

New debugging concept for symmetric multiprocessing (SMP)

Lauterbach, the leading manufacturer world-wide of high-quality debuggers and real-time trace tools, is presenting its new debugging concept for SMP systems at the ESC 2008. SMP systems use an operating system to distribute processes dynamically to several cores or hardware threads. SMP Linux on a MIPS 34K will be demonstrated live at the fair.

Hardware parallelization

To achieve higher processor performance for complex applications and to save energy, more and more parallelization of tasks is used. The most popular method is to provide several identical execution units that can all process the same task. To most readers, the term "identical execution units" means symmetric multi-core processors. It is easy to imagine that the kernel of an operating system is running permanently on a core and ensures that the application processes are evenly distributed to all cores (see Fig. 1).

However, for the parallelization of tasks not necessarily a multi-core processor is required. Hardware multithreading, for example, is an approach that enables parallelization also for single-core processors. Here, we deal with a basic problem of cores

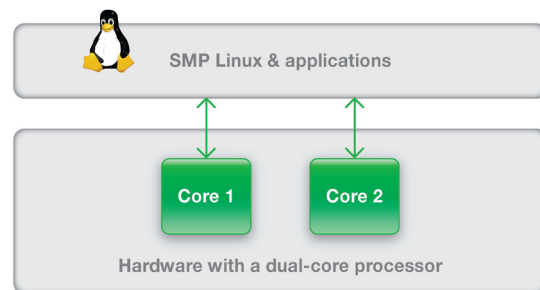


Fig. 1: SMP Linux on a dual-core processor

with pipeline architecture: cache misses or data dependencies between the instructions mean that the pipelined instruction processing has to be stalled in order to wait for the availability of required data. The greater the difference between instruction processing time and memory access time, the higher the performance loss.


Hardware multithreading deals with this situation by making the core process several mutually independent tasks quasi-simultaneously. In the case of an operating system, "tasks" and "processes" are equivalent.

In principle, this works as follows: As soon as a task can no longer be further processed because required data is not yet available, the processing of another task is continued.

Fig. 2 shows the processing of 3 tasks on 3 cores with simple pipeline architecture (at the top) and then (at the bottom) shows the execution on a multithreaded core.

In order for hardware multithreading to work, fast switching between tasks must be possible. This can be done simply by giving each task a separate register set for its context. With this method, it is easy to deduce that the number of register sets provided by the core defines how many tasks can be processed in parallel. For the MIPS 34K the number is five. In this way, one >>

Processing on 3 cores with simple pipeline architecture



Processing on a multithreaded core



Fig. 2: No more stalls in the pipeline, thanks to the quasi-simultaneous processing of several tasks

New debugging concept for symmetric multiprocessing (SMP)

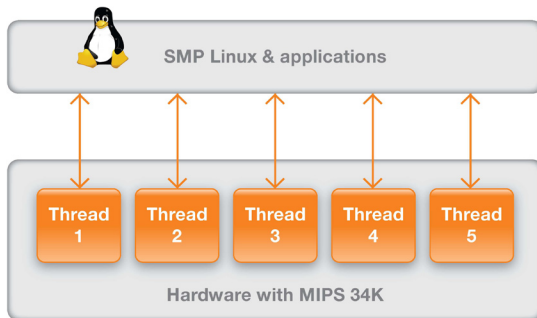


Fig. 3: SMP Linux on a multithreaded core – here MIPS 34K

“execution unit” (a so-called hardware thread) derives from each register set (see Fig. 3).

SMP operating systems

At the hardware level, the parallelization of processes is implemented by symmetric multi-core processors or multithreaded cores. Now software is needed to organize the parallelization. An operating system is normally used for this purpose.

One variant, which is primarily suitable for hardware with identical execution units, is operating systems that implement symmetric multiprocessing (SMP). The central feature of SMP operating systems is the dynamic distribution of task/processes to the available cores or hardware threads at program runtime.

Other features:

- One instance of the operating system operates all cores or hardware threads.
- Each application process can run on each core or hardware thread. As a rule, only the kernel is rigidly assigned to a core or hardware thread.
- All cores or hardware threads have equal rights to request and use resources (e.g.: memory, external interfaces, external devices).
- The operating system provides the functions for distributing resources to the cores or hardware threads.

New debug concept

At present, the TRACE32 concept enables just one core to be debugged with one instance of the debugger (*Core View*). Up to now, this concept has

also worked perfectly for multi-core processors since these processors were also operated as asymmetric multiprocessing systems (AMP). Asymmetric multiprocessing means: An independent instance of an operating system runs on each core. For AMP systems, this always statically defines which process is running on which core. In this way, the debug information can also be uniquely assigned to the corresponding core.

