

```

*****
41057 Mon May 5 11:11:06 2014
new/usr/src/uts/common/io/iprb/iprb.c
4778 iprb shouldn't abuse ddi_get_time(9f)
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Garrett D'Amore <garrett@damore.org>
*****
unchanged_portion_omitted

258 int
259 iprb_attach(dev_info_t *dip)
260 {
261     iprb_t        *ip;
262     uint16_t      w;
263     int           i;
264     mac_register_t *macp;

266     ip = kmem_zalloc(sizeof(*ip), KM_SLEEP);
267     ddi_set_driver_private(dip, ip);
268     ip->dip = dip;

270     list_create(&ip->mcast, sizeof(struct iprb_mcast),
271               offsetof(struct iprb_mcast, node));

273     /* we don't support high level interrupts, so we don't need cookies */
274     mutex_init(&ip->culock, NULL, MUTEX_DRIVER, NULL);
275     mutex_init(&ip->ru lock, NULL, MUTEX_DRIVER, NULL);

277     if (pci_config_setup(dip, &ip->pcih) != DDI_SUCCESS) {
278         iprb_error(ip, "unable to map configuration space");
279         iprb_destroy(ip);
280         return (DDI_FAILURE);
281     }

283     if (ddi_regs_map_setup(dip, 1, &ip->regs, 0, 0, &acc_attr,
284                           &ip->regsh) != DDI_SUCCESS) {
285         iprb_error(ip, "unable to map device registers");
286         iprb_destroy(ip);
287         return (DDI_FAILURE);
288     }

290     /* Reset, but first go into idle state */
291     PUT32(ip, CSR_PORT, PORT_SEL_RESET);
292     drv_usecwait(10);
293     PUT32(ip, CSR_PORT, PORT_SW_RESET);
294     drv_usecwait(10);
295     PUT8(ip, CSR_INTCTL, INTCTL_MASK);
296     (void) GET8(ip, CSR_INTCTL);

298     /*
299      * Precalculate watchdog times.
300      */
301     ip->tx_timeout = TX_WATCHDOG;
302     ip->rx_timeout = RX_WATCHDOG;
301     ip->tx_timeout = drv_usec2ohz(TX_WATCHDOG * 1000000);
302     ip->rx_timeout = drv_usec2ohz(RX_WATCHDOG * 1000000);

304     iprb_identify(ip);

306     /* Obtain our factory MAC address */
307     w = iprb_eeprom_read(ip, 0);
308     ip->factaddr[0] = w & 0xff;
309     ip->factaddr[1] = w >> 8;
310     w = iprb_eeprom_read(ip, 1);
311     ip->factaddr[2] = w & 0xff;
312     ip->factaddr[3] = w >> 8;

```

```

313     w = iprb_eeprom_read(ip, 2);
314     ip->factaddr[4] = w & 0xff;
315     ip->factaddr[5] = w >> 8;
316     bcopy(ip->factaddr, ip->curraddr, 6);

318     if (ip->resumebug) {
319         /*
320          * Generally, most devices we will ever see will
321          * already have fixed firmware. Since I can't verify
322          * the validity of the fix (no suitably downrev
323          * hardware), we'll just do our best to avoid it for
324          * devices that exhibit this behavior.
325          */
326         if ((iprb_eeprom_read(ip, 10) & 0x02) == 0) {
327             /* EEPROM fix was already applied, assume safe. */
328             ip->resumebug = B_FALSE;
329         }
330     }

332     if ((iprb_eeprom_read(ip, 3) & 0x3) != 0x3) {
333         cmn_err(CE_CONT, "?Enabling RX errata workaround.\n");
334         ip->rxhangbug = B_TRUE;
335     }

337     /* Determine whether we have an MII or a legacy 80c24 */
338     w = iprb_eeprom_read(ip, 6);
339     if ((w & 0x3f00) != 0x0600) {
340         if ((ip->miih = mii_alloc(ip, dip, &iprb_mii_ops)) == NULL) {
341             iprb_error(ip, "unable to allocate MII ops vector");
342             iprb_destroy(ip);
343             return (DDI_FAILURE);
344         }
345         if (ip->canpause) {
346             mii_set_pauseable(ip->miih, B_TRUE, B_FALSE);
347         }
348     }

350     /* Allocate cmds and tx region */
351     for (i = 0; i < NUM_TX; i++) {
352         /* Command blocks */
353         if (iprb_dma_alloc(ip, &ip->cmds[i], CB_SIZE) != DDI_SUCCESS) {
354             iprb_destroy(ip);
355             return (DDI_FAILURE);
356         }
357     }

359     for (i = 0; i < NUM_TX; i++) {
360         iprb_dma_t *cb = &ip->cmds[i];
361         /* Link the command blocks into a ring */
362         PUTCB32(cb, CB_LNK_OFFSET, (ip->cmds[(i + 1) % NUM_TX].paddr));
363     }

365     for (i = 0; i < NUM_RX; i++) {
366         /* Rx packet buffers */
367         if (iprb_dma_alloc(ip, &ip->rxbuf[i], RFD_SIZE) != DDI_SUCCESS) {
368             iprb_destroy(ip);
369             return (DDI_FAILURE);
370         }
371     }
372     if (iprb_dma_alloc(ip, &ip->stats, STATS_SIZE) != DDI_SUCCESS) {
373         iprb_destroy(ip);
374         return (DDI_FAILURE);
375     }

377     if (iprb_add_intr(ip) != DDI_SUCCESS) {
378         iprb_destroy(ip);

