- **Regan Robertson** Mod • 14 days ago

  We hope everyone is enjoying the second day of the summit. We are going into the lunch break and will be back at 1pm for the next session.

  - •
  - Reply
  - •
  - Share ›

- **Todd Carpenter** • 14 days ago

  Hello, everyone! Thanks for attending my talk today.

  This presentation discusses using seL4 capability to provide the static, cyclic scheduling that Gernot mentioned yesterday. He's been encouraging me to switch to MCS for at least a couple years. I thought the domain scheduler was still worth discussing since it is so ridiculously easy to use and is rock solid. It has enough similarities to ARINC653 that it's great for rapid prototyping. We're building control systems, not web browsers. Once MCS is solid, we'll switch.

  So, to questions, observations, etc. Let the show begin!

  - 1
  - •
  - Reply
  - •
  - Share ›

    - **Jason H Li** Todd Carpenter • 14 days ago

      Looking forward to it!

      - 1
      - •
      - Reply
      - •

- Share ›

Jerry Dussault • 14 days ago

My audio volume controls are at max, and still difficult to hear. Any way to increase the gain on the audio?

- 1
- •
- Reply
- •
- Share ›

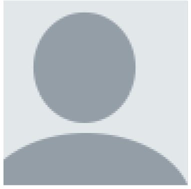

Todd Carpenter Jerry Dussault • 14 days ago

Drat. If you're on Ubuntu, go to settings, and enable over-amplification?

- 1
- •
- Reply
- •
- Share ›

Jerry Dussault Todd Carpenter • 14 days ago

That helps A LOT! Much better. Thanks Todd!

- 1
- •

- Reply
- Share ›

Todd Carpenter Jerry Dussault • 14 days ago

Yay!

- Reply
- Share ›

Jacob Saina Jerry Dussault • 14 days ago

Same; all of the streams have been very low volume for me.

- Reply
- Share ›

Jason H Li Jerry Dussault • 14 days ago

LOL - your first question has to be about IT support!

- 1

- Reply
- •
- Share ›

Todd Carpenter Jason H Li • 14 days ago

This is how it works.

But it's relevant - it's part of the control problem.

- •
- Reply
- •
- Share ›

Jerry Dussault Jason H Li • 14 days ago

Everything else so far has been crystal clear ;)

- •
- Reply
- •
- Share ›

Todd Carpenter Jerry Dussault • 14 days ago

Ouch!

- Reply
- Share ›

Aleksey Nogin • 14 days ago

I am wondering - how is this work related to the Trusted Build and related timing analysis that Mike Whalen (UMN) and others on the (then-) Rockwell Collins team developed under HACMS a few years back? Is this something new/separate, or a continuation of the same work?

- Reply
- Share ›

Ihor Kuz Aleksey Nogin • 14 days ago

Evolution. HAMR is the successor to Trusted Build from HACMS.

- 2
- Reply
- Share ›

Todd Carpenter Ihor Kuz • 14 days ago

What Ihor said. :)

- Reply
- Share ›

Todd Carpenter Aleksey Nogin • 14 days ago

Hi Aleksey! Great question. One of the funding vehicles is DARPA CASE, which is a follow-on to HACMS. Dr. Whalen's work on HACMS was provided a conceptual proof-of-concept for some of this. Many enhancements are now available over that earlier prototype work. I strongly encourage you to attend KSU's presentation tomorrow - you'll hear much more about to today's approach to the earlier "trusted build." It's way cool.

- 1
- Reply
- Share ›

Darren Cofer Aleksey Nogin • 14 days ago

It supersedes the HACMS Trusted Build work. KSU and Adventium are part of our CASE team, and developing this new and improved HAMR code generation approach. John Hatcliff will talk about this tomorrow morning.

- 2
- Reply
- Share ›

Todd Carpenter • 14 days ago

Thank you for attending the talk! I plan to monitor questions here, or feel free to shoot me additional questions via email.

- •
- Reply
- •
- Share ›

Jerry Dussault Todd Carpenter • 14 days ago

Thanks Todd. Since I'm not very familiar with timing/scheduling with seL4, I found this presentation to be very helpful.

- 1
- •
- Reply
- •
- Share ›

Todd Carpenter Jerry Dussault • 14 days ago

Thanks, Jerry!

- •
- Reply
- •
- Share ›

Ihor Kuz • 14 days ago

Have you looked at doing similar with MCS?

