

Tracking the Virtual World

Synopsys:

For many years the JTAG interface has been used for ARM-based SoC debugging. With this JTAG style debugging, the developer has been granted the ability to debug software at the high-level. Breakpoints could be set, registers and memory could be accessed, and the processor could be started and stopped at will. Thus the static state of the system was made visible to the engineer. It seemed the only thing missing was visibility to a history of executed cycles. With the introduction of the Embedded Trace Macrocell, commonly known as ETM, this requirement had been fulfilled. Offering a program flow trace by way of a compressed trace port, the developer could now see a history of executed cycles. Ah, was this freedom from static mode debugging?

ETM enabled the capture of virtual addresses where the processor had executed, enabling the program flow trace. This process works great for systems using a statically mapped MMU(memory management unit), where each virtual address uniquely maps to a physical address. But with many of today's popular operating systems, such as Linux or Windows CE, a dynamic MMU is used. In this case the mapping of the virtual address depends on the active process, requiring additional smarts to be built into the debugger.

This article outlines the method which enables the Lauterbach debugger to analyze a program trace, despite these process-dependent virtual addresses. But before tackling the topic at hand, let's take a deeper look at this ETM technology.

The ETM is a core that sits next to the CPU. Its main purpose is to monitor the cycles executed by the CPU and to present, via a dedicated trace port, the program flow information to the used debugger. Thus the ETM only outputs trace information relating to discontinuities in program flow, such as branches or interrupts, which directs the CPU away from its sequential execution path. With this information, and the knowledge of what code existed in memory at the time of execution, a complete trace history can be reconstructed. To display the complete instruction flow, the debugger reads the corresponding instructions from the target memory in order to fill in the gaps between the execution discontinuities that have been recorded (See Fig. 1).

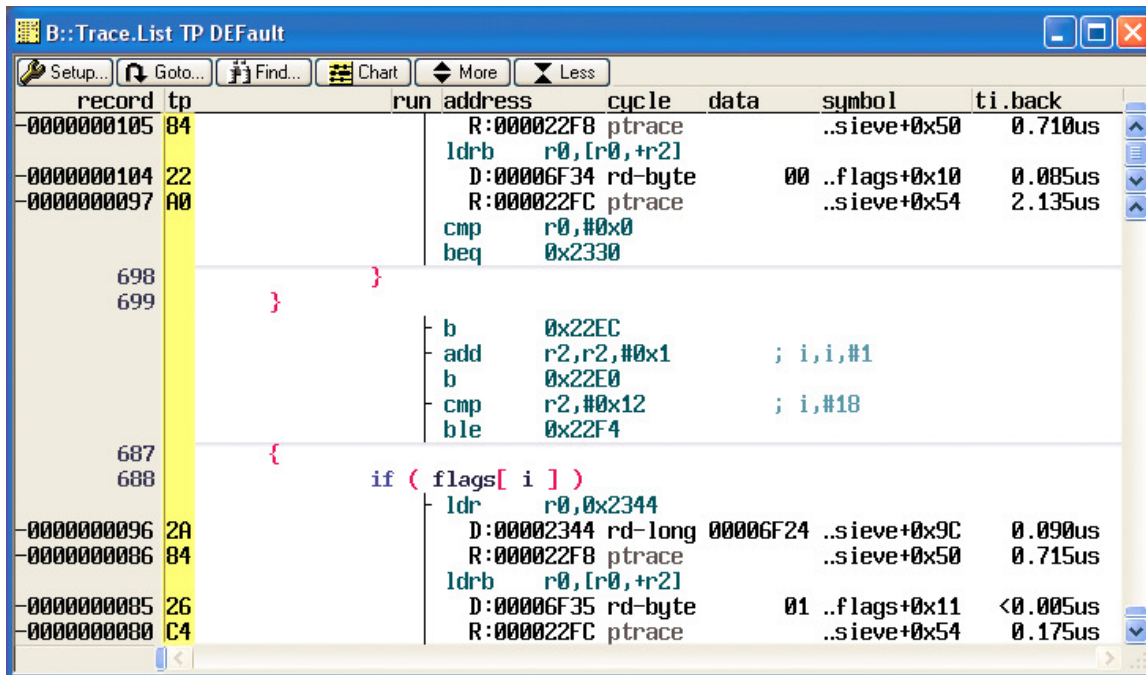


Fig. 1: The yellow area to the left of the screenshot displays the raw trace data as it was recorded by the debugger's trace module. The actual program flow is displayed on the right, as it has been reconstructed by the debugger using the code from the target memory.

