order with respect to the example code, the **DmalsaWithDev()** function is so similar to the previous function that it naturally follows here in the discussion. The primary difference is that the address and count of a buffer in ISA RAM are passed, instead of a packet number with the size contained in its header. On entry, the address and size are validated, with the criterion being that the entire session must be contained in the low 32 KB. In an actual application, this should be streamlined or removed, as was previously discussed. The only other difference is that the count is only adjusted by 1 for the DMAC, instead of 9, since there is no packet header to skip. As was previously mentioned, only DMA channels 2 and 3 are capable of using an ISA memory target for device DMA.

### **DMA From MMU to ISA RAM:**

In order to have data in ISA RAM suitable to DMA to the LPT, the function **DmaPktTolsa()** is called for each of the 3 packets to create a single DMA buffer, which is then sent to the LPT using **DmalsaWithDev()**.

As was previously mentioned, Memory-To-Memory Transfers must always make use of DMA channels 0 and 1, and must always use Block Mode. The **DMA\_CMD\_MEM2MEM\_** bit in the **DMA\_CMD** register must be set. Channel 0 address must be set for the source memory block, and Channel 1 address must be set for the destination. Channel 1 Count must be set for the session size, and it is Channel 1 that will TC when the session is completed. Both channels must be set for Block mode. Both channels should be masked before programming the session, and remain masked when complete. The session is actually started by issuing a "Software DRQ", as opposed to the hardware DRQ used with peripheral devices, on Channel 0 using the **DMA\_REQ** register. As usual, the IRQ enable state is saved/restored, the ISA bus acquired/released, and the **IOBASE** set to **DMA\_IOBASE** to make the DMAC addressable.

### **SCATTER / GATHER DMA**

The SGDMA performs scattering / gathering operations from the MMU to or from the external memory, as well as ISA device transfers to or from MMU. The SGDMA has four DMA channels with each channel having its own set of registers. Each of the four- independent DMA channels may set up to 16 transfers, which can be programmed to occur consecutively. The SGDMA will run the internal 8237-DMA controller alone, once the MCU indicates which packet to transfer. This allows the MCU to handle other operations, thus increasing overall performance. The SGDMA also contains a single PIO engine that permits the MCU to access the ISA bus on a cycle stealing basis with the DMA transfers.

SGDMA memory-to-memory transfer is a special case since both channel 0 and 1 must always be used. The source must be channel 0; the destination must be channel 1. SGDMA only supports memory-to-memory transfers between MMU memory and ISA memory, in either direction (**MEM\_OP** bit in the SGDMA command register indicates whether is MMU or ISA memory operation).

### **MMU Memory to ISA Memory Transfers**

To perform a memory transfer between MMU memory to ISA memory, there are two critical functions that the MCU must perform in order for the transfer to take place. The MCU must clear channel 0 in the SGDMA transfer size high and low registers, as well as to set channel one in the SGDMA ISA Address High and Low Byte registers to the ISA address to be used. Just remember that to write to these registers the corresponding channel must be disabled (**CHANNEL\_ENABLED=0** in register **SGDMA\_CMDx**) as well as no SGDMA transfer operation in progress while trying to attempt a transfer. In addition, the SGDMA will add the size of the completed transfer to channel one in the SGDMA ISA Address High and Low Byte registers after each terminal count (TC). If **PKT\_HDR** bit in channel 0 of the SGDMA Command Register is set to one, the transfer size comes from the MMU packet header. If **PKT\_HDR** bit is cleared, the transfer size is the value in channel one of SGDMA transfer size high and low registers. This only applies if **MEM\_OP** bit in SGDMA Command register is set.

### **ISA Memory to MMU Memory Transfers**

To perform a memory transfer between ISA memory to MMU memory, there are three critical functions that the MCU must perform in order for the transfer to take place. The MCU must set channel one in the SGDMA ISA Address High and Low Byte registers to the ISA address to be used. Next, it must set channel 0 in the SGDMA Transfer Size High and Low Registers to the ISA buffer size, and finally it must set channel 1 in the SGDMA Transfer Size High and Low Registers to the session transfer size. Unlike the MMU to ISA memory transfer, after each TC, the SGDMA will add the session transfer size to channel 0 of the SGDMA ISA Address High and Low Byte registers and subtract the session transfer size from channel 0 of the SGDMA Transfer Size High and Low Registers. The actual transfer size is the lesser of the values in channel 0 and 1 of the SGDMA Transfer Size High and Low Registers. If **PKT\_HDR** bit in channel 1 of the SGDMA Command Register is set to one, then the actual transfer size plus 8 will be written to the packet memory at an offset of 6. **M2M\_INCOMPLETE** bit in channel 0 of the SGDMA Status Register will indicate whether a complete ISA to MMU memory transfer has performed. **M2M\_INCOMPLETE** bit in channel 0 of the SGDMA Status Register will only be set when channel 0 of the SGDMA Transfer Size High and Low Registers reach a value of 0, indicating the ISA buffer has been completely transferred, and channel 1 of the SGDMA Packet Number Start FIFO Register is empty.

In addition, the SGDMA will add the size of the completed transfer to channel one in the SGDMA ISA Address High and Low Byte registers after each terminal count (TC). If **PKT\_HDR** bit in channel 0 of the SGDMA Command Register is set to one, the transfer size comes from the MMU packet header. If **PKT\_HDR** bit is cleared, the transfer

size is the value in channel one of SGDMA transfer size high and low registers. This only applies if MEM\_OP bit in SGDMA Command register is set.

The following is an example on how to program the DMA registers and preload the SGDMA Start FIFO so that when the time comes to receive data, all the function needs to do is release the ISA bus, and the DMA will commence.

Argument: DataIn

If TRUE, program for DMA from DEVICE to HOST (DataIn)

If FALSE, program for DMA from HOST to DEVICE (DataOut)

