

SPHERE: Software architecture for Predictable HETerogeneous REal-time systems

keywords: Time-Sensitive Networking (TSN), Real-time Networks, Industrial and Automotive networks, Heterogeneous platforms, Real-time systems.

The presentation will overview the research activities and scientific results obtained by the Real-time systems and Networks group of the University of Catania relevant to the SPHERE project "Software architecture for Predictable HETerogeneous REal-time systems", funded by the Italian Ministry of Research through the PRIN 2017 Program.

The advent of commercial-of-the-shelf (COTS) heterogeneous multi-core platforms offers great opportunities for developing high-performance embedded computing systems in many application scenarios. However, safety-critical applications require crucial properties, such as safety, security, and temporal predictability, which cannot be so easily guaranteed on such platforms. At the same time, the increased complexity of emerging systems (e.g., autonomous driving cars, humanoid robots, augmented reality, and virtual interactive environments) requires the integration of different subsystems and functionalities, thus imposing additional constraints to embedded systems designers and highlighting the limits of current development frameworks.

The SPHERE project aims at providing an integrated operating system framework (validated through a pioneering autonomous driving system for the automotive domain) to abstract the hardware complexity of cutting-edge multi-core platforms and simplify the management of heterogeneous computational resources. In particular, the presentation will discuss the desirable properties of software transmission mechanisms to enable real-time communications among multiple distributed control units handled by a hypervisor.

The SPHERE framework needs a communication network able to handle different traffic flows with different requirements, both internal to the heterogeneous platforms and external with other devices. Time-Sensitive Networking (TSN) is an emerging set of standards suitable for this context [1]. In fact, the TSN family standards can be combined to design network architectures able to meet the requirements of clock synchronization [2], fault-tolerance [3], and to enable the support of different traffic classes [4] with different temporal requirements.

Heterogeneous platforms typically have a limited number of Ethernet ports that can be lower than the number of virtual machines. This entails contention on the Ethernet port or the need for exclusive access from one virtual machine. Moreover, as TSN transmissions have to follow specific scheduling rules, a scheduling algorithm has to be applied to schedule transmissions among multiple virtual machines when a hypervisor is adopted. In the presentation, we will discuss three different approaches [5] to handle TSN real-time transmissions (at the application level, virtual machine level and Hypervisor level, respectively) in a heterogeneous platform with virtualization.

The application level transmission offers a flexible solution, thanks to the per-flow scheduling granularity. However, suitable scheduling policies have to be used to limit jitter.

The virtual machine level transmission provides the application with transparent transmission scheduling. However, using this approach the TxScheduler task is scheduled regardless of the priority of the flows to be transmitted, thus losing the per-flow scheduling granularity.

The Hypervisor level transmission offers the same scheduling with per-flow granularity as the application level, but providing better performance. However, such a solution increases the Hypervisor complexity.

A detailed overview of the different scheduling approaches will be provided. Moreover, we will present a performance assessment of such approaches, providing some insights on the limits and the advantages of the investigated architectures.

[1] L. Lo Bello, W. Steiner, "A Perspective on IEEE Time-Sensitive Networking for Industrial Communication and Automation Systems", Proceedings of the IEEE, vol. 107, no. 6, pp. 1094-1120, June 2019.

[2] "IEEE Standard for Local and Metropolitan Area Networks--Timing and Synchronization for Time-Sensitive Applications," in IEEE Std 802.1AS-2020 (Revision of IEEE Std 802.1AS-2011) , vol., no., pp.1-421, 19 June 2020

[3] "IEEE Standard for Local and metropolitan area networks--Frame Replication and Elimination for Reliability," in IEEE Std 802.1CB-2017 , vol., no., pp.1-102, 27 Oct. 2017.

[4] "IEEE Standard for Local and Metropolitan Area Network--Bridges and Bridged Networks," in IEEE Std 802.1Q-2018 (Revision of IEEE Std 802.1Q-2014) , vol., no., pp.1-1993, 6 July 2018.

[5] L. Leonardi, L. Lo Bello, G. Patti, "Towards Time-sensitive Networking in Heterogeneous Platforms with Virtualization", IEEE International Conference on Emerging Technologies and Factory Automation (ETFA), pp. 1155-1158, Vienna, Austria, Sept. 2020.