

AMXTM Timing Guide and Data

for

AMX for ARM Multitasking Executive

First Printing: August 1, 1997

Last Printing: April 1, 2003

Copyright © 1997 - 2003

KADAK Products Ltd.

206 - 1847 West Broadway Avenue

Vancouver, BC, Canada, V6J 1Y5

Phone: (604) 734-2796

Fax: (604) 734-8114

**Copyright © 1997-2003 by KADAK Products Ltd.
All rights reserved.**

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without the prior written permission of KADAK Products Ltd., Vancouver, B.C., CANADA.

DISCLAIMER

KADAK Products Ltd. makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability and fitness for any particular purpose. Further, KADAK Products Ltd. reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of KADAK Products Ltd. to notify any person of such revision or changes.

TRADEMARKS

AMX in the stylized form and KwikNet are registered trademarks of KADAK Products Ltd. AMX, AMX/FS, InSight, *KwikLook* and *KwikPeg* are trademarks of KADAK Products Ltd. Microsoft, MS-DOS and Windows are registered trademarks of Microsoft Corporation. ARM is a trademark of Advanced RISC Machines Ltd. All other trademarked names are the property of their respective owners.

AMX TIMING GUIDE and DATA

Table of Contents

	Page
1. AMX Timing Issues	1
2. Task Switching Measurements	5
2.1 Task Wait/Wake Synchronization	5
2.2 Task Mailbox Synchronization	6
3. Clock Service Measurements	7
3.1 Clock Interrupt and Timer Service	7
3.2 Task Wait (Delay) for Timed Interval	8
4. Interrupt Response Measurements	9
4.1 Interrupt Latency	9
4.2 Handler Latency	10
5. Timing Data Sheets	11
ARM7TDMI and ARM Ltd SD Toolkit C (ARM state)	11
ARM7TDMI and ARM Ltd SD Toolkit C (Thumb state)	13
Atmel AT91 and ARM Ltd ADS C (ARM state)	15
Atmel AT91 and ARM Ltd ADS C (Thumb state)	17
ARM966E and ARM Ltd ADS C (ARM state)	19
ARM966E and ARM Ltd ADS C (Thumb state)	21
XScale 80321 and ARM Ltd ADS C (ARM state)	23
XScale 80321 and ARM Ltd ADS C (Thumb state)	25

This page left blank intentionally.

1. AMX Timing Issues

The current release of KADAK's AMX™ Multitasking Executive is the fifth in a long series spanning more than 20 years of development. With AMX coded in C and with the increasing use of RISC processors, kernel timing metrics have become increasingly difficult to define and measure.

It is the purpose of this AMX Timing Guide to provide actual AMX timing information and to introduce the factors which affect these performance figures. Since instruction counts and cycle counts can no longer be used to measure execution times, it is necessary to provide specific measurements and to define, in detail, the conditions under which they were made. The remainder of this chapter will introduce the major factors which make timing measurement so difficult.

Chapter 2 introduces the timing tests used to measure typical task switching times in commonly occurring task to task synchronization scenarios. Chapter 3 defines the processing overhead necessary to support AMX timing services. Chapter 4 introduces the topic of interrupt latency and illustrates typical interrupt service timing.

Timing data sheets are provided in Chapter 5. The data sheets provide specific timing figures for all of the measurements described in Chapters 2 to 4. The data sheets identify the exact characteristics of the hardware on which the measurements were made and describe the C compiler options used to compile AMX.

Processor Effects

Today's RISC and superscalar microprocessors use many techniques to boost instruction execution speeds. Internal instruction queues are used to allow instruction fetch and decode operations to proceed while previously decoded instructions are executing. Multiple execution units are often employed to permit simultaneous, parallel execution of instructions. Register scoreboarding may occur so that execution only stalls when a particular instruction needs a result from another incomplete instruction.

Branch prediction may be used to force the processor to assume a particular execution path which is the expected norm. A time penalty is then incurred only in the rare cases when the prediction is in error.

Processors may use multiple register sets or register windowing schemes to improve procedure parameter passing and to minimize the register preservation requirements of procedures.

Memory Effects

The processor's interface to memory is the single greatest impediment to full speed instruction execution. Unless instructions can be fetched from memory with zero delay, instruction execution will stall.

Various techniques are used to improve memory access. The most obvious is the use of high speed memory devices which are matched to the processor clock speed. Unfortunately, cost often precludes this preferred solution.

