

```

*****
34729 Sun Mar  2 12:33:18 2014
new/usr/src/uts/i86pc/io/pcplusmp/apic.c
4665 pcplusmp open-codes register operations
*****
unchanged portion omitted
529 #endif /* DEBUG */

531 /*
532  * platform_intr_enter
533  *
534  * Called at the beginning of the interrupt service routine to
535  * mask all level equal to and below the interrupt priority
536  * of the interrupting vector.  An EOI should be given to
537  * the interrupt controller to enable other HW interrupts.
538  *
539  * Return -1 for spurious interrupts
540  *
541  */
542 /*ARGSUSED*/
543 static int
544 apic_intr_enter(int ipl, int *vectorp)
545 {
546     uchar_t vector;
547     int nipl;
548     int irq;
549     ulong_t iflag;
550     apic_cpus_info_t *cpu_infop;

552     /*
553      * The real vector delivered is (*vectorp + 0x20), but our caller
554      * subtracts 0x20 from the vector before passing it to us.
555      * (That's why APIC_BASE_VECT is 0x20.)
556      */
557     vector = (uchar_t)*vectorp;

559     /* if interrupted by the clock, increment apic_nsec_since_boot */
560     if (vector == apic_clkvect) {
561         if (!apic_oneshot) {
562             /* NOTE: this is not MT aware */
563             apic_hrttime_stamp++;
564             apic_nsec_since_boot += apic_nsec_per_intr;
565             apic_hrttime_stamp++;
566             last_count_read = apic_hertz_count;
567             apic_redistribute_compute();
568         }

570         /* We will avoid all the book keeping overhead for clock */
571         nipl = apic_ipls[vector];

573         *vectorp = apic_vector_to_irq[vector + APIC_BASE_VECT];

575         apic_reg_ops->apic_write_task_reg(apic_ipltopri[nipl]);
576         apic_reg_ops->apic_send_eoi(0);
577         if (apic_mode == LOCAL_APIC) {
578             #if defined(__amd64)
579                 setcr8((ulong_t)(apic_ipltopri[nipl] >>
580                     APIC_IPL_SHIFT));
581             #else
582                 if (apic_have_32bit_cr8)
583                     setcr8((ulong_t)(apic_ipltopri[nipl] >>
584                         APIC_IPL_SHIFT));
585                 else
586                     LOCAL_APIC_WRITE_REG(APIC_TASK_REG,
587                         (uint32_t)apic_ipltopri[nipl]);
588             #endif
589         }

```

```

586         LOCAL_APIC_WRITE_REG(APIC_EOI_REG, 0);
587     } else {
588         X2APIC_WRITE(APIC_TASK_REG, apic_ipltopri[nipl]);
589         X2APIC_WRITE(APIC_EOI_REG, 0);
590     }

592     return (nipl);
593 }

595     cpu_infop = &apic_cpus[psm_get_cpu_id()];

597     if (vector == (APIC_SPUR_INTR - APIC_BASE_VECT)) {
598         cpu_infop->aci_spur_cnt++;
599         return (APIC_INT_SPURIOUS);
600     }

602     /* Check if the vector we got is really what we need */
603     if (apic_revector_pending) {
604         /*
605          * Disable interrupts for the duration of
606          * the vector translation to prevent a self-race for
607          * the apic_revector_lock.  This cannot be done
608          * in apic_xlate_vector because it is recursive and
609          * we want the vector translation to be atomic with
610          * respect to other (higher-priority) interrupts.
611          */
612         iflag = intr_clear();
613         vector = apic_xlate_vector(vector + APIC_BASE_VECT) -
614             APIC_BASE_VECT;
615         intr_restore(iflag);
616     }

618     nipl = apic_ipls[vector];
619     *vectorp = irq = apic_vector_to_irq[vector + APIC_BASE_VECT];

621     apic_reg_ops->apic_write_task_reg(apic_ipltopri[nipl]);
622     if (apic_mode == LOCAL_APIC) {
623         #if defined(__amd64)
624             setcr8((ulong_t)(apic_ipltopri[nipl] >> APIC_IPL_SHIFT));
625         #else
626             if (apic_have_32bit_cr8)
627                 setcr8((ulong_t)(apic_ipltopri[nipl] >>
628                     APIC_IPL_SHIFT));
629             else
630                 LOCAL_APIC_WRITE_REG(APIC_TASK_REG,
631                     (uint32_t)apic_ipltopri[nipl]);
632         #endif
633     } else {
634         X2APIC_WRITE(APIC_TASK_REG, apic_ipltopri[nipl]);
635     }

637     cpu_infop->aci_current[nipl] = (uchar_t)irq;
638     cpu_infop->aci_curipl = (uchar_t)nipl;
639     cpu_infop->aci_isr_in_progress |= 1 << nipl;

641     /*
642      * apic_level_intr could have been assimilated into the irq struct.
643      * but, having it as a character array is more efficient in terms of
644      * cache usage.  So, we leave it as is.
645      */
646     if (!apic_level_intr[irq]) {
647         apic_reg_ops->apic_send_eoi(0);
648         if (apic_mode == LOCAL_APIC) {
649             LOCAL_APIC_WRITE_REG(APIC_EOI_REG, 0);
650         } else {
651             X2APIC_WRITE(APIC_EOI_REG, 0);

