

Re: [PATCH 1/4] Blackfin: arch patch for 2.6.18

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-09/msg07302.html>

- *From:* Aubrey <aubreylee@xxxxxxxxxx>
 - *Date:* Tue, 26 Sep 2006 11:42:25 +0800
-

On 9/26/06, Arnd Bergmann <arnd@xxxxxxxxxx> wrote:

On Monday 25 September 2006 17:39, Aubrey wrote:

- > 1) Timer interrupt will call `do_irq()`, then `return_from_int()`.
- >
- > 2) `return_from_int()` will check if there is interrupt pending or signal pending, if so, it will call `schedule_and_signal_from_int()`.
- >
- > 3) `schedule_and_signal_from_int()` will jump to `resume_userspace()`
- >
- > 4) `resume_userspace()` will call `_schedule` to run the user task.

I have a little trouble reading your assembly code, but your `return_from_int()` function should normally not call `schedule_and_signal_from_int()` when the interrupt happened in kernel context (like in the `idle` function):

```
+ /* if not return to user mode, get out */
+ p2.l = lo(IPEND);
+ p2.h = hi(IPEND);
+ r0 = [p2];
+ r1 = 0x17(Z);
+ r2 = ~r1;
+ r2.h = 0;
+ r0 = r2 & r0;
+ r1 = 1;
+ r1 = r0 - r1;
+ r2 = r0 & r1;
+ cc = r2 == 0;
+ if !cc jump 2f;
```

This looks a lot like your `user_mode()` function, so you jump over `schedule_and_signal_from_int()` here.

What you described would be a preemptive kernel (`CONFIG_PREEMPT`), but you clearly don't have that enabled.

Re: [PATCH 1/4] Blackfin: arch patch for 2.6.18

No, `schedule_and_signal_from_int` will be called.

The above code is checking if there are at least two bits set on, if so, `schedule_and_signal_from_int` will be called.

Blackfin supports 3 processor mode: (1) user mode (2) supervisor mode (3) emulation mode. In the kernel space, the processor should be in the supervisor mode. To keep the processor in the supervisor mode, we raise the lowest priority interrupt event. Kernel actually in the interrupt handler of the lowest priority interrupt event. See `arch/blackfin/mach-bf53x/head.S`.

```
=====
/* This section keeps the processor in supervisor mode
 * during kernel boot. Switches to user mode at end of boot.
 * See page 3-9 of Hardware Reference manual for documentation.
 */
```

```
/* EVT15 = _real_start */
```

```
p0.l = lo(EVT15);
p0.h = hi(EVT15);
p1.l = _real_start;
p1.h = _real_start;
[p0] = p1;
csync;
```

```
p0.l = lo(IMASK);
p0.h = hi(IMASK);
p1.l = IMASK_IVG15;
p1.h = 0x0;
[p0] = p1;
csync;
```

```
raise 15;
p0.l = .LWAIT_HERE;
p0.h = .LWAIT_HERE;
reti = p0;
rti;
```

```
=====
So, in the kernel space, there is always one bit in the IPEND register is set. And if there comes a timer interrupt event, in the timer interrupt handler, there should be two bits set in the IPEND register. Therefore, schedule happens in the return_from_int.
```

So, I still say there is no latency here.

–Aubrey

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to `majordomo@xxxxxxxxxxxxxxxxx`

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

Re: [PATCH 1/4] Blackfin: arch patch for 2.6.18