Most high performance processors provide an autonomous memory control unit capable of independent execution. High speed, multi-ported memory buffers may be used to allow burst transfers of small blocks of data to or from memory in anticipation of the processor's needs. In some cases, the block transfer size is configurable. The manner in which instructions straddle these block boundaries can directly affect performance. Minor shifts of code up or down in memory may noticeably affect execution time.

The processor can minimize the effects of slow data memory if instructions with no external memory references can be executed while previous memory accesses complete. However, this technique depends on judicious ordering of instructions to meet the memory access constraints, a requirement which cannot always be met. The instruction ordering burden falls on the assembly language programmer or on the high level language code generator.

Caches

Another common technique used to improve execution speed is instruction and/or data caching. The instruction cache is high speed memory internal or external to the processor which is used to hold copies of instructions as they are fetched from memory. The data cache is high speed memory internal or external to the processor which is used to buffer data on its way to or from external memory.

The entire address space maps to the cache in some processor dependent fashion. The cache includes control information which identifies whether each instruction or data element in the cache is valid. When the processor fetches an instruction or data element that is already marked as valid in the cache, the cache copy is used with no time penalty and no external memory interference.

The instruction cache is usually internal to the processor and hence fixed in size. Some processors also include an internal data cache. When an external cache is used, the size of the cache is usually determined by the manner in which it is interfaced to the processor.

Some processors allow the sizes of the internal instruction and data caches to be adjusted to meet the needs of specific applications. Some processors also allow specific blocks of instructions to be copied to the cache and locked in place for guaranteed fast access to particular pieces of code.


Finally, most processors allow the application to dynamically enable or disable the instruction and/or data caches. It is also usually possible to invalidate all or portions of the instruction and/or data caches.

Code and Data Location

Instruction and data caching can show dramatic improvement in the execution time of many applications in which the path of execution is single threaded. Much research has been done to show that for many applications, both instruction and data accesses are frequently very localized and hence benefit immensely from caching.

However, in a multitasking environment, the thread of execution is not so linear. Task switches require stack switching and context saving and restoral. These switches, often in response to interrupt activity, interfere with the sequential instruction execution process thereby negating some of the benefits of instruction caching.

The position of code and data in memory can also affect the caching performance as the following example illustrates. Assume that two code modules, A and B, are located in memory at addresses `0x1000` and `0x2000` respectively as shown in the diagram below. Assume that each module occupies less than 1024 bytes, the size of the instruction cache. Also assume that addresses map directly to the cache using the least significant 10 bits of the address.



If procedure `Bhigh` in module B repeatedly calls procedure `Alow` in module A, both procedures will end up in cache and benefit accordingly.

However, if procedure `Bhigh` in module B repeatedly calls procedure `Ahigh` in module A, the procedures will compete with each other for residence in the cache and performance will suffer.

Few applications are small enough to permit the fine tuning necessary to achieve optimal code location for maximum cache benefit. It is more likely that your code, and the procedures which it calls, will simply end up in memory at locations which vary with every minor code change. Move any block of code up or down in memory with respect to the other blocks of code with which it interacts and you can expect execution timing differences with caching enabled.

Compiler Effects

Earlier releases of AMX for many CISC processors were coded in assembly language. The reasons were historical. The most significant benefits were execution speed and code invariability. For these older products, the AMX code executed by the processor today is exactly the same code tested by KADAK and in use by developers world wide since the particular AMX product was first released. Only documented code changes to eliminate faults or add enhancements have been made.

Today, AMX is coded in C to satisfy the demands of developers not familiar with assembly language code and to ease the porting of AMX by KADAK to new and diverse processor architectures. However, the benefits of C have not been without a cost.

The greatest cost has been the sacrificing of AMX code invariability. The quality and efficiency of the AMX kernel is now dependent upon the code generation capabilities of the C compiler with which AMX is compiled. It is unlikely that any two C compilers will produce identical AMX code sequences. In fact, different releases of the same C compiler may very well produce different copies of AMX code. Fortunately, the testing methods developed by KADAK over many years now lead to almost fault free product releases.

There are other factors introduced by the C compiler which also may affect execution time. The parameter setup and procedure entry and exit code sequences can vary greatly among C compilers. The register preservation requirements of a particular C compiler, over and above those dictated by the processor, may affect AMX overhead.

Finally, beware of code optimizations touted by the compiler vendors. AMX source code does not include the kind of C code sequences which often confuse optimizing C compilers. Despite this fact, several C compilers have, with the simplest level of optimization enabled, generated invalid instruction sequences for perfectly valid C statements.

2. Task Switching Measurements

2.1 Task Wait/Wake Synchronization

AMX allows tasks to synchronize using a very simple wait/wake mechanism. The following example uses this feature to measure the time taken to start and end a task and to switch between tasks.