```

```

650     }
620 }

622 #ifdef  DEBUG
623     APIC_DEBUG_BUF_PUT(vector);
624     APIC_DEBUG_BUF_PUT(irq);
625     APIC_DEBUG_BUF_PUT(nipl);
626     APIC_DEBUG_BUF_PUT(psm_get_cpu_id());
627     if ((apic_stretch_interrupts) && (apic_stretch_ISR & (1 << nipl)))
628         drv_usecwait(apic_stretch_interrupts);

630     if (apic_break_on_cpu == psm_get_cpu_id())
631         apic_break();
632 #endif /* DEBUG */
633     return (nipl);
634 }
    unchanged_portion_omitted

650 /*
651  * Any changes made to this function must also change X2APIC
652  * version of intr_exit.
653  */
654 void
655 apic_intr_exit(int prev_ipl, int irq)
656 {
657     apic_cpus_info_t *cpu_infop;

659     apic_reg_ops->apic_write_task_reg(apic_ipltopri[prev_ipl]);
690 #if defined(__amd64)
691     setcr8((ulong_t)(apic_ipltopri[prev_ipl] >> APIC_IPL_SHIFT));
692 #else
693     if (apic_have_32bit_cr8)
694         setcr8((ulong_t)(apic_ipltopri[prev_ipl] >> APIC_IPL_SHIFT));
695     else
696         apicadr[APIC_TASK_REG] = apic_ipltopri[prev_ipl];
697 #endif

661     APIC_INTR_EXIT();
662 }
    unchanged_portion_omitted

686 /*
687  * Mask all interrupts below or equal to the given IPL.
688  * Any changes made to this function must also change X2APIC
689  * version of setspl.
690  */
691 static void
692 apic_setspl(int ipl)
693 {
694     apic_reg_ops->apic_write_task_reg(apic_ipltopri[ipl]);
732 #if defined(__amd64)
733     setcr8((ulong_t)(apic_ipltopri[ipl] >> APIC_IPL_SHIFT));
734 #else
735     if (apic_have_32bit_cr8)
736         setcr8((ulong_t)(apic_ipltopri[ipl] >> APIC_IPL_SHIFT));
737     else
738         apicadr[APIC_TASK_REG] = apic_ipltopri[ipl];
739 #endif

696     /* interrupts at ipl above this cannot be in progress */
697     apic_cpus[psm_get_cpu_id()].aci_ISR_in_progress &= (2 << ipl) - 1;
698     /*
699     * this is a patch fix for the ALR QSMP P5 machine, so that interrupts
700     * have enough time to come in before the priority is raised again
701     * during the idle() loop.
702     */

```

```

703     if (apic_setspl_delay)
704         (void) apic_reg_ops->apic_get_pri();
705 }
    unchanged_portion_omitted

```