

Tracing Linux Kernel Events via STM

Michael Trimarchi <michael at amarulasolutions.com>

Stefan Kolbinger <Stefan.Kolbinger at lauterbach.com>

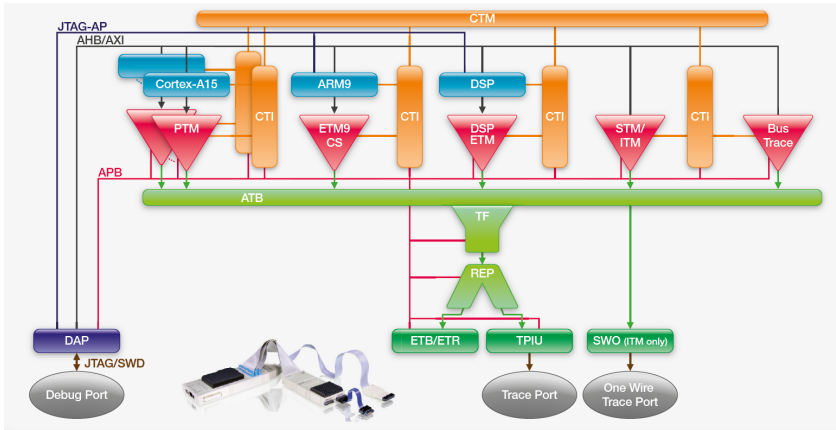
This work was done in collaboration with Lauterbach using
Lauterbach Power Trace

June 27, 2014

Intro ARM Coresight STM

- Coresight is a standard debug interface for debug components in an embedded System on Chip
- STM is a Coresight component that implements a low latency and high bandwidth printf style debug capability
- Scalable solution enabling multi-processors and processes to access STM without being aware of others; STM supports 65,536 channels enabling significant scalability

ARM Coresight STM



Why do we need to use tracing over STM?

- The Software development costs increase compared to hardware one Modern SoC are very complex system
- Optimization and software quality are critical to ensure product success
- Traditional debug and tracing mechanism implement in modern OS are invasive

Linux Tracer with tracepoints


- We use tracepoints
- A tracepoint placed in code provides a hook to call a function (probe) that you can provide at runtime.
- tracepoint can be "off" it has no effect, except for adding a tiny time penalty
- Tracepoints "on", the function you provide is called each time the tracepoint is executed, in the execution context of the caller. You can put tracepoints at important locations in the code.

Tracing Linux Kernel Events via STM

- Use Linux Kernel tracing architecture
- Add new tracepoints hooks
- Send trace data over STM channel
- Max event size is 8 Words (32bit)
- Use one channel x CPU for increase parallelism
- Make it portables to more SoC and with limited changes to the Linux Kernel
- Define a new standard for hardware debugger that can read data from STM channels and can be ported even to no open-source Operating System

Trace with Linux

- PowerTrace II system
- Dual-core cortex A9 system
- STM's compatible core
- Widely used platform
- Compatible with OMAP5 platform
- Supported in mainline



What are we tracing?

- Process/Thread Creation
- Process/Thread Switch
- Process/Thread Deletion
- Process Rename
- Thread State Change
- Interrupt Exception Entry/Exit
- ISR Entry/Exit
- ASID Assign

Tracing Linux kernel events via STM: example

This is an example of the execution of a tracepoint. In this case we have a Process Context Switch, that generate the event to the STM channel

The image displays three windows from the PowerView tool. The top window, titled '[B:~SystemTrace.List /]', shows a list of trace records with columns for record, run address, cycle, data, and ti.back. The second window, titled '[B:~SystemTrace.CustomTrace.stm_message.List/finding spare spare t:back /Track]', shows a list of system trace messages with columns for record, spare, CPU, event name, and timing. The bottom window shows the source code for the tracepoint function, including the function signature and the code that writes data to the trace buffer.

```
void tracing_sched_wakeup_trace(struct trace_array *ta, struct task_struct *task, struct task_struct *prev, unsigned long flags)
```

```
pc = preempt_count();
local_irq_save(flags);
cpu = raw_smp_processor_id();
data = per_cpu_ptr(ctx_trace->trace_buffer.data, cpu);
if (likely(!atomic_read(&data->dfbsltd)))
    tracing_sched_switch_trace(ctx_trace, prev, next, flags, pc);
local_irq_restore(flags);
```

PowerView can show the event on a graphics interface

Pending problems

- tracing is slow (up to 24uS event)
- tracing work on STM 1.0 from TI
- Stm Clock is not optimized (NDA required with TI)
- code is written on top of an old TI driver (cleanup is not sufficient)

Next steps...

- Find someone that has time.. (1)
- Rebase to a newer linux kernel and Coresight implementation (3.16 and Patrick Pratel patchset <https://lkml.org/lkml/2012/12/19/331>)
- Export more tracepoint events or create pluggable backend to the tracer
- Create a graphical visualization of the events on Lauterbach TRACE32(R) Combine hardware tracing information with Linux Kernel Tracing over STM
- Mainline kernel submission??? refer to point (1)