Assume that low priority task B is running and triggers higher priority task A. Since task B is preempted by the trigger, AMX suspends task B and starts task A. Task A then unconditionally waits. AMX suspends task A and resumes task B which then wakes task A. AMX suspends task B and resumes task A which then ends. AMX sets task A idle and resumes task B which also ends.

The code for these two tasks is as follows.

```
void CJ_CCPP taskA(void)
{
    cjtkwait();
}

void CJ_CCPP taskB(void)
{
    cjtktrigger(taskAid);
    cjtkwake(taskAid);
}
```

On a time line, the thread of execution appears as illustrated in the following diagram (not to scale). Measurements for times T_i are provided on separate AMX timing data sheets.



2.2 Task Mailbox Synchronization

AMX allows tasks to synchronize by sending messages to a mailbox. The following example uses this feature to measure the time taken to send and receive AMX messages and to switch between tasks.

Assume that low priority task B is running and triggers higher priority task A. Since task B is preempted by the trigger, AMX suspends task B and starts task A. Task A then unconditionally waits at a mailbox for a message to arrive. AMX suspends task A and resumes task B which then sends a message to the mailbox. Since task B is preempted by the message transmission, AMX suspends task B, delivers the message from the mailbox to task A and resumes task A which then ends. AMX sets task A idle and resumes task B which also ends.

The code for these two tasks, ignoring message construction and interpretation, is as follows.

```
void CJ_CCPP taskA(void)
{
    struct cjxmsg recvmsg;

    cjmbwait(mailboxid,
             &recvmsg, 0, 0);
}

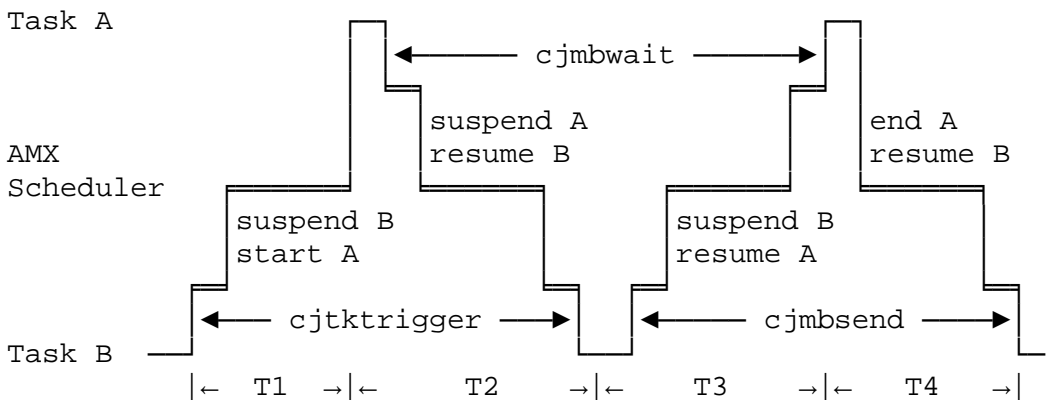
void CJ_CCPP taskB(void)
{
    struct cjxmsg sendmsg;

    cjtktrigger(taskAid);
    cjmbsend(mailboxid,
            &sendmsg, CJ_NO);
}
```

Note

The default AMX message size is 12 bytes. Task A receives a local copy of the 12 byte message ready for processing.

On a time line, the thread of execution appears as illustrated in the following diagram (not to scale). Measurements for times T_i are provided on separate AMX timing data sheets.



3. Clock Service Measurements

3.1 Clock Interrupt and Timer Service

AMX is delivered with a clock driver for one or more counter/timer devices. The clock driver responds to clock interrupts, dismisses the interrupt request and calls the AMX Clock Handler. If an active AMX timer expires, the AMX Clock Handler preempts the interrupted task and forces the AMX Kernel Task to service the application timer.

The following example measures the clock overhead for each AMX tick and for servicing an expired periodic application timer.

Assume that the AMX clock tick is set to exactly match the hardware clock interrupt rate. Assume that task A creates and starts a periodic AMX timer with a period of two AMX ticks. The Timer Procedure associated with the timer does absolutely nothing. Task A goes compute bound forever in order to illustrate task preemption for clock service.


The code for the Timer Procedure and task A is as follows.

```
void CJ_CCPP timerproc(CJ_ID timerid, void *unused)
{
}

void CJ_CCPP taskA(void)
{
    CJ_ID    timerid;

    cjtmcreate(&timerid, "TIMR", (CJ_TMRPROC)timerproc, 2, CJ_NULL);
    cjtmwrite(timerid, 2);
    for (;;) ;
}
```

On a time line, the thread of execution appears as illustrated in the following diagram (not to scale). Measurements for times T_i are provided on separate AMX timing data sheets.



