Safe use of Multi-Core and Manycore Processors

This 100th Safety-Critical Systems Club (SCSC) seminar was held both face to face and online on 11th Nov 2021. The topics centred on approaches for using multi and many core processors (MCP) in safety-related and safety-critical applications. This is a new field for industry, and there are many challenges to produce a suitable assurance argument.

Mike Parsons introduced the seminar speakers and opened by noting that the use of such processors has been around for a while. However, the concept of using (and proving!) them in safety-critical applications is new.

Multi and Manycore Safety Working Group (MCWG)

Lee Jacques from Leonardo (and co-chair of the working group) gave an overview of the group covering past, present and future plans.

The past

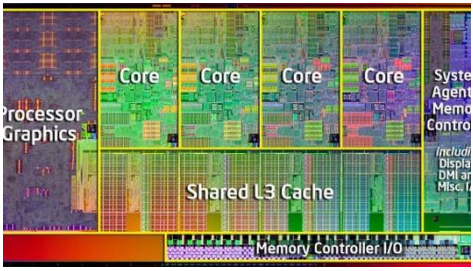

Lee explained the (relatively young) history of the group and that it was created to discuss the challenges around multicore certification and the creation of CAST-32A – a multicore position paper by the Certification Authorities Software Team (CAST). The group has made some good progress in creating a common ontological model, and a number of sub groups have shared knowledge and information thus creating some strong networks.

Present

Whilst initially making good progress Lee highlighted that, the group has started to struggle, as the scope of the sub groups was too wide and introduced significant overlap. The ontology group however, was able to maintain a good pace as it has a defined and bounded scope of activities, which provided the team with a clear focus. Other challenges related to resource availability and the sharing of specific Intellectual Property (IP) were noted.

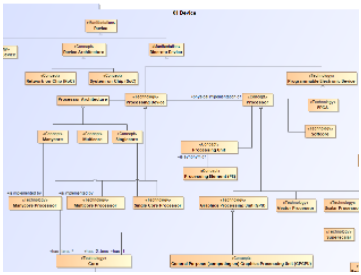

Future

Lee explained that the group was intending to refresh its approach and take a leaf from the ontology group be defining a clear set of well-defined tasks. The current approach was trying to cover a set of objectives that were too broad and a defined focus on clear smaller tasks would be more effective.

Lee asked the group for ideas and challenges that can be used to build a backlog of tasks, which the group can tackle in this new approach.

The Safety and Security Considerations for the Use of Multi-Core Processors

Mike Standish from DSTL and Mark Hadley from Atkins, presented how safety and security should be considered when undertaking certification of MCPs.

They noted a number of challenges including an inability to make an "in service" case, as this technology is not generally in use within a safety-critical component yet. Also, the complexity of these devices means that it is critical that the user understands the device, its architectural performance, boundary, longevity and specific configuration. Without understanding all of these parameters, the ability to make a reasoned assurance case is challenging.

They introduced the "wheel of qualification" (Hadley & Standish 2019) and emphasised the need to make a diverse assurance case relying on various types of evidence and not just (for example) a measure of worst-case execution time (which although perfectly valid, cannot be used alone to justify assurance).

Although a challenge for MCP certification, Safety is relatively well-known in terms of process and assurance. Security on the other hand, introduces additional challenges and does not live in harmony with safety. For example, security patches, whilst necessary, can invalidate (or at least cause a re-evaluation) of an assurance case.

Mike and Mark's closing note was that MCP certification is not just about the MCP, it's about

the whole ecosystem around the solution. This extends from the hardware setup and development environment through to the supply chain.

Incremental Assurance of Multicore Integrated Modular Avionics

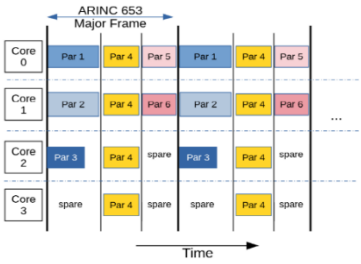

Guillam Bernat from Rapita, discussed the challenges of certifying Integrated Modular Avionics (IMA) with MCP solutions and techniques for performing performance analysis. Using Rapita tools, it is possible to monitor all aspects of MCP performance using embedded RapiDaemons. The challenge with a traditional IMA solution means a significant amount of recertification is required when changing one item.

