

Re: [PATCH 6/9] clockpro-clockpro.patch

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-12/msg08554.html>

- *From:* Marcelo Tosatti <marcelo.tosatti@xxxxxxxxxxxxx>
 - *Date:* Sat, 31 Dec 2005 01:27:02 -0200
-

Hi Rik!

On Fri, Dec 30, 2005 at 08:22:12PM -0500, Rik van Riel wrote:

> On Fri, 30 Dec 2005, Marcelo Tosatti wrote:

>

>> I think that final objective should be to abstract it away completely,
>> making it possible to select between different policies, allowing
>> further experimentation and implementations such as energy efficient
>> algorithms.

>

> I'm not convinced. That might just make vmscan.c harder to read ;)

Are you serious or just joking? :)

Sure it might make it harder to read, but allowing selectable policies is very interesting. Peter's patches go half-way into that direction.

Lets say, if CLOCK-Pro underperforms for a given workload (take into account that its simply optimizing reclaim for a subset of all existing access patterns, ie. heuristics), it would be easier for people to develop/use different policies.

>> About CLOCK-Pro itself, I think that a small document with a short
>> introduction would be very useful...

>

> <http://linux-mm.org/AdvancedPageReplacement>

I meant something more like Documentation/vm/clockpro.txt, for easier reading of patch reviewers and community in general.

>>> The HandCold rotation is driven by page reclaim needs. HandCold in turn
>>> drives HandHot, for every page HandCold promotes to hot HandHot needs to
>>> degrade one hot page to cold.

>>

>> Why do you use only two clock hands and not three (HandHot, HandCold and
>> HandTest) as in the original paper?

>

> Because the non-resident pages cannot be in the clock.
> This is both because of space overhead, and because the

Re: [PATCH 6/9] clockpro-clockpro.patch

> non-resident list cannot be per zone.

I see – that is a fundamental change from the original CLOCK-Pro algorithm, right?

Do you have a clear idea about the consequences of not having non-resident pages in the clock?

> I agree though, Peter's patch could use a lot more
> documentation.

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxx

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

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

- ***Follow-Ups:***

- ◆ ***Re: [PATCH 6/9] clockpro-clockpro.patch***
 ◇ *From: Rik van Riel*

- ***References:***

- ◆ ***[PATCH] vm: page-replace and clockpro***
 ◇ *From: Peter Zijlstra*
- ◆ ***[PATCH 6/9] clockpro-clockpro.patch***
 ◇ *From: Peter Zijlstra*
- ◆ ***Re: [PATCH 6/9] clockpro-clockpro.patch***
 ◇ *From: Marcelo Tosatti*
- ◆ ***Re: [PATCH 6/9] clockpro-clockpro.patch***
 ◇ *From: Rik van Riel*

- Prev by Date: ***Re: recommended mail clients [was] [PATCH] ati-agp suspend/resume support (try 2)***
- Next by Date: ***Re: [patch 00/2] improve .text size on gcc 4.0 and newer compilers***
- Previous by thread: ***Re: [PATCH 6/9] clockpro-clockpro.patch***
- Next by thread: ***Re: [PATCH 6/9] clockpro-clockpro.patch***
- Index(es):
 - ◆ ***Date***
 - ◆ ***Thread***