3.2 Task Wait (Delay) for Timed Interval


AMX allows tasks to delay for one or more AMX clock ticks. The following example uses this feature to measure the time taken for a task to resume execution after a timed delay.

Assume that the AMX clock tick is set to exactly match the hardware clock interrupt rate. Task A calls AMX to wait for one AMX tick. AMX suspends task A and resumes some other lower priority task X. When the next clock tick occurs, the clock driver dismisses the interrupt request and calls the AMX Clock Handler. The AMX Clock Handler forces AMX to suspend task X and start the AMX Kernel Task to service the expired task delay. AMX then resumes task A.

The code for task A is as follows.

```
void CJ_CCPP taskA(void)
{
    cjtkwaitm(1);
}
```

On a time line, the thread of execution appears as illustrated in the following diagram (not to scale). Measurements for times T_i are provided on separate AMX timing data sheets.



4. Interrupt Response Measurements

4.1 Interrupt Latency

Interrupt latency, sometimes called interrupt response time, means different things to different people. Within the AMX environment, KADAK uses the term interrupt latency to refer to the time taken from the leading edge of an external interrupt request signal to the processor to the fetch of the first interrupt service instruction by the processor.

Interrupt latency consists of the sum of two components: processor interrupt latency and software interrupt latency.

Processor Latency

Many factors affect interrupt latency. The first is the delay generated by the processor itself in its recognition of the external interrupt request. This delay may be fixed for some processors and variable for others. For RISC and superscalar processors, the delay may range from 3 or 4 processor clock cycles to tens or even hundreds of clock cycles.

The longest delays are usually associated with interrupt requests which occur while the processor is executing complex instructions which, for processor reasons, must be allowed to complete (or at least reach some internal state at which the instruction can be resumed) before the interrupt request can be acknowledged. Multiply and divide are instructions of this type, as are looping instructions such as block moves and string compares.

Additional delay may stem from instruction and/or data caching operations. Response will be faster if the interrupt service instructions happen to already be in the cache as a consequence of a previous interrupt.

Since the processor's interrupt response always requires the saving of some processor state information, the processor architecture can dramatically affect latency. Expect longer delays if the processor automatically saves some or all registers or switches to alternate register sets. The delays will be even longer if the saving is done to memory.

Software Latency

Most processors permit some or all internal and/or external device interrupts to be inhibited under software control. The longest interval for which any software module disables interrupts is called software latency.

The AMX 4-ARM and AMX 4-Thumb kernels use this feature from time to time to briefly inhibit all interrupt sources while they execute some critical section of code. The longest interval for which AMX disables interrupts is called the **AMX interrupt latency**.

KADAK has identified and measured the worst case (longest) intervals during which AMX inhibits interrupts. AMX interrupt latency measurements are provided on separate AMX timing data sheets.

4.2 Handler Latency

Of equal importance to interrupt latency is what KADAK calls handler latency, the time from the processor's first response to an interrupt request signal through to the first useful instruction in the interrupt service procedure.

The term useful is nebulous but can be characterized as follows. Useful instructions are those which deal with the interrupting device and its related application. Useful instructions can not usually be executed until a few processor registers have been saved freeing them for use. In some cases, a switch to an alternate stack may be necessary. Once these housekeeping operations have been completed, useful work can be accomplished.

Handler latency is therefore the sum of the processor latency and the housekeeping execution time.


Interrupt Service

In the AMX 4-ARM and AMX 4-Thumb environments, handler latency is easily measured since the AMX Interrupt Supervisor performs all housekeeping actions.

The following example measures the interrupt overhead for two AMX application Interrupt Service Procedures (ISPs): an ISP coded in assembly language and an ISP coded in C. In each case, an external device interrupt is used to produce an IRQ interrupt. Only this one device is connected to the IRQ pin. Each ISP simply clears the device interrupt request and dismisses the interrupt.

On entry to an AMX ISP coded in assembly language or C, only the few registers which C considers alterable are free for use. All registers, if any, required by C are setup and ready for use.

On a time line, the thread of execution appears as illustrated in the following diagram (not to scale). Measurements for times T_i are provided on separate AMX timing data sheets.



AMX™ 4-ARM Multitasking Executive Timing Data

for

20 MHz ARM7TDMI (ARM state) and ARM Ltd C

November 19, 1997

Processor: ARM7TDMI test processor on the ARM Development Board (ARM7TDMI Version) operating at 20 MHz.