**static void ProgramSgDma(uint8 DataIn) reentrant**

```
{
    uint8 PacketNumber;
    uint8 x;
    uint8 TmpIoBase = IOBASE ;

    /* ----- Load up the SGDMA FIFO ----- */

    /* Acquire the ISA bus and DMA transfer will be held off until the bus is released */

    intrpt_disable() ;
    BUS_REQ &= ~BUS_REQ_HLDA_ ;
    while (BUS_REQ & BUS_REQ_AEN_ ) ;
    intrpt_enable() ;

    if ( DataIn )
    {

        /* Lock & Load SGDMA FIFO */

        for (x=0;x<MAX_DMA_PACKETS;x++)
        {

            /* Allocate a packet */

            MMUCR = MMUCR_ALLOCATE_ ;
            while (ARR&ARR_FAILED_ ) ;
            PacketNumber = ARR&ARR_PN_MASK_ ;

            /* Load packet into the SGDMA FIFO */

            SGDMA_START_FIFO3 = PacketNumber ;

            /* Put the endpoint number in the first byte of the header */

            PNR = PacketNumber;
            PRL = 0;
            PRH = PRH_PNR_ | PRH_WRITE_ | PRH_AUTO_INCR_ ;
            MMU_DATA = BULK_IN_ENDPOINT;

            /* Put the length in the sixth byte of the header */

            PRL = 6;
            PRH = PRH_PNR_ | PRH_WRITE_ | PRH_AUTO_INCR_ ;
            MMU_DATA = PKTHDRSZ + MAXPKTSZ;
            MMU_DATA = 0;
        }
    }
}
```

```

    }

}

/* ----- Program the 8237 DMA controller ----- */

IOBASE = DMA_IOBASE ;
DMA_MSTR_CLR = 0xFF ;

/* Read the DMA_STS Register to clear any exiting TC's */

DMA_CMD = DMA_STS ;

/* Set the command register to go as fast as possible */

DMA_CMD = DMA_CMD_COMP_TIME_ ;

/* Set the DMA Mode:
   DMA_DEMAND_MODE_ -> DREQ is used by device to regulate the transfer
   if DataIn, DMA_WRITE_XFER_ -> Source = I/O, Destination = Memory
   if DataOut, DMA_READ_XFER_ -> Source = Memory, Destination = I/O
   DMA_CH3_ -> Use DMA channel 3 */

if ( DataIn )
    DMA_MODE = DMA_DEMAND_MODE_ | DMA_WRITE_XFER_ | DMA_CH3_ ;
else
    DMA_MODE = DMA_DEMAND_MODE_ | DMA_READ_XFER_ | DMA_CH3_ ;

/* ----- Program the SGDMA controller ----- */

/* Disable the SGDMA channel before programming it */

SGDMA_CMD3 = ~SGDMA_ENABLE_ ;

/* Program the SGDMA Command register
   * SGDMA_PKT_HDR_ -> Destination is a packet in the MMU with a packet header
   * SGDMA_MEMOP_ -> Starting address & length per packet is determined by the packet & its header */

SGDMA_CMD3 = SGDMA_PKT_HDR_ | SGDMA_MEMOP_ ;

/* Set the SGDMA size registers. According to the 102 spec, these are not used when the channel is set with
   SGDMA_PKT_HDR_ */

SGDMA_SZLO3 = 0x40 ;
SGDMA_SZHI3 = 0x00 ;

/* Enable the SGDMA channel, the first packet in the SGDMA Start FIFO is loaded into the SGDMA "chamber"
   ready to go as soon as the bus is released and the device starts robbing DREQ */

SGDMA_CMD3 = SGDMA_ENABLE_ | SGDMA_PKT_HDR_ | SGDMA_MEMOP_ ;

/* ----- Set up end-of-transfer conditions ----- */

/* Gate Channel3 TC off the isa bus for scattering */

```

```

intrpt_disable() ;
IMR_0 &= ~(INT0_ISADMA_) ;
BUS_MASK = BUS_STAT_CH3RQ_ ;
BUS_REQ |= BUS_REQ_INH_TC3_ ;

```

```
intrpt_enable() ;
```

*/\* don't release the isa bus ... DMA is now armed. When ready to start DMA, release the bus\*/*

*/\*restore the callers IOBASE \*/*

```

IOBASE = TmpIoBase ;
}

```

The following example will demonstrate a SGDMA flush after end of a data transfer. After the device has interrupted us, indicating it is done with the DMA transfer, the flush process begins.

*/\* flush the SGDMA done fifo to get all packets out to the host\*/*

```

while (!(SGDMA_STS3 & SGDMA_DONE_FIFO_EMPTY_))
{

```

*/\* RX Book keeping \*/*

```

    pnr = SGDMA_DONE_FIFO3 ;
    Wrapper.CSW.DataResidue -= GetPacketSize(pnr) ;
    mmu_deallocate(pnr) ;

```

```
}
```

*/\* flush SGDMA START by manually starting/stopping xfers \*/*

```

while (!(SGDMA_STS3 & SGDMA_START_FIFO_EMPTY_))
{
    SGDMA_CMD3 = SGDMA_PKT_HDR_ | SGDMA_MEMOP_ ;
    SGDMA_CMD3 = SGDMA_PKT_HDR_ | SGDMA_MEMOP_ | SGDMA_ENABLE_ ;
}

```

*/\* jog the last one over onto the done fifo \*/*

```

SGDMA_CMD3 = SGDMA_PKT_HDR_ | SGDMA_MEMOP_ ;
SGDMA_CMD3 = SGDMA_PKT_HDR_ | SGDMA_MEMOP_ | SGDMA_ENABLE_ ;
SGDMA_CMD3 = SGDMA_PKT_HDR_ | SGDMA_MEMOP_ ;

```

*/\* all packets still in play are now in the SGDMA done fifo... flush it \*/*

```

RequestLength = Wrapper.CSW.DataResidue ;
AvailableLength = 0 ;

```

```

while (!(SGDMA_STS3&SGDMA_DONE_FIFO_EMPTY_))
{
    pnr = SGDMA_DONE_FIFO3 ;
    RequestLength -= GetPacketSize(pnr) ;
    mmu_deallocate(pnr) ;
}

```

```

intrpt_disable() ;

/* Clear out the pending DMA & USB TX interrupts */

flg = ISR_0;
flg &= ~(INT0_ISADMA_ | INT0_RX_PKT_);
ISR_0 = flg;

/* Read to disable any pending ISADMA irqs */

flg = BUS_STAT ;

/* Re-enable the interrupts */

msk = IMR_0 ;
msk &= ~(INT0_ISADMA_ | INT0_RX_PKT_);
IMR_0 = msk;

/* Reacquire the ISABUS */

BUS_REQ &= ~BUS_REQ_HLDA_ ;
while (BUS_REQ & BUS_REQ_AEN_);
intrpt_enable() ;
}

```



## CHAPTER 6 - GETTING ON THE BUS

This chapter describes everything that is needed to implement a fully Chapter 9 compliant USB Device using the USB97C102. In addition to the USB Specification, this text makes frequent reference to the example code, which is contained in the following files:

**Chpt6.C** -- application code  
**Chpt6.H** -- application data structures, especially USB descriptors  
**UsbCore.C** -- core USB code  
**UsbCore.H** -- interface between Core code and application code  
**Usb.H** -- USB related data structures (descriptors, Setup packets, etc.)

The example code implements a complete USB device, which passes UsbCheck V2.6 for Chapter 9 compliance. The device is limited to a single USB Configuration, but supports an arbitrary number of Interfaces, as defined by **cUSB\_NUM\_IFs** in **Chpt6.H**. Each Interface supports an arbitrary number of alternate settings, as defined by **cUsbNumAltIFs []** in **Chpt6.c**. Each alternate setting contains an arbitrary combination of endpoints, as defined by **usbTotalCfgDesc** in **Chpt6.h**. The basic architecture is that all application dependent code is contained in **Chpt6.C** and **Chpt6.H**, while all USB related code is contained in the other files, which act as a Kernel or Core to which new applications can be linked. In the case of most EP0 Transfers, application dependent information is obtained by having the Core code perform a function call into the Application code in order to retrieve the information; a description of each individual Transfer is included in a following section.

A number of other application dependent values are defined in Chpt6.X:

- The choice of MCU and DMA clocks; GPIO function, direction and value; MEM\_BANK setting, etc. (using our old friend `Usb97C102HwInit()`)
- All USB String Descriptors
- USB VID, PID, Class, Subclass, Protocol, etc.
- EP0 max. packet size

In order to execute this code on the EVB hardware, the jumpers must be set as follows. **On Assy 6126 Rev. B, set JP29 3-5 and 4-6, JP19 2-3, J27 inserted.**

### Device States:

For a more complete description of USB Device States, see section 9.1.1 in the USB Specification.