In addition to the program flow, the ETM also has the ability to supply information relating to data accesses. However, this can cause problems with the program flow information do to the large amount of information that is required (i.e. a long access would require a 32-bit address and 32-bit data). Too many data accesses within a short period of time can cause an overloading of the ETM trace port, resulting in gaps in the recorded program flow.

Since the ETM is assigned directly to the CPU, it always delivers virtual addresses. Which is appropriate since the code is executing in this address space, as well as all of the variables, pointers, stack, etc... In this case we do not care which physical address a given process runs. As the trace and memory are presented in the same way the processor sees them, which is most desirable.

But what happens if MMU is reprogrammed at run time causing the processor's view to change during execution? In this scenario, some virtual addresses can become invalid, others become valid, and yet others could be mapped to different code than when execution began. Since the ETM only displays the virtual branch destination addresses, and not the code that has actually been executed, it now becomes difficult to reconstruct the program flow.

Operating systems utilizing a dynamically managed MMU are typical examples of this scenario. These include, Windows CE, Linux, QNX, LynxOS and Symbian OS. With these embedded operating systems, each process receives a separate virtual address space. Physically, these obviously reside in different memory areas. But virtually, every process uses the same address

space, which typically begins with address zero. Therefore the processor only works in the address space of the currently active process. During a process change, the processor's MMU is reprogrammed such that it then sees and handles a completely different process (code and data) within the same virtual address space.

So what happens if a trace recording includes one or more process switches?

Let's assume the system first runs in process "A." In this case the ETM records the program execution using the virtual addresses of process "A." Now assume a process switch occurs causing the MMU to be reprogrammed so that process "B" is now visible to the CPU, and process "B" is now being executed. The ETM now records the program execution with the virtual addresses of process "B." If we now stop the processor along with the trace recording and attempt to reconstruct the program flow, two problems arise:

Problem 1: Since the two processes use the same virtual address space, both processes will appear in the trace buffer to be executing from the same address space. Although, process "A" and process "B" are really two different processes, each with its own unique code. It is now the task of the debugger to distinguish between which parts of the trace belong to which process.

Problem 2: As mentioned previously, the ETM only records branch destination addresses, and not the code that was actually executed. Thus the debugger depends on its ability to access the executed program from target memory in order to reconstruct the program flow. However, at the time of the reconstruction, the processor is in the context of process "B" and cannot directly access the code in process "A." Therefore reconstruction of process "A" would not be possible. Thus the debugger must be allowed to access the code from process "A," even though the current context is in process "B."

Let us first address problem 1. To correctly correlate the collected branch trace history back to the correct piece of executed code, it is important the debugger knows exactly where in the trace history the process switches had occurred. This requires the trace system to include a mechanism by which the process switches can be collected during real-time execution. Typically, each operating system has a variable that shows which process is currently active at a given time. This process variable is written with a new value (relating to the new process) whenever a process switch takes place. By recording the write accesses to this process variable, the debugger has the ability to decipher which sections of the trace history correlate to which process. From the point of one of these process variable writes, until the next time it is written, everything runs within the context of the described process (See Fig. 2).

