

## Re: spinaphore conceptual draft

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-05/6524.html>

---

**From:** Andi Kleen ([ak\\_at\\_muc.de](mailto:ak_at_muc.de))

**Date:** 05/30/05

Date: 30 May 2005 21:28:26 +0200  
Date: Mon, 30 May 2005 21:28:26 +0200  
To: Kyle Moffett <[mrmacman\\_g4@mac.com](mailto:mrmacman_g4@mac.com)>

On Mon, May 30, 2005 at 02:04:36PM -0400, Kyle Moffett wrote:

> >I suspect any attempt to use time stamps in locks is a bad  
> >idea because of this.  
>  
> Something like this could be built only for CPUs that do support that  
> kind of cycle counter.

That gets you into a problem with binary distribution kernels.  
While binary patching works to some extent, it also becomes  
messy pretty quickly.

> >My impression is that the aggressive bus access avoidance the  
> >original poster was trying to implement is not that useful  
> >on modern systems anyways which have fast busses. Also  
> >it is not even clear it even saves anything; after all the  
> >CPU will always snoop cache accesses for all cache lines  
> >and polling for the EXCLUSIVE transition of the local cache line  
> >is probably either free or very cheap.  
>  
> The idea behind these locks is for bigger systems (8-way or more) for  
> heavily contended locks (say 32 threads doing write() on the same fd).

Didn't Dipankar & co just fix that with their latest RCU patchkit?  
(assuming you mean the FD locks)

> In such a system, cacheline snooping isn't practical at the hardware  
> level, and a lock such as this should be able to send several CPUs to

Why not? Cache snooping has to always work with low overhead, otherwise the  
machine is not very useful coherent. I assume that any bigger system  
has a cache directory anyways, which should minimize the traffic;  
and for smaller setups listening to broadcasts works fine.

-Andi

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in

Linux-Kernel: Re: spinaphore conceptual draft

the body of a message to [majordomo@vger.kernel.org](mailto:majordomo@vger.kernel.org)

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

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