When a device is first attached to the USB after POR, it must be dormant to all bus activity, including Address0/EP0. In **UsbCore.H**, this is the **DEVICE\_DEFAULT** State.

After it receives its first `UsbReset`, it must respond at Address0/EP0 only. **UsbCore.H** does not assign a new value to the Software State at this point because it is not needed for anything. In this State, the Host can query the device for its USB Descriptors, etc. and will eventually send it a `SET_ADDRESS` command, at which point the device must move from Address0 to whichever USB Address the Host assigns; this is the **DEVICE\_ADDRESS** State.

Eventually, the Host will send a `SET_CONFIGURATION` command; a Configuration of zero is the Unconfigured State, which means that only EP0 remains active. A non-zero Configuration is a Running State; this is the **DEVICE\_CONFIGURED** State. For the most part, device software is really only concerned with whether the device is in CFG1 or not, since software behavior is mostly the same for all of the other States.

From any State it is possible for the Host to suspend the device if it is bus powered; this is the **DEVICE\_SUSPENDED** State. While in Suspend, it is possible for the Host to Resume the device and, for devices that support the feature, it is possible for the device to Remote Wakeup the Host. Although the example code does not implement these features, the following sections describe the techniques involved.

### USB Suspend/Resume:

The Host can instruct a USB device to Suspend by stopping all traffic (i.e., establishing a J state) to the device for a period of at least 3 mS. The Host can then Resume the device through a variety of mechanisms, each of which involves a transition out of the J state. The SIE provides the ability for the software to detect when SOF's arrive, either by triggering an MCU timer from the SOF (using the **UTIL\_CONFIG** register bits), or by issuing an IRQ when the SOF arrives (using the **INT1\_SOF\_** bit in **IMR\_1**), or by issuing an IRQ when bus is idle for 3ms (using the

**SUSPEND\_** bit in the **WU\_SRC\_1**), and either triggering a timer or reading a free running timer, etc. From any of these techniques, the software can determine when the device is supposed to enter a Suspend State.

Entering the Suspend State consists of first suspending the SIE, which places it into a low power state by stopping its clock; this is set in the **SIE\_CONFIG** register. The next step is to suspend the DMAC by stopping its clock, switching the MCU to its ring oscillator and stopping the MCU clock, and enabling the ring oscillator to be stopped with the **PCON** register LSB; all of these are accomplished in the **CLOCK\_SEL** register. Finally, setting the LSB in the **PCON** register high stops the MCU ring oscillator. The following code sequence illustrates the technique of Suspending:

```
SIE_CONFIG |= SIE_CONFIG_SIE_SUSPEND_;
CLOCK_SEL = CLKSEL_SLEEP_ | CLKSEL_ROSC_EN_;
PCON |= 0X01;
```

When the Resume signaling arrives, the MCU will begin executing its **Isr2()** function. In that function, the **WK1\_RESUME\_** bit in the **WU\_SRC\_1** register must be cleared by writing a “1” to the corresponding bit in **WU\_SRC\_1** register (otherwise the ISR will keep getting re-entered, since the IRQ condition has not been cleared). It is also the author’s preference to restore the **CLOCK\_SEL** register in the ISR, but that could be done in the foreground instead if preferred. After returning from the ISR, execution will continue in the foreground thread at the instruction immediately after the **PCON** register LSB was set to 1. Following is some example **Isr2()** code:

```
ISR = ISR_1;
ISR_1 = 0xFF; /* clear bits in read */

if (ISR & INT1_PWR_MNG_) {
    SRC = WU_SRC_1;
    WU_SRC_1 = 0xFF; /* clear bits in read */

    if (SRC & WK1_RESUME_) { /* just like in Usb97c102HwInit() */
        CLOCK_SEL = (BYTE)(CLKSEL_ROSC_EN_ | cClocks);
        CLOCK_SEL |= CLKSEL_MCUCLK_SRC_;
        CLOCK_SEL &= ~CLKSEL_ROSC_EN_;
    }
}
```

In order to enable the wake from Resume event, it is necessary to unmask the corresponding bit in the **WU\_MSK\_1** register; it is convenient to do this early in **\_main()** when the other interrupt related masks are configured. **Note that if this is not done, then the device will get stuck in the Suspend State!**

```
EX1 = 1;
IMR_1 = (BYTE)~(INT1_PWR_MNG_);
WU_MSK_1 = (BYTE)~(WK1_RESUME_ | WK1_USB_RESET_);
```

### USB Remote Wakeup

Implementing this feature involves doing the Suspend/Resume activity above but, in addition to enabling wakeup from USB Resume or USB Reset, a wakeup is also enabled from one or more GIRQ signals using the bits in the **WU\_SRC\_2** register. When the external device asserts the corresponding GIRQ signal, the MCU wakes up with the same type of **PWR\_MNG** IRQ as for a USB Resume; of course, the software can determine the source of the wakeup event by reading the **WU\_SRC\_X** registers. The Resume code sequence is similar, except that the software must also cause the SIE to issue Resume signaling to the Host by setting the **SIE\_CONFIG\_USB\_RESUME\_** bit in the **SIE\_CONFIG** register.

It is common that the same peripheral device and GIRQ used for Remote Wakeup is also used for normal operation. When this happens, the **WU\_SRC\_2** register should be unmasked to enable the wakeup just before Suspending, and it should be masked again as part of the Resume. The opposite procedure should be done with the corresponding bit in the **IMR\_0** register if the peripheral is interrupt driven while the device is not Suspended.

## EPCTRL Registers

The USB97C102 contains a separate EPCTRL Register for each USB Endpoint. This set of registers permits software to define the state of each RxEP and TxEP as being either Disabled, Enabled, Busy or Stalled. In addition, each EP can be defined as being Isochronous or not. The behavior of an EP in each of these modes is dependent upon the state of the EP and the type of USB traffic addressed to it. The following discussion of behavior is for the case in which the MMU and RXFIFO are both not full. Note that allowing either of them to fill should not be permitted in any USB application, since every USB device must be capable of receiving a Setup Packet at any time (discussed in a later section), which is not possible if either the MMU or RXFIFO is full.

### Non-ISO OUT EP's

**RX\_ENABLE\_:** Both SETUP and OUT packets are received, regardless of CRC or Data Toggle. For bad CRC, there is no handshake, otherwise the handshake is ACK. It is the responsibility of firmware to discard packets with bad CRC or Data Toggle. In addition, the **EPCTRL\_RX\_TOGGLE\_** bit is read-only, so it is the responsibility of software to maintain a data toggle bit to be used for rejecting duplicate packets on Bulk, IRQ and Control EP's. The requirements for initializing the toggle bit vary with the EP type and are discussed in the following sections.

**RX\_STALL\_:** For OUT packets, same as **RX\_ENABLE\_** except that STALL handshake is sent instead of ACK. It is the responsibility of firmware to discard the packet. For SETUP packets, the packet is received and NO handshake is sent; it is the responsibility of firmware to discard the packet and clear the **RX\_STALL\_** condition so that the retransmitted Setup packet will subsequently be received and ACK'd.

**RX\_BUSY\_:** For OUT packets, the packet is not received and the handshake is NAK. For SETUP packets, the same situation as **RX\_STALL\_** above.