Memory: 32K of "on-chip" 32-bit synchronous SRAM with no wait states. This on board memory simulates on-chip SSRAM in an ARM7TDMI core. An additional 128K of 32-bit static RAM with two (2) wait states is also used.

Cache: The ARM7TDMI test chip has no internal instruction or data caches. External (secondary) instruction and/or data caches are not present.

Code/Data Location: The AMX kernel code and data and the application code and data have been located in SSRAM and SRAM without any attempt to locate either for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Ltd Software Development Toolkit C/C++ v2.11 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-otime* switch optimizes code for speed
- no function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4 code only
- code executes in **ARM state only** (no Thumb code execution)
- AMX 4-ARM procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-ARM task switch, excluding the events leading to the switch, takes 10.6 μ s. This value was derived by averaging 100 task switch measurements. The task switch time is completely determined by the access constraints of the memory system.

**AMX™ 4-ARM Multitasking Executive
Timing Data**

for

20 MHz ARM7TDMI (ARM state) and ARM Ltd C

November 19, 1997

Measurement (note 1)	Timing Diagram	T1 µs	T2 µs	T3 µs	T4 µs
Synchronization (note 2)					
Wait/Wake	(page 5)	24.2	18.3	20.8	16.9
Mailbox	(page 6)	24.2	36.9	50.0	19.0
Clock Service (note 2)					
Timers	(page 7)	11.3	31.5	17.4	
Task Delay	(page 8)	26.4	54.9		
Interrupt Service (note 2)					
Handler Latency	(page 10)	5.4	2.5	5.4	2.5

Interrupt Latency: The worst case AMX 4-ARM interrupt latency is 6.5 µs.

The worst case handler latency for an AMX ISP coded in assembly language is 5.4 µs.

The worst case handler latency for an AMX ISP coded in C is 5.4 µs.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Each test is performed 100 times and the average of all measured values is reported.

AMX™ 4-Thumb Multitasking Executive Timing Data

for

20 MHz ARM7TDMI (Thumb state) and ARM Ltd C

November 19, 1997

Processor: ARM7TDMI test processor on the ARM Development Board (ARM7TDMI Version) operating at 20 MHz.

Memory: 32K of "on-chip" 32-bit synchronous SRAM with no wait states. This on board memory simulates on-chip SSRAM in an ARM7TDMI core. An additional 128K of 32-bit static RAM with two (2) wait states is also used.

Cache: The ARM7TDMI test chip has no internal instruction or data caches. External (secondary) instruction and/or data caches are not present.

Code/Data Location: The AMX kernel code and data and the application code and data have been located in SSRAM and SRAM without any attempt to locate either for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Ltd Software Development Toolkit C/C++ v2.11 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-Otime* switch optimizes code for speed
- no function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4T code only
- code executes in **ARM and Thumb states**
- AMX 4-Thumb and application code compiled with interworking
- AMX 4-Thumb procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-Thumb task switch, excluding the events leading to the switch, takes 11.0 μ s. This value was derived by averaging 100 task switch measurements. The task switch time is completely determined by the access constraints of the memory system.

**AMX™ 4-Thumb Multitasking Executive
Timing Data**

for

20 MHz ARM7TDMI (Thumb state) and ARM Ltd C

November 19, 1997

Measurement (note 1)	Timing Diagram	T1 µs	T2 µs	T3 µs	T4 µs
Synchronization (note 2)					
Wait/Wake	(page 5)	25.1	19.6	22.9	18.2
Mailbox	(page 6)	25.1	41.8	56.1	19.8
Clock Service (note 2)					
Timers	(page 7)	11.4	36.7	21.0	
Task Delay	(page 8)	31.1	65.1		
Interrupt Service (note 2)					
Handler Latency	(page 10)	5.4	2.5	5.4	2.5

Interrupt Latency: The worst case AMX 4-Thumb interrupt latency is 9.2 µs.

The worst case handler latency for an AMX ISP coded in assembly language is 5.4 µs.

The worst case handler latency for an AMX ISP coded in C is 5.4 µs.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Each test is performed 100 times and the average of all measured values is reported.

AMX™ 4-ARM Multitasking Executive Timing Data

for

32 MHz Atmel AT91M42800A (ARM state) and ARM Ltd ADS C

November 1, 2002

Processor: Atmel AT91M42800A processor on the Atmel AT91EB42 Evaluation Board operating at 32 MHz.

Memory: 256 Kbyte of 16-bit static RAM with one (1) wait state.

Cache: The AT91M42800A processor has no internal instruction or data caches. External (secondary) instruction and/or data caches are not present.