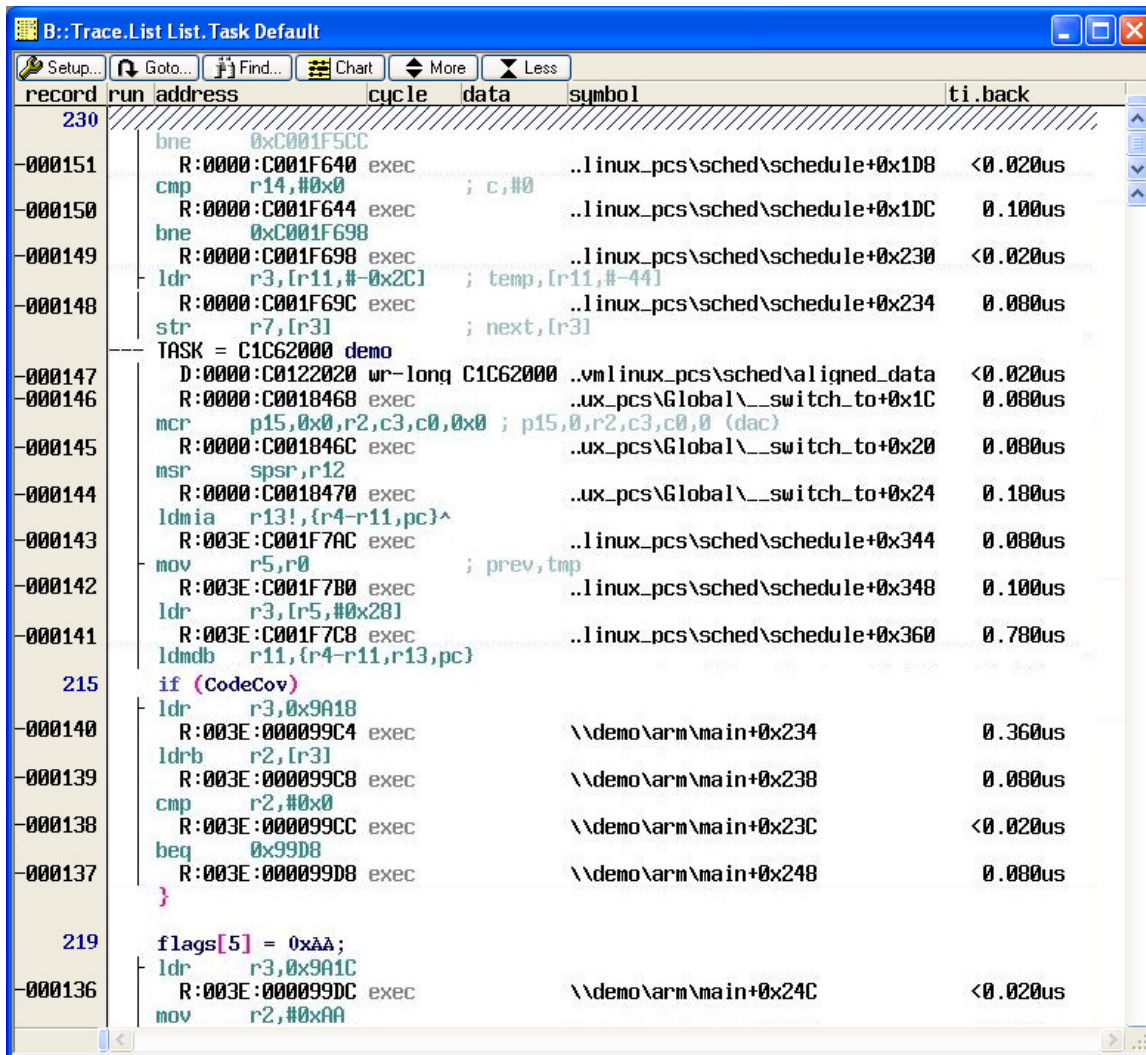


Fig. 2: This screenshot shows a trace recording in which a process switch has occurred.

So in order to gain visibility to these context switches, the ETM must be programmed in such a way that the write accesses to these process variables are recorded in addition to the normal program flow information. It must also be possible to filter out the other, non-process related, data accesses. As we discussed earlier, the collection of too many data accesses can exceed the transmission capacity of the ETM port.

This method allows the debugger to assign the individual trace entries to particular processes. However, there are a couple of disadvantages. Firstly, in order to assign a segment of trace to a known process, the process variable must already have been written/known. So the segment of trace information from the start of the trace collection until the first write to the process variable cannot be assigned to a known process. As there is no way of telling which process was executing prior to the first recorded process switch. Secondly, there must be an available,

correctly programmed, trace filter that can be used for the sole purpose of collecting the accesses to the process variable.

With the introduction of ETM version 1.2, hooks were built into the trace system, enabling yet another method by which the process switches can be tracked. ETM v1.2 includes an additional register, “context ID,” which contains the process ID. It is the responsibility of the operating system to update this “context ID” register whenever there is a process switch. The contents of this register can be collected along with each target branch address, as part of the program flow trace information (See Fig. 3). Given this additional information, now the debugger knows the corresponding process ID for each trace entry, without having to first search for an access to the process variable. Always having knowledge of the process ID, even the beginning of the trace capture will be assigned to its respective process.