**RX\_DISABLE\_:** The EP is completely disabled; no packets are received and no handshakes are sent.

### Non-ISO IN EP's:

**TX\_ENABLE\_:** when the TxFIFO for the EP is empty, IN tokens are responded to with a NAK handshake; if the TxFIFO is not empty, the packet is sent. The transmission is only considered complete when the ACK handshake from the Host is received, so a Host no handshake (from a bad CRC at the Host end) or a dropped ACK handshake will result in automatic retransmission in response to subsequent IN tokens from the Host. The **EPCTRL\_TX\_TOGGLE\_** bits are writable, and define the DATAx PID to be used on the next transmission from the corresponding EP. The DATA PID toggles automatically with subsequent transmissions in response to ACK handshakes from the Host. It is the responsibility of software to initialize the **EPCTRL\_TX\_TOGGLE\_** bit appropriately for the EP type (Bulk, IRQ or Control) and state, which is described in a following section. Note that, since RX flow control bits are in the same register as the writable Tx data toggle bit, **Tx EP's cannot be used at the same EP address as an RxEP** that will require concurrent flow control (i.e., BULK; IRQ, ISO, and Control EPs are OK) because doing so would overwrite the TxToggle in the course of doing the read-modify-write for Rx flow control. However, this restriction still provides a minimum of 15 (typically more) unidirectional pipes in addition to the Default Pipe, which should be plenty considering that there are only 32 pages of Packet Memory to share between all of the pipes anyway.

**TX\_STALL\_:** IN tokens get a STALL handshake; no packet is sent even if the TxFIFO is not empty.

**TX\_BUSY\_:** IN tokens get a NAK handshake; no packet is sent even if the TxFIFO is not empty.

**TX\_DISABLE:** the EP is completely disabled; no packets are transmitted and no handshakes are sent.

### ISO EP's

Software should only set the **xx\_ENABLE\_** and **xx\_DISABLE\_** values. Disabled EP's do not send or receive any packets. ISO EP's never send any handshakes.

Rx packets are received regardless of CRC errors, DATA PID or EP STALL condition. It is the responsibility of software to discard packets for CRC or STALL, but to mark their time. According to the USB Specification, DATA PID should be ignored on ISO Rx packets.

### Endpoint Command Register:

The endpoint command register allows the dynamic modification and configuration of specific endpoints. This register allows the MCU to write to each individual bit field within the existing register endpoint set without having to read, modify, and write operations.

The TX/RX bit on Endpoint Command Register will allow the command specified in bits 6-4 of same register to control the TX endpoint when set to one. When this bit is cleared, the command will control the RX endpoint. In other words, if bit is set to one, the command will control all TX signals described above on EPCTRL register. When bit is cleared the command will control all RX signals described above on EPCTRL register. The lower four bits 3-0 of the eight bit Endpoint command register are used to select desired the endpoint 0 thru 15 (0000 endpoint 0, 0001 endpoint 1, 0010 endpoint 2, etc).

The following is an example on how to define endpoints using the Endpoint Command Register as described above.

```
#define _endpoint_set_tx_toggle(_mcr_ndp);
{
    _intrpt_disable();
    x_epcmd = kbm_epcmd_tx | kbm_epcmd_settog | (_mcr_ndp);
    _intrpt_restore();
}

#define _endpoint_clr_tx_toggle(_mcr_ndp);
{
    _intrpt_disable();
    x_epcmd = kbm_epcmd_tx | kbm_epcmd_clrtog | (_mcr_ndp);
    _intrpt_restore();
}

#define _endpoint_tx_enable(_mcr_ndp);
{
    _intrpt_disable();
    x_epcmd = kbm_epcmd_tx | kbm_epcmd_enable | (_mcr_ndp);
    _intrpt_restore();
}

#define _endpoint_tx_disable(_mcr_ndp);
{
    _intrpt_disable();
    x_epcmd = kbm_epcmd_tx | kbm_epcmd_disable | (_mcr_ndp);
    _intrpt_restore();
}

#define _endpoint_tx_busy(_mcr_ndp);
{
    _intrpt_disable();
    x_epcmd = kbm_epcmd_tx | kbm_epcmd_busy | (_mcr_ndp);
    _intrpt_restore();
}

#define _endpoint_tx_stall(_mcr_ndp);
{
```

```

    _intrpt_disable();
    x_epcmd = kbm_epcmd_tx | kbm_epcmd_stall | (_mcr_ndp);
    _intrpt_restore();
}

```

### NonControl Endpoint Register:

NonControl Endpoint Register 1 and 2 allow the designer to setup each individual endpoint as a non-control endpoint. This is needed since the USB Spec V1.1 states that “If a non-control endpoint receives a Setup PID, it must ignore the transaction and return no response”. Each bit of the NonControl Endpoint Registers will correspond to the associated endpoint. Bits 0-7 of the NonControl Endpoint 1 register will correspond to Endpoint 7-15. Bits 0-7 of the NonControl Endpoint 2 will correspond to Endpoint 0-7. The MCU will write to these registers. When setting a bit to one, the endpoint will not respond to a Setup PID meaning is a non-control endpoint. When clearing a bit, the endpoint will respond to a Setup PID meaning is a control endpoint. The following code sequence illustrates how to utilize the NonControl Endpoint Register:

The control endpoints are EP0, EP4, EP8 and EP12

```

{
    :
    :
    #define sie_init_ex();

    {
        NONCTRLEP_HI = 0xEE;
        NONCTRLEP_LO = 0xEE;
    }

    #else

    #define sie_init_ex();

    #endif
    :
    :
}

```

## RESETS

### POR

At POR time, the Core code begins execution in `_main()`, which calls `Usb97102Hwlnit()` which was described in a previous chapter. After this, `ApHwlnit()` in **Chpt6.C** is called to allow the application to initialize its hardware state; this consists of calling `Siolnit()`, which was described in a previous chapter. At this point, `DBGPRINT`'s are operational. The core code then calls `ApSwlnit()` in **Chpt6.C**, which causes the application to initialize its software state; this consists of calling `Ep1RxQuelnit()` and `Ep3RxQuelnit()` to initialize the receive queues, in addition to marking the ISA queues empty, the DMAC as unowned, and no ISR errors are pending. It should be noticed at this point that most of the data and functions in **Chpt6.C** are marked as static, since they are private to the application and should not be referenced from outside the file; by marking these items as static, the compiler's scoping rules will guarantee that no external references occur.

At this point, the core code sets up the IRQ masks, and enables the SIE and IRQ's. Finally, `ApUsbAttach()` in **Chpt6.C** is called to cause the application to attach the device to the USB; this is provided for application hardware that supports electronic attachment to the bus. For the EVB hardware, attaching to the bus is accomplished by setting the GPIO7 pin high. Since the `GPIOA_OUT` register is often shared with ISR's, the usual dance with EA is done.

Core execution continues with an infinite loop that handles USB Reset, EP0 activity, etc. as described below.