- 1
- •
- Reply
- •
- Share ›

Todd Carpenter Ihor Kuz • 14 days ago

Hi Ihor! Thanks for the question. MCS has all this capability and more. So yes, it's possible. One of the funding programs currently has a constraint to stick with the verified kernel, so we don't want to upset things.

MCS has a ton of features, so I'm excited to get into it.

The first thing I plan to evaluate, once we get the time, is how easy it is to use MCS to provide the minimal temporal partitioning shown here. The domain scheduler is truly simple: one C file with the schedule, and in your top level camkes file you specify what is in which domain. Couldn't be much easier to use.

- •
- Reply
- •
- Share ›

Todd Carpenter Todd Carpenter • 14 days ago

So stay tuned. A friend told me that the MCS verification will be done any day now. :)

One of the reasons I'd like the verification complete is that will likely keep the subsequent interface changes to a minimum. That will make it easier to build tools to target it.

- Reply
- Share ›

Ihor Kuz Todd Carpenter • 14 days ago

I was specifically curious whether you've already tried to implement such a system using MCS, Or not yet.

- Reply
- Share ›

Todd Carpenter Ihor Kuz • 14 days ago

Not yet, although the temptation is great. Might have to squeeze in an experiment or two sooner than later.

- Reply
- Share ›

Jason H Li • 14 days ago

Todd - with this approach, do you have some time guarantees calculated already or still in progress for mission critical systems?

- Reply
- Share ›

Todd Carpenter Jason H Li • 14 days ago

Hi Jason! Great question. This approach allows you to pre-calculate certain temporal properties, such as end-to-end latency, jitter, etc., as mentioned in the slides. Several of those properties can be extracted almost directly from the schedule. Note, however, that this relies on other prior analysis, such as Worst Case Execution Time. So these temporal properties are at the domain level.

- Reply
- Share ›

Eric Smith • 14 days ago

Todd, thank you for a clear and informative talk! Folks, please put any additional questions for Todd here in the chat.

- 1
- •
- Reply
- •
- Share ›

Gernot • 14 days ago

Thanks for referring to MCS. To clarify, while MCS verification will be done soon, this will be the functional correctness proof. Binary correctness will be verified with the existing toolchain, so will be there as well. However, the security proofs, especially confidentiality, won't be there for a while

- 2
- •
- Reply
- •
- Share ›

Todd Carpenter Gernot • 14 days ago

Thanks for the clarification, Gernot! That distinction is important to consider when building systems.

- •
- Reply
- •
- Share ›

Arslan Khan • 14 days ago

Q: How is asynchronous interrupt handling (hardware not sel4 events) catered in the temporal isolation model?

- •
- Reply
- •
- Share ›

Gernot Arslan Khan • 14 days ago

The confidentiality proofs assume no interrupts, which is one of the things I don't like about the present domain scheduler

- 1
- •
- Reply
- •
- Share ›

Gernot Gernot • 14 days ago

which, btw, is in line with all other separation kernels AFAIK. Our time protection work will be able provide confidentiality even with interrupts enabled, but that's still research...

- •
- Reply
- •
- Share ›

Arslan Khan Gernot • 14 days ago

Thanks... makes sense... i look forward to that work..

- •
- Reply
- •
- Share ›

Arslan Khan Arslan Khan • 14 days ago

i know there are some systems that mask them and schedule them at a later time slot..which i guess would a domain here..but IMO that results in huge interrupt latencies..

- 1
- •
- Reply
- •
- Share ›

Todd Carpenter Arslan Khan • 14 days ago

It's all in what the requirements are for the particular system you're trying to develop. "fast" is nice, but "requirements" are necessary. The "huge" you mention might be irrelevant if your control loop is stable or your response time is sufficient.

One issue we did run up against is tiny hardware buffers. That forced us to statically schedule a high-rate thread to service the interrupts. That's a fairly common approach, and again, you can prove responsiveness properties. The down-side is, of course, relative inefficiency on the processor. Since most of these systems are in the field for decades, and they MUST work, this is an acceptable tradeoff -for these particular systems-. Not all systems, of course.

- •
- Reply
- •
- Share ›

—

Regan Robertson Mod • 14 days ago

Please join us for the next session that starts in 2 minutes. You can either go back to agenda to get to the next session or at the bottom of the page there is a next session button.