

fast ethernet driver MII phy serial clock rates

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2004-04/0210.html>

From: David Burns (*david.burns_at_dugeem.net*)

Date: 04/24/04

Date: Sat, 24 Apr 2004 17:23:59 +1000

To: net@freebsd.org

Hello all,

It appears that quite a few of the "el cheapo" hardware Fast Ethernet drivers (at least rl, sis, ste, vr, wb – these are just the ones I found in /usr/src/sys/pci) have added DELAY(1) statements around MII serial clock ops which will result in a max Management Data Clock (MDC) frequency of 500kHz for the serial management interface. Which means that a mii_readreg (or writereg) operation will take a minimum of 128?s (64?s for mii_sync + 64?s for data read/write). During which time the driver is locked.

NB this assumes that a DELAY(1) is really a delay of 1?s! Which I don't think it is ... :-(

However many Fast Ethernet (ie 100Mb/s) PHYs appear to specify a maximum MDC rate of 2.5MHz.

Whilst at first this appears harmless – the mii_readreg & mii_writereg routines are periodically called by MII bus functions every second:

- With autoneg on there are around 7 mii register ops (0.9ms total)
- With autoneg off there are around 3 mii register ops (0.4ms total)

The serial management access bits are set/cleared via various macros (eg. CLRBIT/SETBIT). Generally a clock bit operation consists of a CSR_READ & CSR_WRITE which are of course PCI read & write operations with minimum clock times of 4 cycles and 3 cycles respectively – or 210 nanoseconds per half cycle (@33MHz) which is a bit slower than 2.5MHz! Of course this assumes PCI 33MHz – which is all this hardware will work with.

So I'd like to propose that these DELAY() statements be removed if testing results are okay. I believe this has already been done with the xl driver some time ago...

For verification I made this change on the ste v1.58 driver and it worked fine – and has resulted in 5–10% network performance improvements. Next up I will test the vr driver.

freebsd-net: fast ethernet driver MII phy serial clock rates

If needs be I can open a PR for this but wanted some feedback first from others who may have previously worked on the driver MII code.

David

freebsd-net@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-net>

To unsubscribe, send any mail to "freebsd-net-unsubscribe@freebsd.org"