## USB Reset

When the SIE detects a USB Reset, it signifies this by activating the **SIE\_STAT\_USB\_RESET\_** bit in the **SIE\_STAT** register. This event can also issue an IRQ using the **WK1\_USB\_RESET\_** bit in the **WU\_SRC\_1** register, which in turn is enabled by the **INT1\_PWR\_MNG\_** bit in **IMR\_1**; if both of these conditions are simultaneously enabled, then a USB Reset will cause an **IRQ\_2**. The sample code enables this condition and calls back to the Application with **ApUsblsrReset()**, but the sample code does not have any work to do at interrupt time, so it just sets a **bUsbResetPending** flag and returns. The foreground code then uses this flag to call **UsbReset()** in **UsbCore.C**, which actually performs the reset activity.

**UsbReset()** sets USB Address0 and disables all EPs, before calling **ApUsbReset()** in **Chpt6.C**. **ApUsbReset()** then calls **ApSwlnit()** to initialize its software state, **DmacReset()** to initialize the DMAC, and then it resets the MMU. When the application code returns to **UsbReset()**, this function waits for the USB Reset to finish before resetting the SIE and enabling EP0. The reason for this wait is that resetting the SIE also resets the counter that it uses to detect USB Reset events; if this is done while the USB Reset is still active, the SIE will detect another USB Reset event, which will then cause a new **IRQ\_2**, etc. and the device will then get **caught in a loop** until the USB Reset is finally over. In order to avoid this condition, the code waits until the USB Reset is completed before resetting the SIE; this results in a single execution of the USB Reset software functions for each actual reset on the bus. Note that the USB Reset is required to last for a minimum of 10 mS, and the device must complete its reset processing within 10 mS after that; in practice, these timing requirements are easily satisfied without taking any special precautions, but the device programmer should certainly be aware of them anyway.

## EP0 Control Transfers

There are three (3) basic types of Control Transfers that can occur, as described in the USB Specification, section 8.5.2; also, Chapter 9 describes the Standard Requests.

### 1) Control Read:

The Host sends a Setup packet, followed by one or more IN tokens. The **wLength** field in the Setup packet contains the maximum number of bytes to return; the device must limit the transfer to the specified size in the event that the requested item is larger. During the Data Stage, the Host sends IN tokens to read the packets from the device; the first data packet has a **DATA1** PID, and subsequent data packets toggle the **DATA** PID. The device issues **NAK** handshakes during the Data Stage if it is busy. After the last data packet, the Host sends a zero-byte **OUT** packet during the Status Stage. The device handshakes this packet with either an **ACK** to indicate success, or a **STALL** for any error. If the device wishes to Stall the transfer, the first (and only) data packet should be zero length so that the Host will skip immediately to the Status Stage.

Standard Requests:

**GET\_STATUS**, **GET\_DESCRIPTOR**, **GET\_CONFIGURATION**, **GET\_INTERFACE**, **SYNCH\_FRAME** (optional)

### 2) Control Write:

The Host sends a Setup packet, followed by one or more **OUT** tokens. The **wLength** field in the Setup packet contains the total number of bytes that the Host will be sending. During the Data Stage, the Host sends **OUT** packets to the device, with the first packet having a **DATA1** PID, and subsequent packets toggle the **DATA** PID. The device issues **NAK** handshakes during the Data Stage if it is busy. If the device wishes to Stall the transfer, the Stall handshake is sent during the Data Stage in response to the first (and possibly only) **OUT** packet. To accept the Transfer, the device receives and **ACKs** the packets during the Data Stage and sends a zero-byte handshake packet during the Status Stage with a **DATA1** PID.

Standard Requests:

**SET\_DESCRIPTOR** (optional)

### 3) No-Data Control:

The Host sends a Setup packet. Since any required information is contained in the Setup packet, there is no Data Stage; this Transfer type is similar to a Control Write with zero data bytes. During the Status Stage, the Host sends IN token(s). The device sends a zero-byte packet with a **DATA1** PID to indicate success, or a Stall handshake otherwise (i.e., to Stall the transfer).

Standard Requests:

**CLEAR\_FEATURE**, **SET\_FEATURE**, **SET\_ADDRESS**, **SET\_CONFIGURATION**, **SET\_INTERFACE**

Note that when handling a SET\_ADDRESS Transfer, it is necessary to not actually assign the new address until AFTER the Status Stage has completed; this is because the Host will try to read the handshake packet at the original address.

As described in sections 8.5.2.2 and 5.5.5 of the USB Specification, there are situations in which, from the device's perspective, the sequencing in a Control Transfer appears to be out of order; this can happen especially as a result of lost handshakes going back to the Host.

### EP0 Stalls

Under a variety of circumstances, the device must Stall a Transfer to EP0. The method for doing this varies with the Transfer Type, as was described above. The previous description of the EP\_CTRL register defines the bit fields that must be set in order to Stall each transfer type, but they are summarized here for convenience:

**Control Read:** Stall the Rx and then queue a zero-length Tx packet with a DATA1 PID; reading this packet causes the Host to advance to the Status Stage. Leave the Rx Stalled until the next Setup packet arrives, and then clear the Rx Stall.

**Control Write:** Stall the Rx and leave it Stalled until the next Setup packet arrives, then clear the Rx Stall.

**No-Data Control:** Stall both the Rx and Tx and leave them Stalled until the next Setup packet arrives, then clear both Stalls.

While it may not seem necessary to Stall the Rx in all cases, or to leave it Stalled until the next Setup packet, this turns out to be true when every case of lost handshakes, etc. is considered.

### EP0 FSM

Since the transfers described above can involve multiple packets spread out over an appreciable amount of time, one suitable way to implement the software is as a Finite State Machine (FSM) that is called in a foreground-polling loop, which is how the example code in this chapter does it. States are defined as follows:

**IDLE:** The device is waiting for a Setup packet to start a transfer.  
**SETUP:** The device has received a Setup packet, but has not yet processed it.  
**RD\_DATA:** The device is in the Data Stage of a Control Read Transfer  
**WR\_DATA:** The device is in the Data State of a Control Write Transfer  
**WR\_STATUS:** The device is in the Status Stage of either a Control Write or a No-Data Control transfer.

The State transitions for each Transfer type are as follows:

CTL-READ: IDLE -> SETUP -> RD\_DATA -> IDLE  
CTL-WRITE: IDLE -> SETUP -> WR\_DATA -> WR\_STATUS -> IDLE  
NoData-CTL: IDLE -> SETUP -> WR\_STATUS -> IDLE

The code is implemented in **UsbCore.C**, using a foreground polling function **Ep0FSM()** and an RxISR called **Ep0RxIsr()**. Their interaction is described in the following sections. Together, these 2 functions comprise the majority of **UsbCore.C**, and are the key to reuse of this code in other applications.

In all cases, **Ep0RxIsr()** takes care of cleaning up after the previous transfer in the event that a Setup packet appears to arrive out of order from the perspective of the Device; this same function also takes care of cleaning up after any Stalls. As a result, whenever a Setup packet arrives, the software is placed into an IDLE State so that it can process the packet. Note that under some circumstances a handshake was not sent to the Host for the Setup packet; when this happens, the packet is dropped and the Host will retransmit it.

For all Transfer types, the transition from IDLE -> SETUP is performed in **Ep0RxIsr()**; this involves copying the packet into the **usbSetupPkt** structure, removing and releasing the packet, and setting the State.