```

```

379         return (DDI_FAILURE);
380     }

382     if ((macp = mac_alloc(MAC_VERSION)) == NULL) {
383         iprb_error(ip, "unable to allocate mac structure");
384         iprb_destroy(ip);
385         return (DDI_FAILURE);
386     }

388     macp->m_type_ident = MAC_PLUGIN_IDENT_ETHER;
389     macp->m_driver = ip;
390     macp->m_dip = dip;
391     macp->m_src_addr = ip->curraddr;
392     macp->m_callbacks = &iprb_m_callbacks;
393     macp->m_min_sdu = 0;
394     macp->m_max_sdu = ETHERMTU;
395     macp->m_margin = VLAN_TAGSZ;
396     if (mac_register(macp, &ip->mach) != 0) {
397         iprb_error(ip, "unable to register mac with framework");
398         mac_free(macp);
399         iprb_destroy(ip);
400         return (DDI_FAILURE);
401     }

403     mac_free(macp);
404     return (DDI_SUCCESS);
405 }

    unchanged portion omitted

689 void
690 iprb_cmd_reclaim(iprb_t *ip)
691 {
692     while (ip->cmd_count) {
693         iprb_dma_t *cb = &ip->cmds[ip->cmd_tail];

695         SYNC_CB(cb, CB_STS_OFFSET, 2, DDI_DMA_SYNC_FORKERNEL);
696         if ((GETCB16(cb, CB_STS_OFFSET) & CB_STS_C) == 0) {
697             break;
698         }

700         ip->cmd_tail++;
701         ip->cmd_tail %= NUM_TX;
702         ip->cmd_count--;
703         if (ip->cmd_count == 0) {
704             ip->tx_wdog = 0;
705         } else {
706             ip->tx_wdog = gethrtime();
707             ip->tx_wdog = ddi_get_time();
708         }
709     }

    unchanged portion omitted

724 int
725 iprb_cmd_submit(iprb_t *ip, uint16_t cmd)
726 {
727     iprb_dma_t *ncb = &ip->cmds[ip->cmd_head];
728     iprb_dma_t *lcb = &ip->cmds[ip->cmd_last];

730     /* If this command will consume the last CB, interrupt when done */
731     ASSERT((ip->cmd_count) < NUM_TX);
732     if (ip->cmd_count == (NUM_TX - 1)) {
733         cmd |= CB_CMD_I;
734     }

736     /* clear the status entry */