record	run	address	cycle	data	symbol	ti.back
-17248414	orr	r2,r1,r2,ls1 #0x8				
-17248414	R:0027:80019EE4	exes			..rn\Global\DoProcessSwitch+0x10	0.040us
-17248414	mcr	p15,0x0,r12,c7,c10,0x4				
-17248414	R:0027:80019EE8	exes			..rn\Global\DoProcessSwitch+0x14	0.040us
-17248414	orr	r12,r12,#0x18				
-17248414	R:0027:80019EEC	exes			..rn\Global\DoProcessSwitch+0x18	0.040us
-17248414	mcr	p15,0x0,r12,c2,c0,0x0				
-17248414	R:0027:80019EF0	exes			..rn\Global\DoProcessSwitch+0x1C	0.040us
-17248410	bx	r14				
-17248410	R:0027:800150A8	exes			...ekern\Global\switch_threads_4	0.040us
-17248404	mcr	p15,0x0,r2,c13,c0,0x1				
task: efile.exe:Main (C804EE28)						
-17248404	R:0009:800150AC	exes			..esched_trampoline_hook_address	0.040us
-17248404	b	0x800150E8 ; resched_trampoline				
-17248404	R:0009:800150E8	exes			..kern\Global\resched_trampoline	0.030us
-17248403	ldr	r1,[r0,#0x120]				
-17248403	D:0009:C8000550	rd-long 8006B7B0			..kern\Global\TheScheduler+0x120	<0.001us
-17248397	R:0009:800150EC	exes			..\Global\resched_trampoline+0x4	0.030us
-17248396	ldr	r0,[r0,#0x130]				
-17248396	D:0009:C8000560	rd-long C804EE28			..kern\Global\TheScheduler+0x130	<0.001us
-17248404	R:0009:800150AC	exes			..esched_trampoline_hook_address	0.040us

Fig. 3: A process ID is displayed in addition to the branch destination address.

Needless to say, either method still requires the debugger to be informed how to analyze these entries. The Lauterbach debugger is informed of this additional trace information by way of an OS awareness declaration file. This is an operating system specific file that describes the unique characteristics of the used embedded operating system.

Now that we can reliably determine the virtual addresses and corresponding processes of the collected trace, we now have everything necessary to recreate the complete execution history of the processor. Since the ETM only captures the branch destination addresses, and not the entire program flow. The source code can be read from the target memory and used to fill in the gaps in the program flow, that occur between the collected branch messages.

Reading the code of the current process is no problem, since the processor and MMU will already be configured to allow access. However, to read the code of another process, the debugger must somehow bypass the current MMU settings, as the MMU will not allow access to the virtual addresses of an inactive process.

Hidden in the operating system is information relating to the used virtual to physical translation. Reprogramming of this MMU translation is managed by the operating system during each process switch. By scanning the operating system for this information, the debugger can create a conversion table which includes the mapping between the used virtual to physical addresses for all processes. Thus the debugger now knows the mapping between all virtual to physical addresses for all processes. Now, rather than having to change processes to access the required code from an inactive process, this table can be used to determine the related physical address. And by deactivating the MMU, the debugger can directly access the code of an inactive process by way of its equivalent physical address.

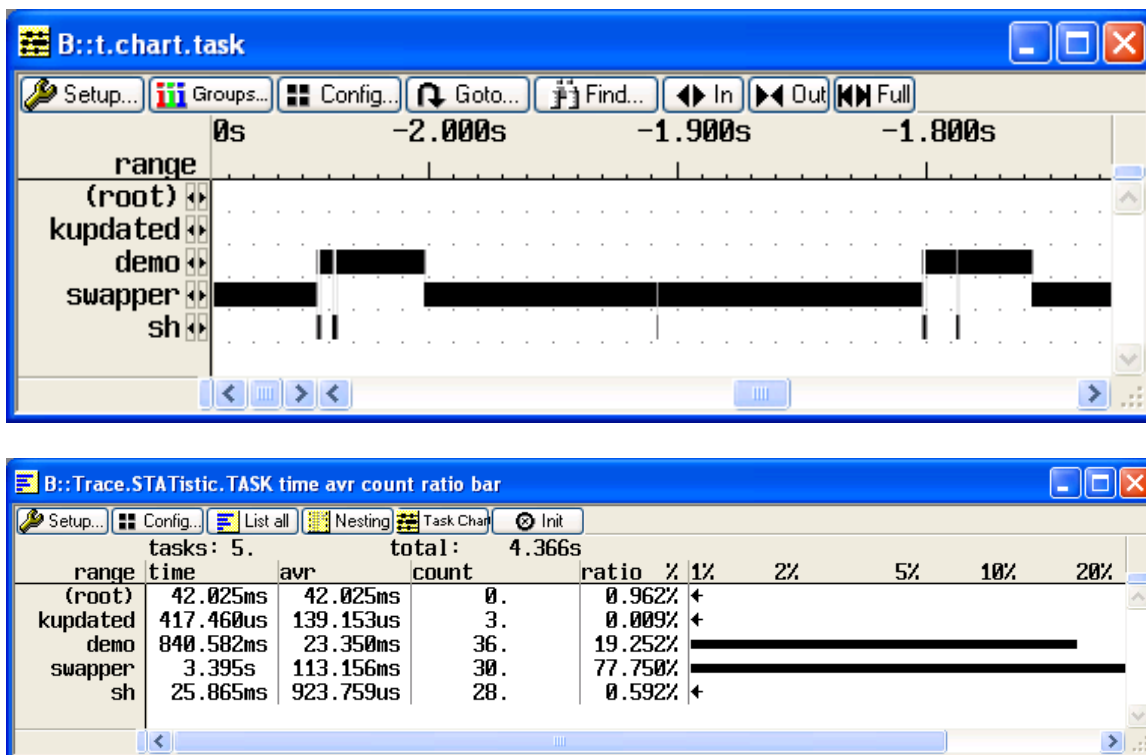


Fig. 4: Graphical and statistical analysis of the process switches recorded in the trace.