Whenever **Ep0FSM()** needs to send any packets to the Host, it makes use of a companion function **Ep0SendPkt()**, which is also contained in **UsbCore.C**. This function takes a generic pointer to a data block to send which, although it makes the code execution relatively slow, permits sending a data block located in any memory space of the MCU. The function also takes a **bytesToSend** argument indicating the packet size; zero is a valid size for this

function because zero-length packets are used as handshakes in some Transfer types. The function allocates a packet from the MMU and saves the packet number in **ep0TxPkt**, since this value is sometimes needed in order to clean up after a transmission. The function does not perform any manipulation of the Data Toggle, which is left to the caller, since this varies depending upon the circumstance.

### CTL-READ Transfers

The transition from SETUP -> RD\_DATA is performed in **Ep0FSM()**. This consists of validating the Setup packet contents, and determining whether to Stall the transfer or not. In the sample code, a desire to Stall the transfer is indicated by a **wEp0BytesToSend** size of zero; the Stall is implemented as described in a previous section and the State is set to IDLE.

Validating the Setup packet for Standard Requests is performed in a large **switch** statement that has cases for each of the Standard Transfers; each case in turn validates the remaining bytes in the Setup packet according to the USB Specification. In most cases, application specific information is obtained by calling **ApUsbXXX()** functions defined in **UsbCore.H** and contained in **Chpt6.C**. A discussion of each Transfer type is contained in a following section. For Vendor or Class Requests, the Setup packet is passed to **ApUsbEp0Read()**, which is defined in **UsbCore.H** and contained in **Chpt6.C**, in order to validate the Transfer.

Assuming success, a data pointer and byte count is available for the data to be returned to the Host. The byte count is limited to the maximum size specified in the Setup Packet, and the State is set to RD\_DATA. The **TX\_TOGGLE** is set to DATA1 PID in preparation for the first packet to be returned.

In the RD\_DATA State, packets are sent to the Host using **cEP0\_MAX\_PKT\_SIZE** (defined in **Chpt6.H**) packets until all of the data is sent. If the data is an exact fit in the packets, then a zero-byte packet is sent last, which the Host may or may not read. In order to avoid having to service IRQ's each time a USB packet is sent or a TxFIFO goes empty (the 2 choices for the device), EP0 transmits are cleaned up at the beginning of **Ep0FSM()** each time it is called. Eventually, the Host will send a handshake packet, which will be received in **Ep0RxIsr()**; this function will clean up any left over Tx packet and set IDLE State.

### No-Data Control Transfers

Everything is very similar to the previous section, except that the State transition is to WR\_STATUS if the Transfer is not Stalled, a **bResult** value of FALSE is used to request a Stall instead of **wEp0BytesToSend** being zero, and Vendor or Class requests are passed to **ApUsbEp0NoData()** instead of **ApUsbEp0Read()**. As before, most Transfers involve calls to **ApUsbXXX()** for application dependent data and validation. Any Stall is implemented as described in a previous section. Success is indicated in WR\_STATUS by queuing a zero-byte Tx packet with a DATA1 PID; this transmission is cleaned up at the beginning of **Ep0FSM()** the next time it executes after the Host reads the packet. In the event that the Host handshake to the packet is lost, or if the next Setup packet arrives before the next **Ep0FSM()** execution, then the Tx is cleaned up when the next Setup packet arrives in **Ep0RxIsr()**.

### Control-Write Transfers

Although none are used by the sample code, the framework is in place in case an application should need them for Vendor or Class transfers. In **Ep0FSM()** in the SETUP State, the Setup packet is passed to **ApUsbEp0Write()** in order to determine whether to Stall the Transfer or not. Any Stall is implemented as described previously, the State is set to IDLE, and the Stall is cleared in the **Ep0RxIsr()** when the next Setup packet arrives.

Assuming that the Application chooses to accept the Transfer, the State is advanced to WR\_DATA, and subsequent OUT packets are queued in **GP\_FIFO1** by **Ep0RxIsr()** and passed to **ApUsbEp0WriteNextPkt()** by **Ep0FSM()**. The end of the Transfer is detected in **Ep0FSM()** by receiving a less than **cEP0\_MAX\_PKT\_SIZE** packet, at which time the State is advanced to WR\_STATUS. In this State, a zero-byte handshake packet is queued, and the State is set to IDLE, as was previously described. Note that in a real application, the software should also keep track of the total bytes received from the Host during the transfer, and should Stall the transfer if it ever exceeds **wLength** from the Setup packet.

**Ep0RxIsr()** takes care of Data Toggles by initializing **bEp0RxTog** when the Setup packet arrives, and then rejecting packets and updating the toggle as each subsequent data packet arrives. It also makes the Rx BUSY as each packet arrives, and checks for overflows of **GP\_FIFO1**. In the foreground, **Ep0FSM()** removes the busy whenever **GP\_FIFO1** is emptied of packets in order to permit more packets to arrive.



## EP0 Application Callbacks

As mentioned above, most Transfers involve application specific data, so callback functions are defined in **UsbCore.H** and the code is contained in **Chpt6.C**.

TRANSFER	TYPE	APPLICATION FUNCTION
GET_DESCRIPTOR	CTRL-READ	ApUsbGetDesc()
GET_CONFIGURATION	CTRL-READ	usbCfg (shared var.)
GET_INTERFACE	CTRL-READ	ApUsbGetIF()
GET_STATUS(DEVICE)	CTRL-READ	usbDevStatus (shared var.)
GET_STATUS(INTERFACE)	CTRL-READ	None (reserved transfer)
GET_STATUS(ENDPOINT)	CTRL-READ	ApUsbIsEpValid()
SET_ADDRESS	NO-DATA-CTRL	None (Ap doesn't care)
SET_CONFIGURATION	NO-DATA-CTRL	ApUsbSetCfg()
SET_INTERFACE	NO-DATA-CTRL	ApUsbSetIF()
CLEAR_FEATURE(EP_STALL)	NO-DATA-CTRL	ApUsbIsEpValid/ApUsbCFES()
CLEAR_FEATURE(WAKEUP)	NO-DATA-CTRL	ApUsbRemoteWakeDisable()
SET_FEATURE(EP_STALL)	NO-DATA-CTRL	ApUsbIsEpValid/ApUsbSFES()
SET_FEATURE(WAKEUP)	NO-DATA-CTRL	ApUsbRemoteWakeEnable()

For all Standard Transfers, the Core code handles validating the Setup packet contents, as well as confirming that the device is in a valid State for the Transfer (e.g., many Transfers are only valid when the device is in a CONFIGURED State). Following this, the Application callbacks are used to validate any application dependent values (e.g., whether a specified Alternate Setting is valid for a given application).

For each of the GET\_XXX Transfers, the information is entirely application specific, so the Core software just calls the appropriate Application function, which then either provides the required information, or returns a zero or FALSE to indicate that the Transfer is not supported and should be Stalled.

For any Transfer that involves an EP as a Target, the Core code calls **ApUsbIsEpValid()**, since it is application specific which EP's are valid for the current combination of Alternate Interface Settings.

For SET\_ADDRESS, the Core code handles setting the new USB Address after the Host has read the handshake packet. This can be seen when Tx packets are cleaned up near the entry to **Ep0FSM()**.

For SET\_CONFIGURATION(0), the Core code handles disabling all EP's; the Application code is responsible for setting up any EP's for a non-zero configuration. The Application code is also responsible for configuring EP's for SET\_INTERFACE, since any such settings are application dependent.