Code/Data Location: The AMX kernel code and data and the application code and data have been located in SRAM without any attempt to locate either for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Developer Suite C/C++ v1.2 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-otime* switch optimizes code for speed
- default *-o2* function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4 code only
- code executes in **ARM state only** (no Thumb code execution)
- AMX 4-ARM procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-ARM task switch, excluding the events leading to the switch, takes 16.8 μ s. This value was derived by averaging 100 task switch measurements. The task switch time is completely determined by the access constraints of the memory system.

AMX™ 4-ARM Multitasking Executive Timing Data

for

32 MHz Atmel AT91M42800A (ARM state) and ARM Ltd ADS C

November 1, 2002

Measurement (note 1)	Timing Diagram	T1 µs	T2 µs	T3 µs	T4 µs
Synchronization (note 2)					
Wait/Wake	(page 5)	37.1	28.8	32.5	26.0
Mailbox	(page 6)	37.1	58.1	78.6	29.1
Clock Service (note 2)					
Timers	(page 7)	17.8	49.1	28.4	
Task Delay	(page 8)	41.6	86.0		
Interrupt Service (note 2)					
Handler Latency	(page 10)	8.4	4.0	8.4	4.0

Interrupt Latency: The worst case AMX 4-ARM interrupt latency is 8.1 µs.

The worst case handler latency for an AMX ISP coded in assembly language is 8.4 µs.

The worst case handler latency for an AMX ISP coded in C is 8.4 µs.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Each test is performed 100 times and the average of all measured values is reported.

AMX™ 4-Thumb Multitasking Executive Timing Data

for

32 MHz Atmel AT91M42800A (Thumb state) and ARM Ltd ADS C

November 1, 2002

Processor: Atmel AT91M42800A processor on the Atmel AT91EB42 Evaluation Board operating at 32 MHz.

Memory: 256 Kbyte of 16-bit static RAM with one (1) wait state.

Cache: The AT91M42800A processor has no internal instruction or data caches. External (secondary) instruction and/or data caches are not present.

Code/Data Location: The AMX kernel code and data and the application code and data have been located in SRAM without any attempt to locate either for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Developer Suite C/C++ v1.2 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-Otime* switch optimizes code for speed
- default *-O2* function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4T code only
- code executes in **ARM and Thumb states**
- AMX 4-Thumb and application code compiled with interworking
- AMX 4-Thumb procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-Thumb task switch, excluding the events leading to the switch, takes 17.1 μ s. This value was derived by averaging 100 task switch measurements. The task switch time is completely determined by the access constraints of the memory system.

**AMX™ 4-Thumb Multitasking Executive
Timing Data**

for

32 MHz Atmel AT91M42800A (Thumb state) and ARM Ltd ADS C

November 1, 2002

Measurement (note 1)	Timing Diagram	T1 µs	T2 µs	T3 µs	T4 µs
Synchronization (note 2)					
Wait/Wake	(page 5)	36.6	28.6	32.8	27.1
Mailbox	(page 6)	36.6	55.6	76.3	29.4
Clock Service (note 2)					
Timers	(page 7)	17.9	51.3	29.8	
Task Delay	(page 8)	40.8	89.6		
Interrupt Service (note 2)					
Handler Latency	(page 10)	8.6	4.0	8.6	4.0

Interrupt Latency: The worst case AMX 4-Thumb interrupt latency is 8.1 µs.

The worst case handler latency for an AMX ISP coded in assembly language is 8.6 µs.

The worst case handler latency for an AMX ISP coded in C is 8.6 µs.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Each test is performed 100 times and the average of all measured values is reported.

AMX™ 4-ARM Multitasking Executive Timing Data

for

80 MHz ARM966E (ARM state) and ARM Ltd ADS C

November 1, 2002

Processor: ARM966E processor operating clock doubled at 80 MHz on the 40 MHz ARM Integrator/AP Board.

Memory: 1 Mbyte of 32-bit synchronous SRAM with zero wait states. The internal 32 Kbyte tightly coupled SRAM is not used.

Cache: The ARM966E processor has no internal instruction or data caches. External (secondary) instruction and/or data caches are not present.

Code/Data Location: The AMX kernel code and data and the application code and data have been located in SSRAM without any attempt to locate either for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Developer Suite C/C++ v1.2 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-otime* switch optimizes code for speed
- default *-o2* function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4 code only
- code executes in **ARM state only** (no Thumb code execution)
- AMX 4-ARM procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-ARM task switch, excluding the events leading to the switch, takes 6.0 μ s. This value was derived by averaging 100 task switch measurements. The task switch time is completely determined by the access constraints of the memory system.