```

```

737     PUTCB16(ncb, CB_STS_OFFSET, 0);

739     /* suspend upon completion of this new command */
740     cmd |= CB_CMD_S;
741     PUTCB16(ncb, CB_CMD_OFFSET, cmd);
742     SYNC_CB(ncb, 0, 0, DDI_DMA_SYNC_FORDEV);

744     /* clear the suspend flag from the last submitted command */
745     SYNC_CB(lcb, CB_CMD_OFFSET, 2, DDI_DMA_SYNC_FORKERNEL);
746     PUTCB16(lcb, CB_CMD_OFFSET, GETCB16(lcb, CB_CMD_OFFSET) & ~CB_CMD_S);
747     SYNC_CB(lcb, CB_CMD_OFFSET, 2, DDI_DMA_SYNC_FORDEV);

750     /*
751     * If the chip has a resume bug, then we need to try this as a work
752     * around. Some anecdotal evidence is that this will help solve
753     * the resume bug. Its a performance hit, but only if the EEPROM
754     * is not updated. (In theory we could do this only for 10Mbps HDX,
755     * but since it should just about never get used, we keep it simple.)
756     */
757     if (ip->resumebug) {
758         if (iprb_cmd_ready(ip) != DDI_SUCCESS)
759             return (DDI_FAILURE);
760         PUT8(ip, CSR_CMD, CUC_NOP);
761         (void) GET8(ip, CSR_CMD);
762         drv_usecwait(1);
763     }

765     /* wait for the SCB to be ready to accept a new command */
766     if (iprb_cmd_ready(ip) != DDI_SUCCESS)
767         return (DDI_FAILURE);

769     /*
770     * Finally we can resume the CU. Note that if this the first
771     * command in the sequence (i.e. if the CU is IDLE), or if the
772     * CU is already busy working, then this CU resume command
773     * will not have any effect.
774     */
775     PUT8(ip, CSR_CMD, CUC_RESUME);
776     (void) GET8(ip, CSR_CMD); /* flush CSR */

778     ip->tx_wdog = gethrtime();
778     ip->tx_wdog = ddi_get_time();
779     ip->cmd_last = ip->cmd_head;
780     ip->cmd_head++;
781     ip->cmd_head %= NUM_TX;
782     ip->cmd_count++;

784     return (DDI_SUCCESS);
785 }

    unchanged portion omitted

1009 void
1010 iprb_update_stats(iprb_t *ip)
1011 {
1012     iprb_dma_t *sp = &ip->stats;
1013     hrtime_t tstamp;
1013     time_t tstamp;
1014     int i;

1016     ASSERT(mutex_owned(&ip->culock));

1018     /* Collect the hardware stats, but don't keep redoing it */
1019     tstamp = gethrtime();
1020     if (tstamp / NANOSPEC == ip->stats_time / NANOSPEC)
1021     if ((tstamp = ddi_get_time()) == ip->stats_time) {

```

```

1021         return;
1021     }

1023     PUTSTAT(sp, STATS_DONE_OFFSET, 0);
1024     SYNCSTATS(sp, 0, 0, DDI_DMA_SYNC_FORDEV);

1026     if (iprb_cmd_ready(ip) != DDI_SUCCESS)
1027         return;
1028     PUT32(ip, CSR_GEN_PTR, sp->paddr);
1029     PUT8(ip, CSR_CMD, CUC_STATSBASE);
1030     (void) GET8(ip, CSR_CMD);

1032     if (iprb_cmd_ready(ip) != DDI_SUCCESS)
1033         return;
1034     PUT8(ip, CSR_CMD, CUC_STATS_RST);
1035     (void) GET8(ip, CSR_CMD); /* flush wb */

1037     for (i = 10000; i; i -= 10) {
1038         SYNCSTATS(sp, 0, 0, DDI_DMA_SYNC_FORKERNEL);
1039         if (GETSTAT(sp, STATS_DONE_OFFSET) == STATS_RST_DONE) {
1040             /* yay stats are updated */
1041             break;
1042         }
1043         drv_usecwait(10);
1044     }
1045     if (i == 0) {
1046         iprb_error(ip, "time out acquiring hardware statistics");
1047         return;
1048     }

1050     ip->ex_coll += GETSTAT(sp, STATS_TX_MAXCOL_OFFSET);
1051     ip->late_coll += GETSTAT(sp, STATS_TX_LATECOL_OFFSET);
1052     ip->ufl0 += GETSTAT(sp, STATS_TX_UFLO_OFFSET);
1053     ip->defer_xmt += GETSTAT(sp, STATS_TX_DEFER_OFFSET);
1054     ip->one_coll += GETSTAT(sp, STATS_TX_ONECOL_OFFSET);
1055     ip->multi_coll += GETSTAT(sp, STATS_TX_MULTCOL_OFFSET);
1056     ip->collisions += GETSTAT(sp, STATS_TX_TOTCOL_OFFSET);
1057     ip->fcs_errs += GETSTAT(sp, STATS_RX_FCS_OFFSET);
1058     ip->align_errs += GETSTAT(sp, STATS_RX_ALIGN_OFFSET);
1059     ip->norcvbuf += GETSTAT(sp, STATS_RX_NOBUF_OFFSET);
1060     ip->oflo += GETSTAT(sp, STATS_RX_OFLO_OFFSET);
1061     ip->runt += GETSTAT(sp, STATS_RX_SHORT_OFFSET);

1063     ip->stats_time = tstamp;
1064 }
    unchanged_portion_omitted

1156 mblk_t *
1157 iprb_rx(iprb_t *ip)
1158 {
1159     iprb_dma_t    *rfd;
1160     uint16_t      cnt;
1161     uint16_t      sts;
1162     int           i;
1163     mblk_t        *mplist;
1164     mblk_t        *mpp;
1165     mblk_t        *mp;

1167     mplist = NULL;
1168     mpp = &mplist;

1170     for (i = 0; i < NUM_RX; i++) {
1171         rfd = &ip->rx_b[ip->rx_index];
1172         SYNCRFD(rfd, RFD_STS_OFFSET, 2, DDI_DMA_SYNC_FORKERNEL);
1173         if ((GETRFD16(rfd, RFD_STS_OFFSET) & RFD_STS_C) == 0) {
1174             break;