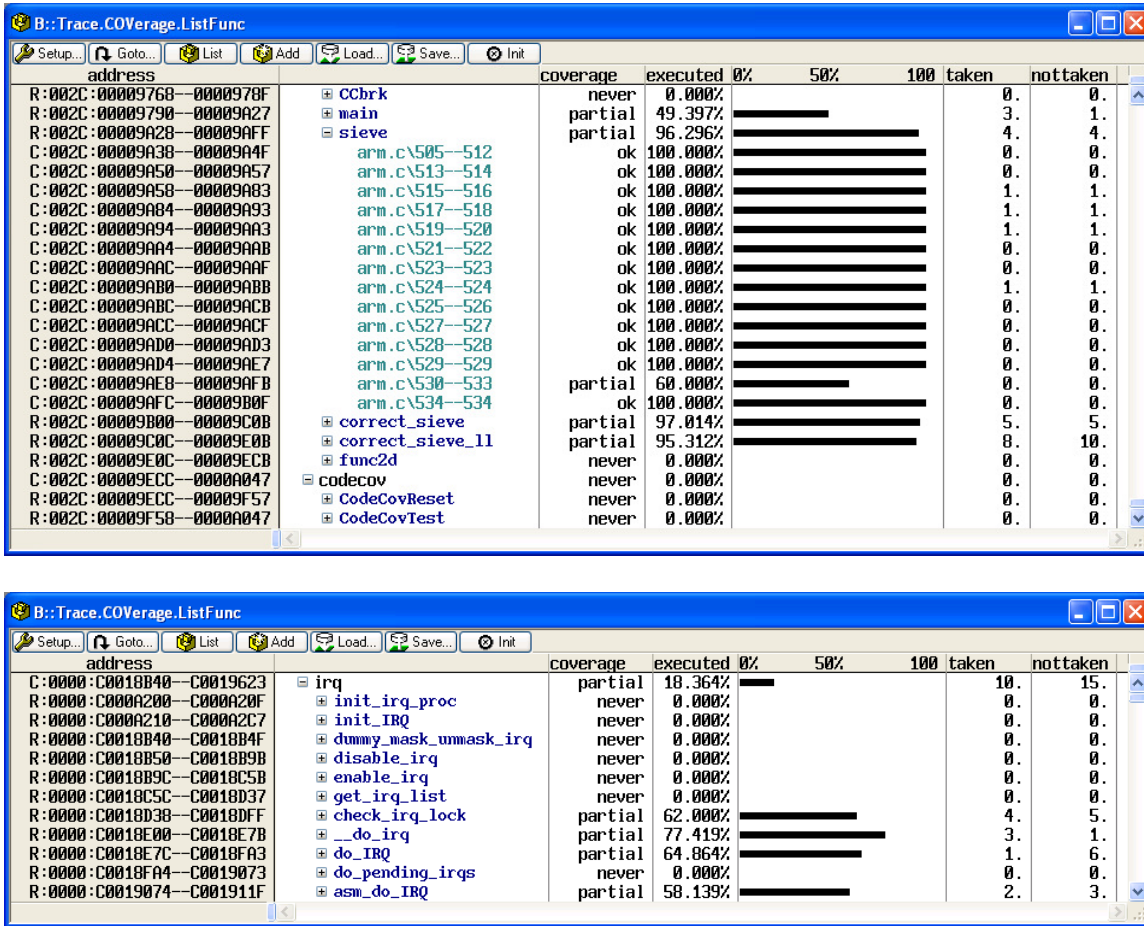


Fig. 5: Code coverage analysis using processes and kernels (segments) based on the trace recording.

Now that the debugger has access to all of the code, including active and inactive processes, the entire program execution can be reconstructed, even through several process switches. With this information, Lauterbach debugging tools can offer the ability to analyze process runtimes (See Fig. 4), cache usage or code coverage (Fig. 5). And with Lauterbach's latest generation of debugging tools, a 2 gigabyte trace memory is available, enabling analysis of quite long run times. For more information, please contact your local Lauterbach representative.

Rudolf Dienstbeck, RTOS Integrations, March 2006