**AMX™ 4-ARM Multitasking Executive
Timing Data**

for

80 MHz ARM966E (ARM state) and ARM Ltd ADS C

November 1, 2002

Measurement (note 1)	Timing Diagram	T1 µs	T2 µs	T3 µs	T4 µs
Synchronization (note 2)					
Wait/Wake	(page 5)	14.6	10.6	12.1	9.9
Mailbox	(page 6)	14.6	22.0	29.5	11.2
Clock Service (note 2)					
Timers	(page 7)	6.6	19.2	11.1	
Task Delay	(page 8)	15.2	33.0		
Interrupt Service (note 2)					
Handler Latency	(page 10)	3.0	1.6	3.0	1.6

Interrupt Latency: The worst case AMX 4-ARM interrupt latency is 3.0 µs.

The worst case handler latency for an AMX ISP coded in assembly language is 3.0 µs.

The worst case handler latency for an AMX ISP coded in C is 3.0 µs.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Each test is performed 100 times and the average of all measured values is reported.

**AMX™ 4-Thumb Multitasking Executive
Timing Data**

for

80 MHz ARM966E (Thumb state) and ARM Ltd ADS C

November 1, 2002

Processor: ARM966E processor operating clock doubled at 80 MHz on the 40 MHz ARM Integrator/AP Board.

Memory: 1 Mbyte of 32-bit synchronous SRAM with zero wait states. The internal 32 Kbyte tightly coupled SRAM is not used.

Cache: The ARM966E processor has no internal instruction or data caches. External (secondary) instruction and/or data caches are not present.

Code/Data Location: The AMX kernel code and data and the application code and data have been located in SSRAM without any attempt to locate either for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Developer Suite C/C++ v1.2 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-otime* switch optimizes code for speed
- default *-o2* function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4T code only
- code executes in **ARM and Thumb states**
- AMX 4-Thumb and application code compiled with interworking
- AMX 4-Thumb procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-Thumb task switch, excluding the events leading to the switch, takes 6.4 μ s. This value was derived by averaging 100 task switch measurements. The task switch time is completely determined by the access constraints of the memory system.

**AMX™ 4-Thumb Multitasking Executive
Timing Data**

for

80 MHz ARM966E (Thumb state) and ARM Ltd ADS C

November 1, 2002

Measurement (note 1)	Timing Diagram	T1 µs	T2 µs	T3 µs	T4 µs
Synchronization (note 2)					
Wait/Wake	(page 5)	15.5	12.0	14.1	11.4
Mailbox	(page 6)	15.5	26.3	35.2	12.3
Clock Service (note 2)					
Timers	(page 7)	6.6	22.7	13.6	
Task Delay	(page 8)	18.5	40.1		
Interrupt Service (note 2)					
Handler Latency	(page 10)	3.0	1.6	3.0	1.6

Interrupt Latency: The worst case AMX 4-Thumb interrupt latency is 3.5 µs.

The worst case handler latency for an AMX ISP coded in assembly language is 3.0 µs.

The worst case handler latency for an AMX ISP coded in C is 3.0 µs.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Each test is performed 100 times and the average of all measured values is reported.

AMX™ 4-ARM Multitasking Executive Timing Data

for

600 MHz XScale 80321 (ARM state) and ARM Ltd ADS C

November 1, 2002

Processor: XScale 80321 processor operating at 600 MHz with an internal bus frequency of 200 MHz on the Intel IQ80321 Evaluation Platform configured to derive all timing from a fundamental 66 MHz input clock.

Memory: 128 Mbyte of 64-bit PC1600 Double Data Rate (DDR) SDRAM effectively providing 200 MHz operation.

Cache: The on-chip 32K instruction cache and the 32K data cache are always enabled, unless otherwise indicated, so that times recorded are typical of real time applications. External (secondary) instruction and/or data caches are not present.

Code/Data Location: No attempt is made to locate the AMX kernel code and data or the application code and data for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Developer Suite C/C++ v1.2 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-otime* switch optimizes code for speed
- default *-o2* function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4 code only
- code executes in **ARM state only** (no Thumb code execution)
- AMX 4-ARM procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-ARM task switch, excluding the events leading to the switch, takes 2.7 μ s. This value was derived by averaging 100 task switch measurements made with instruction and data caches enabled. This task switch time rises to 8.3 μ s if the caches remain enabled but are invalidated prior to each task switch measurement. With the caches disabled, the task switch time rises to 12.9 μ s and is completely determined by the access constraints of the memory system.

