Center for Embedded Systems

An NSF Industry/University Cooperative Research Center

Towards Predictable Execution of Safety-Critical Tasks on Mixed-Criticality Multi-Core Platforms

PIs:

Harini Ramaprasad

Dimitri Kagaris

Southern Illinois University Carbondale

SIU Southern Illinois University CARBONDALE

NATIONAL SCIENCE FOUNDATION

ASU Ira A. Fulton Schools of Engineering ARIZONA STATE UNIVERSITY

Project Overview and Description

- **Project Description**
  - Conduct trade-off study of determinism vs. performance of mixed-criticality tasks on Freescale P4080 multi-core platform.
  - Develop policies to maintain responsiveness of HSS applications under regular and overload conditions.
- **Problem**
  - Deterministic execution of HSS tasks in presence of LSS ones on multi-core architectures is challenging.

Approach

Two stage approach:

- Stage 1: a) Use cache locking and partitioning to improve predictability of HSS tasks; b) Identify/develop suitable policies to apply to end-user scenarios.
- Stage 2: Explore the use of manager partition to dynamically control resource usage of LSS tasks under overload or unexpected situations.

Novelty:

- No study/research on applying cache locking and partitioning mechanisms to mixed-critical workloads executing in virtualized environments.

Potential member company benefits:

- Integration on Multicore -> save space and cost
- Use of hypervisor -> Safe execution of mixed-criticality workloads.
- Allowance of migration of certified HSS tasks to the multi-core platform.

Things needed from IAB:

Sanitized / anonymized workload characteristics for benchmarks.

Project Status

Year 1:

- Explored partitioning of cores among multiple partitions.
- Studied cache behavior under no shared cache (Corenet Platform Cache or L3 cache) partitioning.
- Explored partitioning of CPC among multiple OS partitions and the Physical Memory Area (PMA1) , which is a shared memory region that all partitions access, presumably for system/hypervisor related data storage.
- Results revealed that the real time performance improves with increased CPC assignment to RT partition.

Year 2 progress:

- Research on Cache locking mechanism for the Freescale P4080 platform.
- Developed a kernel module to enable cache locking from the user space.
- Implementing features in the kernel module to load and lock data in the cache lines from user space.

Project Tasks/ Deliverables

	Description	Date	Status
1	Exploration of existing research in the area of cache locking and partitioning.	Q1	In progress
2	Workload characterization and end-use scenario analysis under cache locking and partitioning schemes.	Q2	In progress
3	Exploration of mechanisms to create and configure manager partitions;	Q3	Not yet started
4	Development of strategies for dynamic resource management using manager partitions	Q3	Not yet started
5	Report writing and technology transfer	Q4	Not yet started

Executive Summary

Problem

- Deterministic execution of HSS tasks in mixed-criticality multi-core environment and develop policies to maintain responsiveness of HSS tasks under regular and overload conditions.

Viable solution

- Virtualization (hypervisor) for isolation of HSS & LSS task sets and applying the cache-locking and cache partitioning techniques.
- Hypervisors allow configuring one partition as a manager partition, giving this partition rights to pause and resume other partitions.

State-of-the-art

- Comprehensive trade-off study of *determinism* vs. *performance* of mixed-criticality tasks executing in the Freescale P4080 multi-core platform.

Goals of proposed project

- Conduct trade-off study of determinism versus performance HSS-LSS & HSS-HSS interactions with cache-locking and cache partitioning techniques.
- Explore the use of a manager partition to dynamically control the resource usage of LSS tasks under overload or unexpected situations in an effort to maintain deterministic execution of HSS tasks.