```

```

1175     }

1177     ip->rx_wdog = gethrtime();
1178     ip->rx_wdog = ddi_get_time();

1179     SYNCRFD(rfd, 0, 0, DDI_DMA_SYNC_FORKERNEL);
1180     cnt = GETRFD16(rfd, RFD_CNT_OFFSET);
1181     cnt &= ~(RFD_CNT_EOF | RFD_CNT_F);
1182     sts = GETRFD16(rfd, RFD_STS_OFFSET);

1184     if (cnt > (ETHERMAX + VLAN_TAGSZ)) {
1185         ip->toolong++;
1186         iprb_rx_add(ip);
1187         continue;
1188     }
1189     if (((sts & RFD_STS_OK) == 0) && (sts & RFD_STS_ERRS)) {
1190         iprb_rx_add(ip);
1191         continue;
1192     }
1193     if ((mp = allocb(cnt, BPRI_MED)) == NULL) {
1194         ip->norcvbuf++;
1195         iprb_rx_add(ip);
1196         continue;
1197     }
1198     bcopy(rfd->vaddr + RFD_PKT_OFFSET, mp->b_wptr, cnt);

1200     /* return it to the RFD list */
1201     iprb_rx_add(ip);

1203     mp->b_wptr += cnt;
1204     ip->ipackets++;
1205     ip->rbytes += cnt;
1206     if (mp->b_rptr[0] & 0x1) {
1207         if (bcmp(mp->b_rptr, &iprb_bcast, 6) != 0) {
1208             ip->multircv++;
1209         } else {
1210             ip->brdcstrcv++;
1211         }
1212     }
1213     *mpp = mp;
1214     mpp = &mp->b_next;
1215 }
1216     return (mplist);
1217 }
    unchanged_portion_omitted

1647 void
1648 iprb_periodic(void *arg)
1649 {
1650     iprb_t *ip = arg;
1651     boolean_t reset = B_FALSE;

1653     mutex_enter(&ip->rulock);
1654     if (ip->suspended || !ip->running) {
1655         mutex_exit(&ip->rulock);
1656         return;
1657     }

1659     /*
1660     * If we haven't received a packet in a while, and if the link
1661     * is up, then it might be a hung chip. This problem
1662     * reportedly only occurs at 10 Mbps.
1663     */
1664     if (ip->rxhangbug &&
1665         ((ip->miih == NULL) || (mii_get_speed(ip->miih) == 10000000)) &&
1666         ((gethrtime() - ip->rx_wdog) > ip->rx_timeout)) {