Guillam stated that careful consideration of the partitioning model and use of automated testing is critical to success. Whilst it is possible to automate data gathering activities, there is a significant amount of manual analysis required to interpret the data.

The presentation highlighted potential test solutions using interference generators to mitigate the challenge of identifying and verifying interference paths in a multicore solution. Guillam continued detailing how this could be used in a mixed criticality context, crucial for keeping time and costs down and easing the certification burden.

Multicore Processors usage in Certified Avionics: How Virtualisation Can Help?

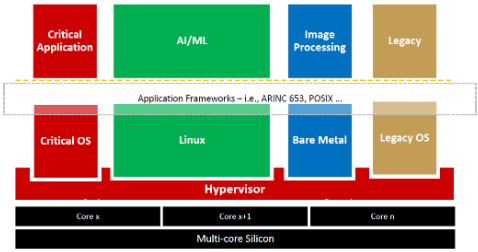

Olivier Charrier from WindRiver, explained how the use of virtual machines can provide the assurance required when partitioning multicore systems. Olivier stressed the importance of considering the architecture and requirements early, and just as importantly as the test strategy. He highlighted that as important as this is, it's also important to consider that assumptions made early on in the process may not fully hold going forward, and that continuous test and evaluation can drive design choices.

The presentation then went on to focus on the potential benefits of using virtualisation to provide a complete partitioning solution (memory, CPU, Cache, etc). Using virtualisation (managed by a Hypervisor) would support an assurance case by providing a well partitioned solution and would address many of the resource usage aspects of CAST-32A. That is of course, assuming you have done the upfront work to determine that virtualisation is an appropriate architecture for your solution!

Telemetry and bare-metal Virtual Machines for Improved Multicore Partitioning

Tim Loveless from Lynx Software Systems gave an overview of Hypervisor technology, which was widely regarded in the seminar as one of the clearest definitions people had seen.

Tim started by noting that most people assume that to ensure a well-partitioned system, you need an ARINC653 based Real-Time Operating System (RTOS). Whilst in many cases this is the correct solution, he noted that for some architectures it is possible to run a bare metal hypervisor to improve performance and reduce complexity'. Do you really need that Ethernet stack and file system?

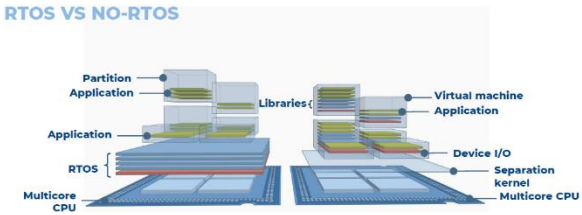

Tim introduced a number of potential patterns, which could be used to deliver a compliant bare metal solution and also introduced the CPU Performance Management Unit (PMU). This is a key component when testing and characterising your CPU as it provides a series of counters measuring everything from number of instructions completed to data misses. He warned though although important for characterisation, they introduce an overhead into the system and the more complex the system, the more integration with the RTOS is required. Think how often, how and when you are going to store or offboard the data whilst trying to maintain multicore, real-time performance…

Multi-core architectures and timing analysis: Their influence on the scheduling of certifiable real-time systems

Iain Bate from the University of York presented some techniques for analysing multicore architectures and building a timing analysis approach to most effectively assess the performance of a system.

As many other presenters noted, it is key that all of this activity needs to be considered up front and that system architecture understanding is key. For example, what CPU resources are being used, what resources are being shared, what partitioning approach should be considered etc. Making design decisions is difficult as the software implementation affects the performance, but at this stage the software has not been written.

Iain presented a 4-step process to introduce some rigour into the process and help guide the identification and mitigation of the key interference factors.

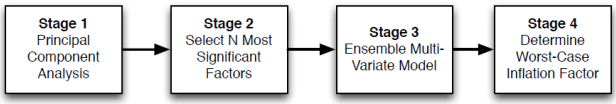

Certification aspects of Multicore

Sam Riley from Frazer Nash (Formerly MAA) gave an insight into the thinking of a certification authority based on first-hand experience and noted that there is no definitive approach and positive engagement with the authority is key.

He concluded by discussing 7 general "lessons" from his experience to help those preparing safety cases, for example, ensuring that evidence is diverse and that key people, such as the design leads and regulator, are taken "along on the journey" from an early stage.

Report by Lee Jacques, MCWG Co-Chair