Technical Detail

- Freescale QorIQ P4080
  - 8 high-performance cores
  - Private L1 & L2, shared L3

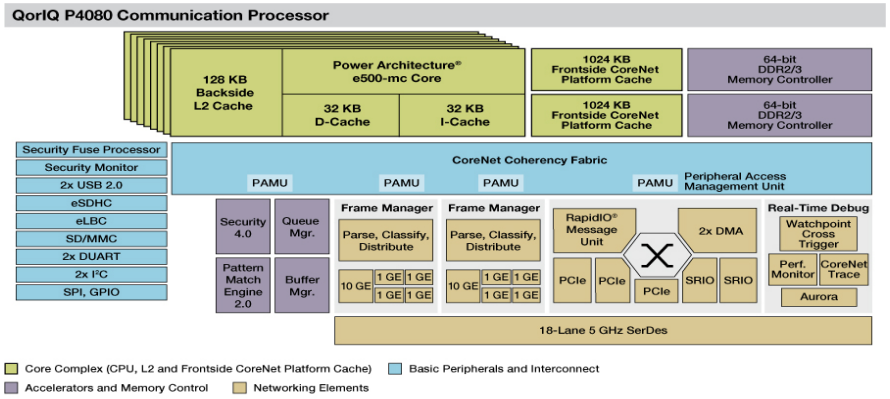

  - Embedded hypervisor
    - •Safe OS partitioning
    - •Takes advantage of hardware mechanisms present in cores
    - •Provides support for partitioning cores, memory, I/O devices
      - –Each OS *only* accesses resources it is authorized to access
      - –Each OS *owns* resources in its partition
    - •Can configure one partition as "*manager partition*"
    - •External interrupts may be directly sent to Oss
  - High-bandwidth communication & coherence infrastructure
    - •Support for prioritization, bandwidth allocation, packet-level queue management and QoS scheduling
  - Software and technical support from Freescale
    - •Complete integrated development environment

Technical Detail

- Suitability of Freescale P4080 platform
  - Scheduling granularity
    - Coarse-grained: only static partitioning of resources allowed
      - *In contrast, Xen allows fine-grained scheduling of multiple OSs on a single core to maximize processor utilization*
    - Suitable for systems with HSS tasks where determinism is paramount
  - Hypervisor design
    - Exploits hardware mechanisms in cores to improve efficiency of virtualization
    - Easier to bound hypervisor interference across OSs
- Benchmarks for creation of mixed-criticality task sets
  - MRTC WCET benchmarks
  - EEMBC benchmarks: LMBench, CoreMark, perf_measure(RCI)

Technical Detail

- Configure one partition as a Manager partition which gives the partition certain rights to manage other partitions.
- Manager Partition has the capability to:
  - Start, stop, restart other partitions (managed partitions) via hypercall APIs.
  - Receive and send doorbell interrupts upon these events:
    - Watchdog expiration
    - Restart request
    - Managed partition state change
      - ✓transition from starting to running
      - ✓transition from stopping to stopped
      - ✓transition from pausing to paused
      - ✓transition from resuming to running

Year 1 results

CPC partitioned between RTOS and GPOS only:

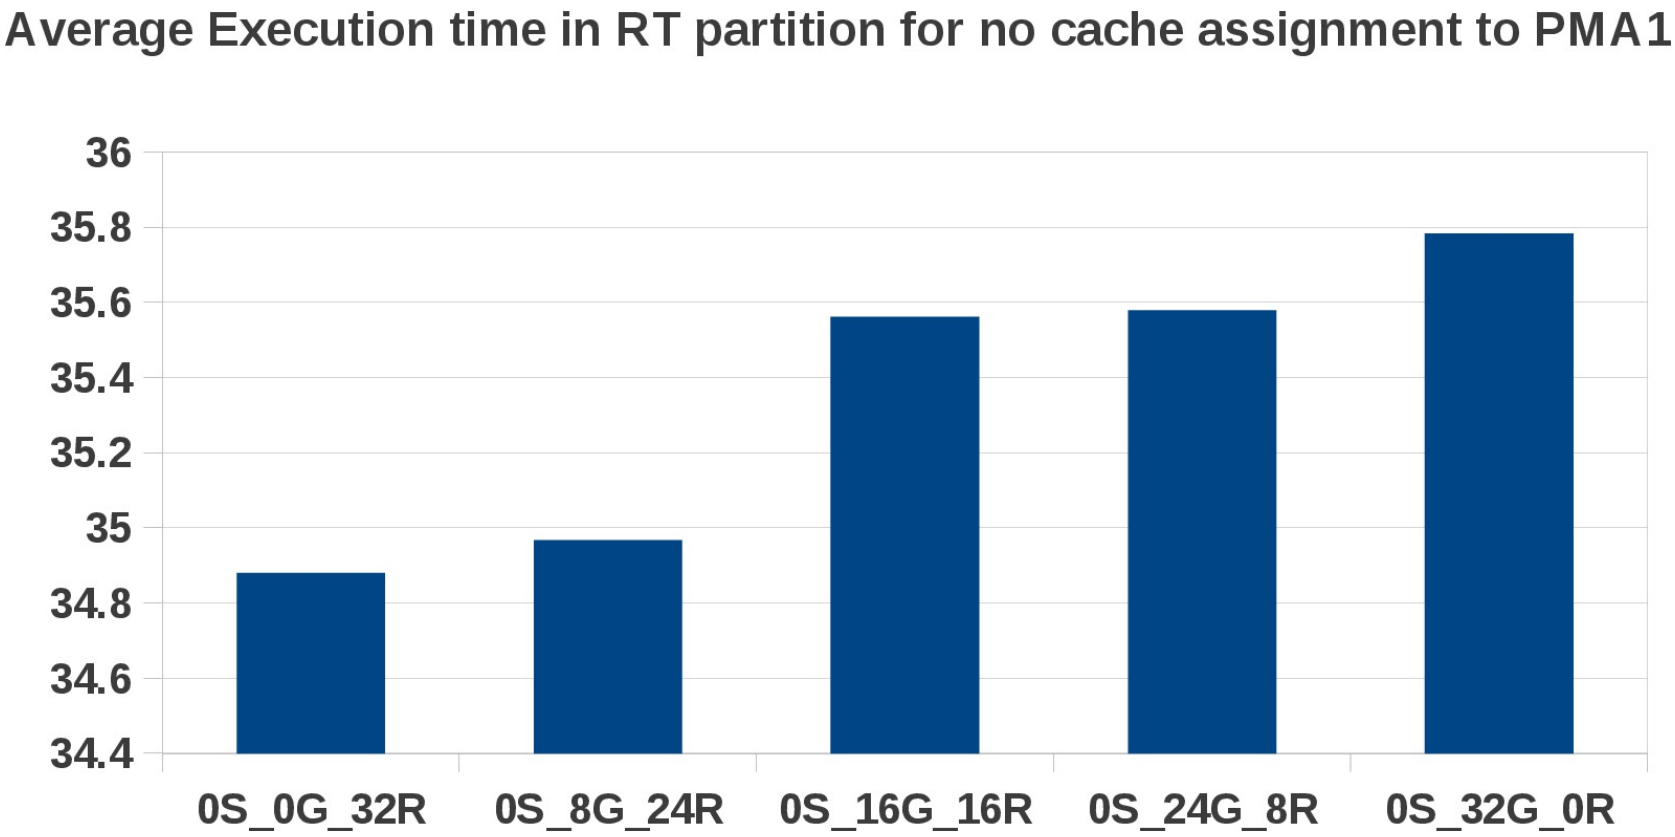

Year 1 results

CPC fully available to RTOS and partially shared with GPOS:

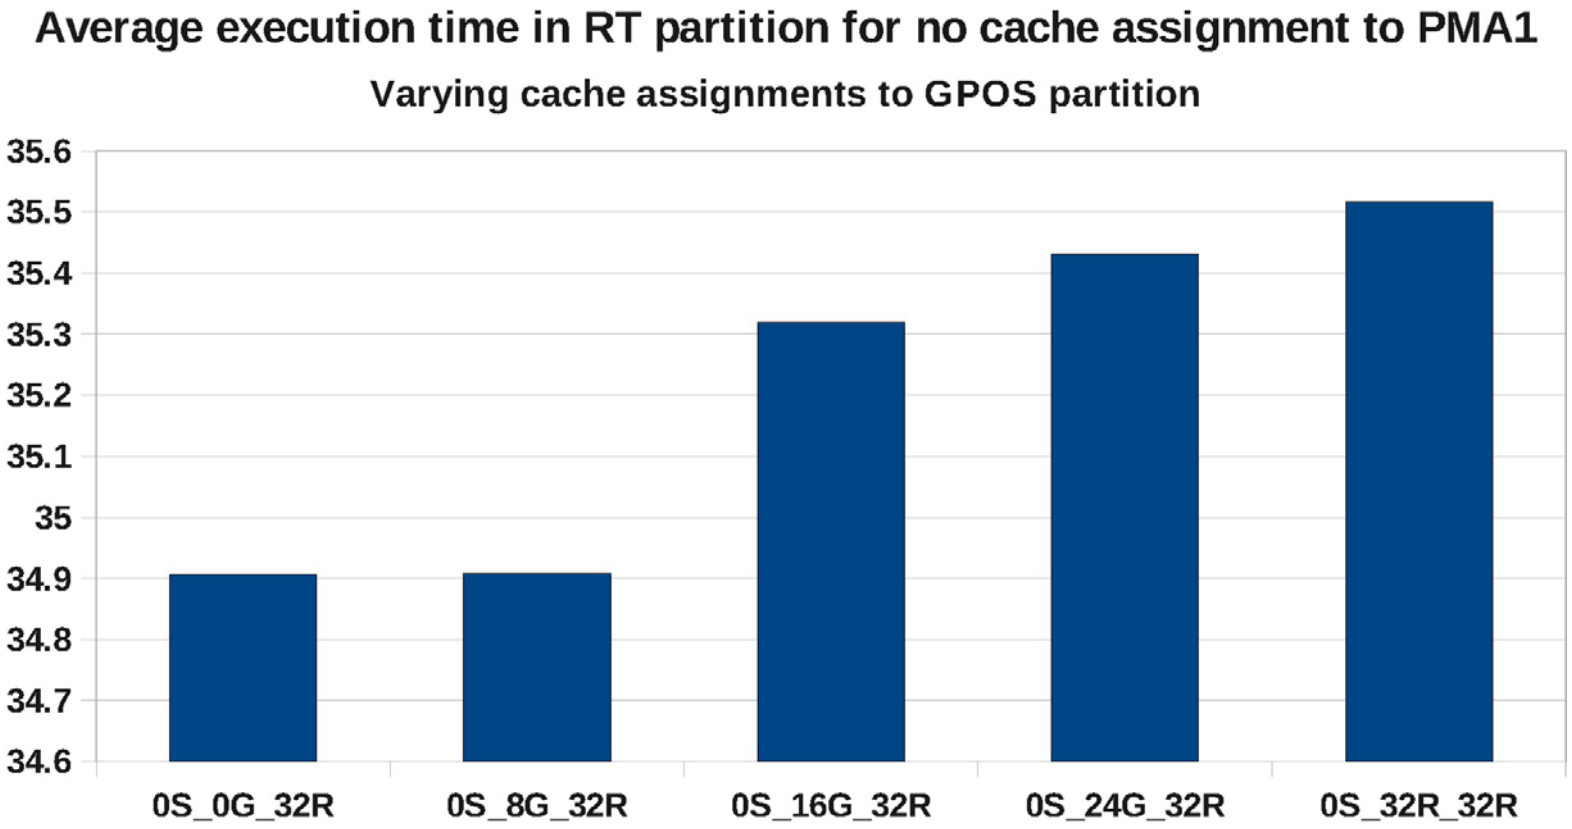

Year 1 results

CPC partitioned between RTOS and PMA1 only:

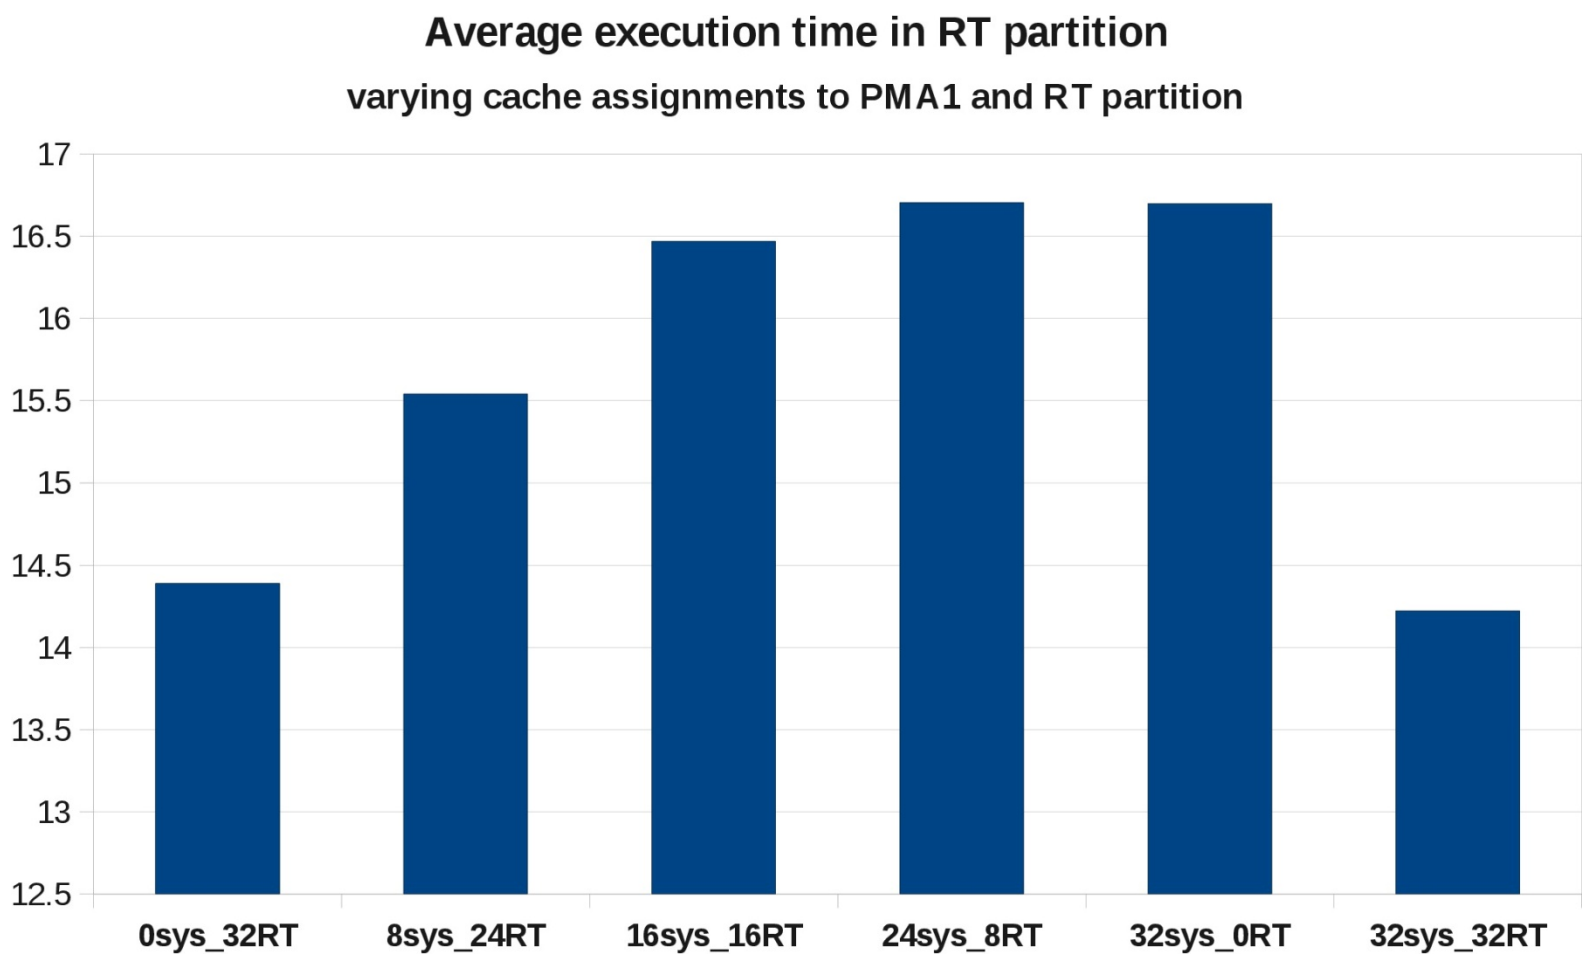

Year 1 results

CPC completely shared between RTOS and PMA1 and partially shared with GPOS:

Average execution time in RT partition for varying cache assignment to GP partition

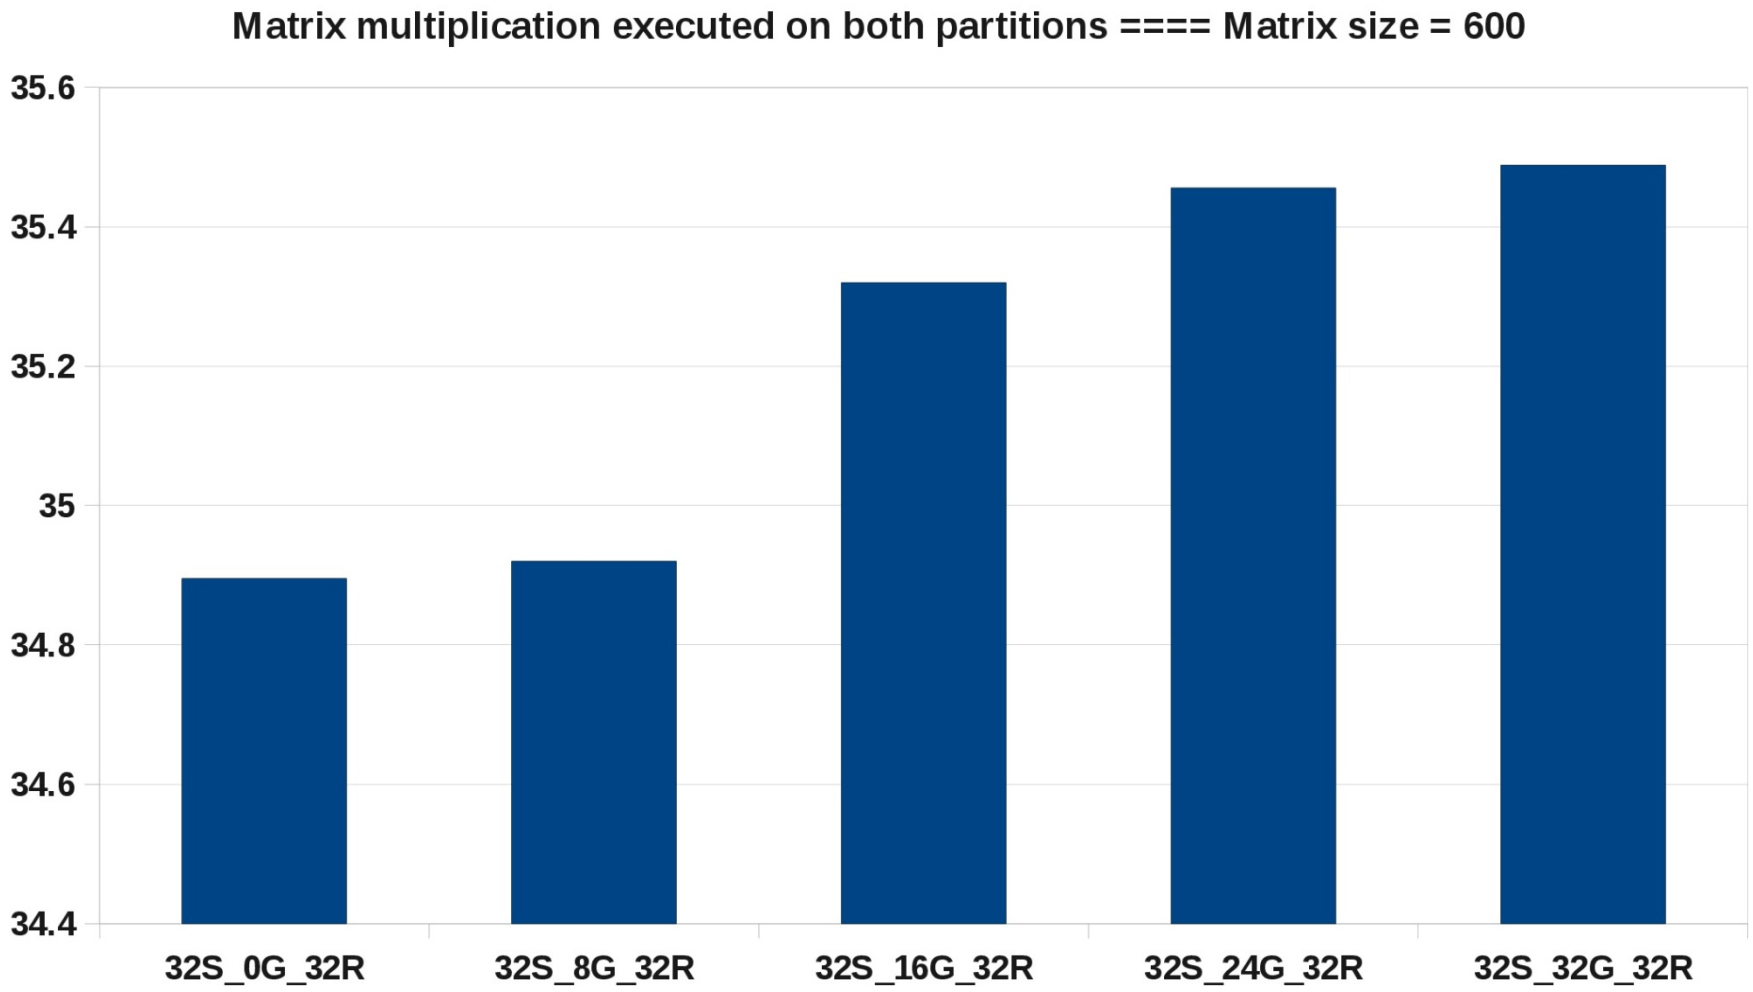

References

- S. Xi, J. Wilson, C. Lu and C. Gill. “RT-XEN: Towards Real-Time Hypervisor Scheduling in Xen.” In Proceedings of the 9th ACM international conference on Embedded Software (EMSOFT), 2011.
- R. Fuchsen. “How to address certification for multi-core based IMA platforms: Current status and potential solutions.” In Proceedings of the Digital Avionics Systems Conference (DASC), 2010.
- C. Ault. “Challenges of safety-critical multi-core systems.” White paper, Wind River Research.
- P. Baltham, et al. “Xen and the art of Virtualization,” In Proceedings of SOSP '03 19th ACM Symposium on Operating systems principles.
- Jun Zhang; Kai Chen; Baojing Zuo; Ruhui Ma; Yaozu Dong; Haibing Guan. “Performance analysis towards a KVM-Based embedded real-time virtualization architecture.” 5th International Conference on Computer Sciences and Convergence Information Technology (ICCIT), 2010 , pp. 421 – 426.