```

```
1666     ((ddi_get_time() - ip->rx_wdog) > ip->rx_timeout)) {
1667         cmn_err(CE_CONT, "?Possible RU hang, resetting.\n");
1668         reset = B_TRUE;
1669     }
1671     /* update the statistics */
1672     mutex_enter(&ip->culock);
1674     if (ip->tx_wdog && ((gethrtime() - ip->tx_wdog) > ip->tx_timeout)) {
1674     if (ip->tx_wdog && ((ddi_get_time() - ip->tx_wdog) > ip->tx_timeout)) {
1675         /* transmit/CU hang? */
1676         cmn_err(CE_CONT, "?CU stalled, resetting.\n");
1677         reset = B_TRUE;
1678     }
1680     if (reset) {
1681         /* We want to reconfigure */
1682         iprb_stop(ip);
1683         if (iprb_start(ip) != DDI_SUCCESS) {
1684             iprb_error(ip, "unable to restart chip");
1685         }
1686     }
1688     iprb_update_stats(ip);
1690     mutex_exit(&ip->culock);
1691     mutex_exit(&ip->rulock);
1692 }
unchanged_portion_omitted
```

```

*****
9904 Mon May 5 11:11:06 2014
new/usr/src/uts/common/io/iprb/iprb.h
4778 iprb shouldn't abuse ddi_get_time(9f)
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Garrett D'Amore <garrett@damore.org>
*****
1 /*
2  * This file and its contents are supplied under the terms of the
3  * Common Development and Distribution License ("CDDL"), version 1.0.
4  * You may only use this file in accordance with the terms of version
5  * 1.0 of the CDDL.
6  *
7  * A full copy of the text of the CDDL should have accompanied this
8  * source. A copy of the CDDL is also available via the Internet at
9  * http://www.illumos.org/license/CDDL.
10 */

12 /*
13  * Copyright 2014 Nexenta Systems, Inc. All rights reserved.
14  * Copyright 2010 Nexenta Systems, Inc. All rights reserved.
15  */

16 #ifndef _IPRB_H
17 #define _IPRB_H

19 /*
20  * iprb - Intel Pro/100B Ethernet Driver
21  */

23 /*
24  * Tunables.
25  */
26 #define NUM_TX          128      /* outstanding tx queue */
27 #define NUM_RX          128      /* outstanding rx queue */

29 /* timeouts for the rx and tx watchdogs (nsec) */
30 #define RX_WATCHDOG     (15 * NANOSEC)
31 #define TX_WATCHDOG     (15 * NANOSEC)
29 #define RX_WATCHDOG     15      /* timeout for rx watchdog (sec) */
30 #define TX_WATCHDOG     15      /* timeout for tx watchdog (sec) */

33 /*
34  * Driver structures.
35  */
36 typedef struct {
37     ddi_acc_handle_t    acch;
38     ddi_dma_handle_t    dmah;
39     caddr_t             vaddr;
40     uint32_t            paddr;
41 } iprb_dma_t;
   unchanged portion omitted

48 typedef struct iprb {
49     dev_info_t          *dip;
50     ddi_acc_handle_t    pcih;
51     ddi_acc_handle_t    regsh;
52     caddr_t             regs;

54     uint16_t            devid;
55     uint8_t             revid;

57     mac_handle_t        mach;
58     mii_handle_t        miih;

60     ddi_intr_handle_t    intrh;

```

```

62     ddi_periodic_t      perh;

64     kmutex_t            culock;
65     kmutex_t            rulock;

67     uint8_t             factaddr[6];
68     uint8_t             curraddr[6];

70     int                 nmcast;
71     list_t              mcast;
72     boolean_t           promisc;
73     iprb_dma_t          cmds[NUM_TX];
74     iprb_dma_t          rxb[NUM_RX];
75     iprb_dma_t          stats;
76     hrtime_t            stats_time;
77     time_t              stats_time;

78     uint16_t            cmd_head;
79     uint16_t            cmd_last;
80     uint16_t            cmd_tail;
81     uint16_t            cmd_count;

83     uint16_t            rx_index;
84     uint16_t            rx_last;
85     hrtime_t            rx_wdog;
86     hrtime_t            rx_timeout;
87     hrtime_t            tx_wdog;
88     hrtime_t            tx_timeout;
84     time_t              rx_wdog;
85     time_t              rx_timeout;
86     time_t              tx_wdog;
87     time_t              tx_timeout;

90     uint16_t            eeprom_bits;

92     boolean_t           running;
93     boolean_t           suspended;
94     boolean_t           wantw;
95     boolean_t           rxhangbug;
96     boolean_t           resumebug;
97     boolean_t           is557;
98     boolean_t           canpause;
99     boolean_t           canmwi;

101     /*
102     * Statistics
103     */
104     uint64_t            ipackets;
105     uint64_t            rbytes;
106     uint64_t            multircv;
107     uint64_t            brdcstrcv;
108     uint64_t            opackets;
109     uint64_t            obytes;
110     uint64_t            multixmt;
111     uint64_t            brdcstxmt;
112     uint64_t            ex_coll;
113     uint64_t            late_coll;
114     uint64_t            uflo;
115     uint64_t            defer_xmt;
116     uint64_t            one_coll;
117     uint64_t            multi_coll;
118     uint64_t            collisions;
119     uint64_t            fcs_errs;
120     uint64_t            align_errs;
121     uint64_t            norcvbuf;

```

new/usr/src/uts/common/io/iprb/iprb.h

3

```
122         uint64_t         oflo;
123         uint64_t         runt;
124         uint64_t         nocarrier;
125         uint64_t         toolong;
126         uint64_t         macxmt_errs;
127         uint64_t         macrcv_errs;
128 } iprb_t;
_____unchanged_portion_omitted_____
```