For CFES, the Core code handles flushing any TxFIFO's that the application does not flush, and it handles the EP\_CTRL register for all EP's (including TxToggles), but the application is responsible for handling Rx Data Toggle initialization and flushing any Rx Queues, since these are application dependent. For SFES, the Core code handles the EP\_CTRL register. For any Transfer with an EP target, the Core code validates the EP using **ApUsbIsEpValid()**, as was previously described.

For both WAKEUP Requests, all processing is handled by the application, which is also responsible for indicating the Feature state in the **usbDevStatus** shared variable.

As was previously described in the section on Control-Read Transfers, the Core code handles limiting the size (i.e., byte count) of the Data Stage, splitting it into multiple USB packets, handshakes, Stalls, etc.

Because of the amount of support provided by the Core functions, most of the Application Functions do absolutely nothing, with the possible exception of returning a pointer and/or a byte count, and the remainder does only a small amount of processing. It is by customizing these functions that new applications can be easily ported to this architecture.

## APPLICATION POLLING

There are 2 polled callbacks to the application. As the name would imply, **ApCfg1Poll()** is called whenever the device State is CONFIGURED, and the Configuration is 1 (i.e., the Configured State). For this application, a total of 4 Endpoints is used, and the polling handler just calls each individual EP handler in turn. The function implemented by each EP handler is described in a later section.

**ApPoll()** is called any time the device is not SUSPENDED; its primary use is for debugging. The example code calls a slow polling function **ApSlowPoll()** either every USB frame while the device is receiving SOF's, or based on a timeout of 6 mS in the Configured State (CFG1), or 2 mS otherwise. **ApSlowPoll()** in turn displays and resets any ISR0 errors, and checks for any keystrokes from the debugging terminal, to which it responds by sending a DBGPRINT message. This code could be expanded to aid in debugging any other application by returning various data and/or register values, or by initiating programmer defined sequences in response to specific keystrokes. This function is instrumented by pulsing GPIO5 high for the duration of its execution. Since the **GPIOA\_OUT** register is shared with the ISR, IRQ's are disabled during the access to this register.

## DATA TRANSFER

### Data Transfer Overview

The application code treats the low 32 KB ISA RAM as a pair of 16 KB circular buffers. It treats the 4 USB EP's as 2 pairs -- 1 Rx and 1 Tx in each pair. Within each pair, packets arriving on the RxEP are written to the ISA RAM, while packets already in the ISA RAM are queued for Tx back to the Host. In this way, each pair of EP's loops back data, with up to 127 USB packets (1 packet less than 8 KB) in each transfer. EP1 and EP3 are the 2 Rx EP's, and are identical except for using separate variables, while EP2 and EP4 are the Tx partners.

**ep1RxQueue[]** and **ep3RxQueue[]** are the Rx packet queues; packets are pushed on the head of each queue using **ep1RxHead** and **ep3RxHead**, and are popped from the tail of each queue using **ep1RxTail** and **ep3RxTail**, as pointers. Empty queue entries are marked with **INVALID\_PN**. Each of these queues is initialized by **ApSwlnit()**, which is called by the Core code during initialization and is also called by **ApUsbReset()** for a USB Reset, using **Ep1RxQueInnit()** and **Ep3RxQueInnit()**. The corresponding queue is flushed by **ApUsbCFES()** using **Ep1RxQueFlush()** and **Ep3RxQueFlush()**.

Since both RxEP's are for BULK packets, data toggles are maintained in **bEp1RxToggle** and **bEp3RxToggle**, as was described in Chapter 4. The data toggles are reset in **ApUsbSetCfg()** and **ApUsbCFES()**.

The memory management is based on what was described in Chapter 4. Each RxEP is allowed to use a peak of 9 MMU pages, and the Threshold is set for 6 (both numerical values are in #define's), which is enforced by making the corresponding RxEP busy in the RxISR, and unbusy in the foreground. Each TxEP is allowed to use a peak of 5 MMU pages, and this is enforced by making sure that the corresponding TxFIFO is not full before attempting to allocate a packet memory. This policy leaves a minimum of 4 free MMU pages for EP0 under peak conditions.

Each pair of Rx and TxEP's share an ISA circular buffer; packets are pushed on the head of each queue using **ep1IsaHead** and **ep3IsaHead**, and packets are popped from the tail of each queue using **ep1IsaTail** and **ep3IsaTail** as pointers. When each queue is empty, its head and tail pointers are equal; a queue is considered full when its head pointer is 1 behind its tail pointer. Each 16 KB buffer is treated as having 128 pages of 128 bytes each, so it can hold up to 127 USB packets. Memory-To-Memory DMA is used to perform the actual transfers to/from each ISA page, using **DmaEntirePktToIsa()** and **DmalsaToEntirePkt()** respectively. The session size is always set to a full 72 bytes, which will handle a maximum size BULK packet, including its 8-byte header. This simplifies the task of handling variable size packets and, recognizing that most packets are maximum size anyway, results in good throughput as well. The details of the DMA functions are as was described in Chapter 5. The ISA queues are marked empty by **ApSwlnit()**, which is called by the Core code during initialization and is also called by **ApUsbReset()** for a USB Reset.

Since all 4 EP's perform Memory-To-Memory DMA, they must all share DMA channels 0 and 1. In order to support this sharing, the **dmaOwner** variable is used; it is set to the corresponding EP [1,4] that currently owns the DMAC, or else it is cleared to zero if no EP is currently using the DMAC, signifying that the DMAC is available for use.

### Data Transfer Details

When packets arrive from the USB, the Core code initially receives the IRQ in its **Isr0()** function, and passes execution to **Aplsr0()**. Note that Register Bank switching is used for speed, as was described in a previous chapter. **Aplsr0()** saves and restores the MMU state on entry and exit of the Rx FIFO loop, since it needs to use

some of the MMU registers in order to handle the packets. If a packet is addressed to EP0, **ApIsr0()** calls back to **Ep0RxIsr()** in the Core code in order to handle the packet. For application packets, **ApIsr0()** validates the packet and discards it for Bad CRC, Stall, or Data Toggle (since these are Bulk packets). Assuming the packet is accepted, it is pushed on the corresponding software queue; if a queue overflows, which should never happen because of the MMP code, it is considered a Fatal Error, and the corresponding EP is Stalled, and the error code is logged in **apIsr0Error**, which will be displayed by the foreground the next time **ApSlowPoll()** is called. As each packet is queued, the packet count for the EP is incremented, and the EP is made BUSY if its count exceeds the defined threshold. For instrumentation, **Isr0()** pulses GPIO2 (see the comment in the GPIO section later in this chapter) high for the duration of its execution, **ApIsr0()** pulses GPIO1 high each time around the Rx FIFO loop, and it also pulses GPIO0 high whenever it discards an EP1 or EP3 packet for bad Data Toggle.

Eventually, the corresponding RxEP handler (1 or 3) will execute in the foreground and will find that its Rx queue is not empty, its ISA queue is not full, and that the DMAC is available. At this point, it will claim ownership of the DMAC by setting the **dmaOwner** to its EP number, and will start a DMA session using **DmaEntirePktTolsa()**. Eventually, it will see that it is the **dmaOwner** and that the session has completed, at which point it will update its queue pointers and mark the **dmaOwner=0** to signify that another EP may use it. Once the DMA is complete, the packet is freed in the MMU, the packet count is decremented, and the EP is made unbusy if the total number of packets has dropped below threshold. Note that the packet count and EPCTRL register are shared with the ISR, so IRQ's are disabled/enabled around these accesses. Also, EP1 raises GPIO3 high when starting a DMA session, and brings it back low when the DMA completes. Similarly, since **GPIOA\_OUT** is shared with the ISR, IRQ's are disabled around accesses to this register.

