

Re: [PATCH] TX algorithms, missetting IFF_OACTIVE and if_timer

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2004-04/0011.html>

From: Ruslan Ermilov (*ru_at_FreeBSD.ORG*)

Date: 04/03/04

Date: Sat, 3 Apr 2004 02:49:52 +0300

To: Bill Paul <wpaul@FreeBSD.ORG>

On Fri, Apr 02, 2004 at 09:03:02AM -0800, Bill Paul wrote:

[...]

> > *To differentiate the case of an empty
> > ring from the full ring, some drivers (ste(4), dc(4), and
> > nge(4)) have the threshold (6 for dc(4), 3 for ste(4), and
> > 2 for nge(4)) to assert the gap between producer and consumer,
> > thus not allow the producer to catch the consumer. (The
> > vr(4) is hairier, and I will not discuss it in detail here.)*

> >

> > *First, could you please explain these magic numbers?*

>

> *Not really, no. Very often, values were chosen because they worked
> (and in some cases, they weren't chosen by me).*

>

Hmm, well, at least I now know (learned the hard way) why the gap is ever necessary — I will just silently join the crew who keep this secret, and don't tell it to anyone. ;)

> > *Also, some drivers use indexes for consumer and producer,
> > where they could use "next" pointers, which should be faster.*

>

> *"Should" be faster? I'm not saying you're wrong, but can you prove
> that it's faster to use lists? I started out using linked lists
> for descriptors, but then I started to encounter chips that used
> producer/consumer indexes internally (like the Adaptec 'starfire'
> chip and the Tigon II). I decided that since I tended to allocate
> all of the descriptors in contiguous chunks anyway, it was simpler
> to just treat them as arrays and use index counters.*

>

I experimented with ste(4) today — except for getting 200 bytes less driver code when converting to use the precomputed pointers, I didn't notice any change in performance, so I threw my changes away. ;)

freebsd-net: Re: [PATCH] TX algorithms, missetting IFF_OACTIVE and if_timer

> > *I also think that using the gap between producer/consumer is
> > suboptimal, but this gap is part of the existing algorithm.*

>

> *Nowhere is it written that you can't change the algorithm. :)*

>

Now I know (I wish you'd tell me it) why the gap is necessary,
but let's keep this secret. ;)

> *Note that if you're looking for approval from me to check in these
> patches, don't bother: I will neither approve nor disapprove. The
> only way for me to know if your changes are correct is to test them,
> and I don't have the time or resources right now to do that. If you
> want to commit them, go ahead. It's your funeral. :)*

>

Understood. This is some ancient code, and you have lot of modern
stuff to play with. ;)

Actually, I was just looking for your advise and your vision.

[...]

> *And then the stork comes, and it's a driver.*

>

LOL

Cheers,

--

Ruslan Ermilov
ru@FreeBSD.org
FreeBSD committer

-
- application/pgp-signature attachment: [stored](#)