- M. Peshave. “High-Assurance Reconfigurable Multicore Processor Based Systems.” In Proceedings of the 13th IEEE International Symposium on High-Assurance Systems Engineering (HASE), 2011.
- M.S. Mollison, J.P. Erickson, J.H. Anderson, S.K. Baruah and J.A. Scoredos. “Mixed-Criticality Real-Time Scheduling for Multicore Systems.” In Proceedings of the 10th IEEE International Conference on Computer and Information Technology (CIT), 2010.
- Freescale Inc. “P4080: QorIQ P4080 Eight-Core Communications Processors with Data Path”, www.freescale.com/webapp/sps/site/prod_summary.jsp?code=P4080.
- devicesolutions.net “Topaz i.MX25 CPU Module Technical Reference Manual.” http://devicesolutions.net/LinkClick.aspx?fileticket=hS62BQTuyhM%3D&tabid=305
- Adventium Labs. “MiCART – Mixed-Criticality, Real-Time Virtualization Support.” www.adventiumlabs.com/?q=productsandservices/micart.
- “Lmbench – Tools for Performance Analysis.” http://www.bitmover.com/lmbench/.
- “MRTC WCET Benchmarks.” www.mrtc.mdh.se/projects/wcet/benchmarks.html.
- Puaut and D. Decotigny. “Low-complexity algorithms for static cache locking in multitasking hard real-time systems.” In In IEEE Real-Time Systems Symposium, pages 114–123, 2002.