Eventually, the corresponding TxEP handler (2 or 4) will execute in the foreground and will find that its Tx queue is not full, its ISA queue is not empty, and that the DMAC is available. At this point, it will claim ownership of the DMAC by setting the **dmaOwner** to its EP number, and will start a DMA session using **DmalsaToEntirePkt()**. Eventually, it will see that it is the **dmaOwner** and that the session has completed, at which point it will update its queue pointers and mark the **dmaOwner=0** to signify that another EP may use it. Before starting the DMA session, it will allocate a packet from the MMU and save the packet number in a local static variable, which it will queue for Tx when the DMA session is complete. For the EP2 handler, GPIO4 is raised high when the DMA session is started, and brought back low when it completes; for EP4, no GPIO's are pulsed. Since **GPIOA\_OUT** is shared with the ISR, IRQ's are disabled around accesses to this register.

### GPIO Summary

The sample code is instrumented as follows:

GPIO7: set high in Ap to connect to USB at full speed (12 Mbps)

GPIO5: pulses high for the duration of ApSlowPoll().

GPIO4: pulses high for the duration of EP2 DMA from ISA to MMU

GPIO3: pulses high for the duration of EP1 DMA from MMU to ISA

GPIO2: pulses high for the duration of Isr0() [see comment below about GPIO2]

GPIO1: pulses high each time around the USB Rx FIFO loop

GPIO0: pulses high for a USB Rx pkt on EP1 or 3 bad toggles

All GPIO's are manipulated in the Application code, with the exception of GPIO2. To be faithful to the architecture description in which all GPIO's are owned by the application, this should be removed from the Core code and either placed in the Application or discarded.

## USB HUB BLOCK

The USB Hub Block consists of a Hub Serial Interface Engine, Hub Repeater, Hub Command Sequencer, and a Hub Control. Hub configuration, Status, Port power control, selective reset on a port-by-port basis, fault recovery frame time logic, are just some of the functions performed by the Hub Block.

The Hub Block consists of 8 registers that are memory mapped into the 8051 MCU memory space. The IdVendor, IdProduct, BcdDevice Low and High Byte registers, HubControl1 register, and reserved register.

As the register name calls it the registers define the particular Vendor ID, Product ID, Device release number Id, except for HubControl1 register which is used to set the Hub in five different modes. One important mode is one, which allows the USB97C102 to be pin and function compatible to the USB97C100, making porting of code very simple. Initialization of the internal USB Hub is a requirement in all modes and conditions relative to a hardware and/or software reset. The requirements for USB Hub initialization are discussed below.

## HUB Initialization

HUB Initialization takes place when the MCU resets the Hub by asserting NhubReset bit in the Hubcontrol1 register. The hub will not respond to any enumeration or device request. This will allow the Hub registers to be initialized prior to enumeration. Initialization of Hub control registers must be accomplished within two (2) ms after reset is de-asserted. After de-assertion of NhubReset bit Hub controller is ready to receive packets from the USB Root Host Controller. Each port will then be enabled and initialized via a control packet from the host.

## HUB Control Register1

HUB Control Register1 allows designer to disconnect / connect ports from Hub by setting one of the 4 HubBypass bits in Hub Control Register1. A total of 5 Hub Bypass modes can be selected. Just remember that when setting the USB97C102 in any HubBypass mode, it is recommended that only one HubBypass bit be set.

Mode 1: "Native Mode / Normal Mode" when HubBypass (2-5) bits in Hub Control Register1 are cleared (0). In this mode no hub bypass is done. Utilizing all up and down stream ports

Mode 2: "USB97C100 Compatibility Mode" when HubBypass2 bit in Hub Control Register1 is set to one, Port 1 and 2 are no longer connected to the Hub. Port 1, which is connected to the rest of the USB97C102, is connected to Port 2. Port 2 becomes the upstream of Port 1.

Mode 3: When HubBypass3 bit in Hub Control Register1 is set to one, Port 1 and 3 are no longer connected to the Hub. Port 1, which is connected to the rest of the USB97C102, is connected to Port 3. Port 3 becomes the upstream of Port 1.

Mode 4: When HubBypass4 bit in Hub Control Register1 is set to one, Port 1 and 4 are no longer connected to the Hub. Port 1, which is connected to the rest of the USB97C102, is connected to Port 4. Port 4 becomes the upstream of Port 1.

Mode 5: When HubBypass5 bit in Hub Control Register1 is set to one, Port 1 and 5 are no longer connected to the Hub. Port 1, which is connected to the rest of the USB97C102, is connected to Port 5. Port 5 becomes the upstream of Port 1.

## Example on how to initialize the HUB:

```
{      :
```

```
/* Initializing Vendor ID */
```

```
    ID_VENDOR_L = 0x24;  
    ID_VENDOR_H = 0x04;
```

```
/* Initializing Product ID; must always be initialized before power up and prior to enumeration */
```

```
    ID_PRODUCT_L = 0x01;  
    ID_PRODUCT_H = 0x00;
```

```
/* Initializing USB device release number */
```

```
    BCD_DEVICE_L = 0x00;  
    BCD_DEVICE_H = 0x02;
```

```
    #if MKF_SELF_POWERED  
        SIE_CTRL2 = HUB_SELF_POWERED_;  
    #endif  
    #if MKF_HUB_BYPASS
```

```
/*setting device to HubBypass2 mode in order to use device as a USB97C100 pin compatible device */
```

```
    HUB_CTRL_1 = HUBCTL1_NhubReset_ | HUBCTL1_HubBypass2_;
```

```

        #else

/* setting hub controller ready to receive packets from root host controller */

        HUB_CTRL_1 = HUBCTL1_NhubReset_;
        #endif
        :
    #endif
        :
        :
}

```

#### **USB97C100 Compatibility (HUB BYPASS 2) Mode**

As mentioned above the USB97C102 can be placed in a mode to emulate the USB97C100 in terms of pin and function compatibility. “HubBypass 2” only needs to be implemented if and only if the designer is using the USB97C102 as a USB97C100 pin compatible device. In other words if USB97C102 is to be used in a design currently utilizing the USB97C100 device, since this change will allow the USB97C102 to be pin replaceable in an existing USB97C100 design.

Porting of firmware from the USB97C100 to the USB97C102 is fairly easy to port if the following two steps are taken into account. First: Designer must initialize Hub (refer to HUB Initialization) before any mode or condition take place. Second: Designer must clear the associated bit(s) of any of the following registers if used: ISR\_0, ISR\_1, WU\_SRC\_1, WU\_SRC\_2. These registers are cleared by writing a “1” to the associated bit(s) in the register. This is essential since bit(s) are automatically cleared each time these registers are read when using USB97C100 device. If writing new firmware for the USB97C102, then one should individually clear the associated bit(s) in the associated “Interrupt Source” registers as the corresponding interrupts are handled. Refer to Application note 8.12 for more information on how to port code.