In contrast, on SMP systems, the processes are not assigned to the cores or hardware threads until run-time. For this reason, it no longer makes sense to start a debugging instance specifically for the debugging of a selected core or hardware thread.

System View

As an alternative to the *Core View*, which works very well for AMP systems, there is now the *System View* for SMP systems.

With the *System View*, only one instance of the TRACE32 debugger is started to debug all cores or hardware threads (see Fig. 4). Here too, root of information display is one core or hardware thread. The new aspect here is that the information »

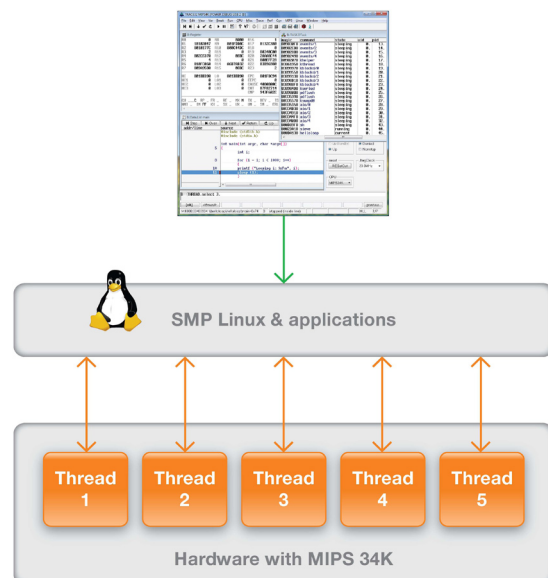


Fig. 4: Only one instance of the TRACE32 debugger is used for debugging all cores or hardware threads.

New debugging concept for symmetric multiprocessing (SMP)

display can be switched by a command to another core or hardware thread.




Fig. 5: The TASK.DTASK window, showing the core or hardware thread to which SMP Linux has assigned a task/process

Process debugging can now be done as follows:

1. The current list of all active task/processes is used to identify the core or hardware thread on which a task/process is running (Fig. 5). Linux uses here the term CPU instead of using core or hardware thread.
2. A command is used to switch the information displayed in the debugger to the required core or hardware thread (Fig. 6).

Common breakpoints

The fact that an SMP operating system does not assign application processes to a core or hardware thread until runtime also has consequences for the setting of on-chip breakpoints.

A simple example: The "sieve" process is to be stopped as soon as it writes to the variable "xyz". In order for the debugger to be able to implement this request, it must program the break logic of all cores or hardware threads for this break condition since it cannot know beforehand the core or hardware thread on which the "sieve" application process will run. This means that even if each core or hardware thread has its own break logic, the debugger programs the breakpoints for the user as if there were only one break logic shared by all.

When setting of software breakpoints, for the implementation of which the original instruction in memory is temporarily overwritten by a break instruction, there are no changes compared to AMP systems. Operating systems protect the address spaces of processes from each other. However, if the same application process is started more than

once, Linux (for example) loads the program code only once. Each instance of the application process thus sees the same program memory as well as the software breakpoints set there. To make sure that the program execution is stopped only in the desired application process, the TRACE32 debugger enables the setting of process-specific software breakpoints.

Summary

The systematic extension of the TRACE32 concept enables Lauterbach to offer its customers simple debugging of embedded designs that use an SMP operating system for controlling several cores or hardware threads. To be able to use this new concept for debugging SMP operating systems, you just need a TRACE32 debugger whose debug cable has not only a license for the processor architecture but also a second license such as a multi-core license.

Following the support for the MIPS 34K, the debugging of SMP systems for ARM and PowerPC architectures is also planned for early 2008.

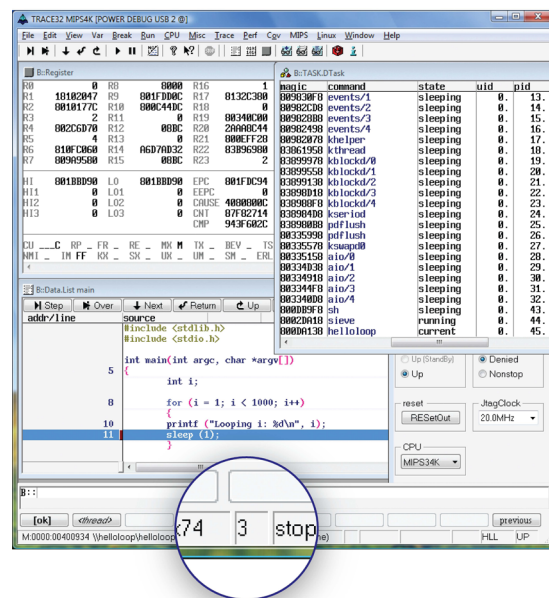


Fig. 6: The TRACE32 window, which always shows only the context of a core or hardware thread. The number of the core or hardware thread is shown in the status line.