AMX™ 4-ARM Multitasking Executive Timing Data

for

600 MHz XScale 80321 (ARM state) and ARM Ltd ADS C

November 1, 2002

Measurement <small>(note 1)</small>	Timing Diagram	T1 μ s	T2 μ s	T3 μ s	T4 μ s
Synchronization <small>(note 2)</small>					
Wait/Wake	(page 5)	6.4	6.7	5.0	3.6
Mailbox	(page 6)	6.4	13.7	13.4	3.7
Clock Service <small>(note 2)</small>					
Timers	(page 7)	3.4	10.2	2.9	
Task Delay	(page 8)	10.2	16.0		
Interrupt Service <small>(note 3)</small>					
Handler Latency	(page 10)	3.6	<0.1	3.9	<0.1

Interrupt Latency: The worst case AMX 4-ARM interrupt latency is 6.5 μ s.

The worst case handler latency for an AMX ISP coded in assembly language is 3.6 μ s.

The worst case handler latency for an AMX ISP coded in C is 3.9 μ s.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Instruction and data caches are enabled. Each test is performed 100 times and the average of all measured values is reported.

Note 3: Instruction and data caches are enabled so that the ISP benefits from caching. However, the caches are invalidated before each measurement so that WORST case handler latencies are measured. Each test is performed 100 times and the average of all measured values is reported.

AMX™ 4-Thumb Multitasking Executive Timing Data

for

600 MHz XScale 80321 (Thumb state) and ARM Ltd ADS C

November 1, 2002

Processor: XScale 80321 processor operating at 600 MHz with an internal bus frequency of 200 MHz on the Intel IQ80321 Evaluation Platform configured to derive all timing from a fundamental 66 MHz input clock.

Memory: 128 Mbyte of 64-bit PC1600 Double Data Rate (DDR) SDRAM effectively providing 200 MHz operation.

Cache: The on-chip 32K instruction cache and the 32K data cache are always enabled, unless otherwise indicated, so that times recorded are typical of real time applications. External (secondary) instruction and/or data caches are not present.

Code/Data Location: No attempt is made to locate the AMX kernel code and data or the application code and data for best performance. Code and data reside in memory according to the order in which modules are encountered by the linker, again typical of real applications.

C Compiler: ARM Developer Suite C/C++ v1.2 with:

- no debug information in code modules
- absolute addresses for all code and data references
- register based argument passing conventions
- *-Otime* switch optimizes code for speed
- default *-O2* function in-lining optimizations
- no optimizations using run-time profiling
- no processor specific code; use generic ARM v4T code only
- code executes in **ARM and Thumb states**
- AMX 4-Thumb and application code compiled with interworking
- AMX 4-Thumb procedure parameter validation enabled
- no AMX application scheduler hooks installed

Task Switch: An AMX 4-Thumb task switch, excluding the events leading to the switch, takes 2.7 μ s. This value was derived by averaging 100 task switch measurements made with instruction and data caches enabled. This task switch time rises to 8.5 μ s if the caches remain enabled but are invalidated prior to each task switch measurement. With the caches disabled, the task switch time rises to 13.0 μ s and is completely determined by the access constraints of the memory system.

**AMX™ 4-Thumb Multitasking Executive
Timing Data**

for

600 MHz XScale 80321 (Thumb state) and ARM Ltd ADS C

November 1, 2002

Measurement (note 1)	Timing Diagram	T1 µs	T2 µs	T3 µs	T4 µs
Synchronization (note 2)					
Wait/Wake	(page 5)	6.4	7.1	5.6	3.6
Mailbox	(page 6)	6.4	14.1	14.2	3.8
Clock Service (note 2)					
Timers	(page 7)	3.7	10.7	3.2	
Task Delay	(page 8)	11.0	16.7		
Interrupt Service (note 3)					
Handler Latency	(page 10)	3.6	0.2	3.9	0.2

Interrupt Latency: The worst case AMX 4-Thumb interrupt latency is 7.1 µs.

The worst case handler latency for an AMX ISP coded in assembly language is 3.6 µs.

The worst case handler latency for an AMX ISP coded in C is 3.9 µs.

Note 1: Refer to the referenced page in the AMX Timing Guide for the applicable timing diagram to which measurements T1, T2, T3 and T4 apply.

Note 2: Instruction and data caches are enabled. Each test is performed 100 times and the average of all measured values is reported.

Note 3: Instruction and data caches are enabled so that the ISP benefits from caching. However, the caches are invalidated before each measurement so that WORST case handler latencies are measured. Each test is performed 100 times and the average of all measured values is reported.