Developments in Flexible Time-Triggered Switched Ethernet

Luís Almeida[1], Zahid Iqbal[1], Paulo Pedreiras[2], Ricardo Marau[2], Luis Silva[2], Mohammad Ashjaei[3], Moris Behnam[3], Thomas Nolte[3], Julian Proenza[4], Manuel Barranco[4], David Gessner[4], Guillermo Rodriguez-Navas[4], Alberto Ballesteros[4], Sinisa Derasevic[4]

[1] IT / DEEC – Faculty of Engineering, University of Porto, Porto, Portugal

[2] IT / DETI – University of Aveiro, Aveiro, Portugal

[3] MRTC – Mälardalen University, Västeras, Sweden

[4] DMI –University of the Balearic Islands, Palma de Mallorca, Spain

I. The FTT Paradigm

The Flexible Time-Triggered Switched Ethernet (FTT-SE) protocol was proposed in 2006 [1] as a realization of the FTT paradigm [2] specifically adapted to switched Ethernet. This paradigm first appeared in 1998 implemented on CAN, the FTT-CAN protocol [3], followed by an implementation on shared Ethernet in 2002 [4]. The FTT paradigm is a framework to develop distributed real-time systems that are reconfigurable and adaptive. It has been successfully used for dynamic Quality-of-Service (QoS) management in control applications [5] and video transmission [6], and for dynamic reconfiguration of real-time systems [7] that are either open or include subsystems that operate occasionally.

Essentially, it consists of carrying out time-triggered scheduling online in an open fashion, thus its name. To support prompt and consistent updates in the scheduling, FTT uses a master/multi-slave architecture where a master node schedules the traffic in Elementary Cycles (EC) and propagates the respective triggers through the system, using Trigger Messages (TM). The messages schedule for each cycle is called the EC-Schedule and it is encoded in the respective TM. Upon receiving the TM, the slave (end) nodes transmit the messages indicated in the EC-Schedule.

The flexibility of FTT goes well beyond the online scheduling of the time-triggered traffic. It also provides a framework that is amenable to combine different traffic classes. In particular, FTT protocols have typically supported synchronous (time-triggered) and asynchronous (event-triggered) messages, each type within a dedicated window with a configurable bound, granting desired guaranteed bandwidths.

In this talk we highlight recent FTT developments in the context of switched Ethernet technology. We start with FTT-SE protocol basics and then provide a few lines on improving performance and robustness with the HaRTES switch that is an FTT-enabled Ethernet switch, as well as on improving scalability and dependability of FTT-SE/HaRTES systems.

II. FTT-SE basics

FTT-SE is a software-only implementation of the FTT paradigm that uses plain COTS switches. A software adaptation layer in the end nodes adapts the transmissions to the FTT controls. The FTT master is attached to the switch and controls the traffic sent to the network so that all switch queues are empty by the end of each EC. This way, on an EC scale, the traffic scheduling is fully controlled by the master, providing an abstraction of dynamic channels that can be created, destroyed or adapted on-line and which enforce strong traffic isolation and strictly bounded latency.

In the realm of real-time Ethernet protocols, its flexibility features with real-time guarantees are unique, possibly comparable to much more costly solutions, only, such as layer 3 or 2.5 carrier-grade switches with RSVP-TE or MPLS, AVBridges or Software-Defined Networks (SDNs).

III. The HaRTES switch

FTT-SE relies entirely on COTS hardware, requiring only a modified device-driver at the end nodes, to conform the traffic to the EC-schedule. Consequently, the sphere of control of the Master is limited and, in particular, does not reach the Ethernet switch. Thought sensible, this design decision has associated limitations in terms of reliability, performance and usability. For example, non-FTT or rogue nodes may jeopardize the whole system timeliness. The requirement that all packets must be scheduled at the ingress and all egress ports in the same EC limits the schedulability of multicast and broadcast traffic, since a single full port may block the transmission of packets in several other output ports with available capacity. In addition, the signaling subsystem, used to notify the Master about pending sporadic transmissions, adds an unwanted latency and has scalability issues.

The HaRTES (Hard Real-Time Ethernet Switching) architecture [16] was developed to address those limitations. The base idea is extending the sphere of control of the master to the packet switching subsystem, enabling per packet processing at the ingress and egress ports, separately. Thus, the switch inspects the Ethernet packets upon arrival and determines their nature, i.e., whether they are FTT, real-time legacy or non-real-time packets, verifies their correctness in terms of size, inter-arrival time, etc., and puts them in appropriate queues. Asynchronous packets are then forwarded as soon as possible in the EC asynchronous window while synchronous packets are forwarded when triggered by the master in the EC synchronous window.

Therefore, the HaRTES switch combines online scheduling of isochronous traffic while handling any arriving traffic patterns and requirements. In particular, it allows incorporating non-FTT nodes, possibly through asynchronous real-time virtual channels, without any kind of tampering. Performance

This work was partially supported by the European Commission through FEDER and Erasmus Mundus Action II project EUROWEB, and by the Portuguese Government through FCT grants SENODS CMU-PT/SAI/0045/2009, Serv-CPS PTDC/EEA-AUT/122362/2010, as well as the Spanish Government through project DPI2011-22992 and grant BES-2012-052040 (MEC).

is improved since packets are scheduled independently across egress ports leading to higher link utilization and there is no need for the signaling mechanism of asynchronous messages. Finally, robustness is also improved since packets are subject to validation rules that check conformance to negotiated parameters, both temporal and semantic.

IV. SCALING UP THE NETWORK

Both FTT-SE and HaRTES systems were initially designed with a single switching network element. However, applying the protocol to large scale systems requires expansion to multi-switch architectures. Three different solutions have been proposed the first of which [8] uses a single master that controls traffic transmissions in the whole network. The second solution [9] uses one master attached to each switch controlling the respective local transmissions. All masters in the network are synchronized to control the traffic across switches. Finally, in the third solution, called hybrid [9], the number of required master nodes in the network is reduced with respect to the second case, by assigning one master to control transmissions within a subgroup of switches.

A comparison study based on analytical models and simulations [9] shows that the multi-master architecture has better bandwidth utilization for large networks than the single-master approach. Moreover, the study reveals that the hybrid architecture can further improve the bandwidth utilization performance in large networks.

V. RECONCILING FLEXIBILITY AND DEPENDABILITY

The centralized nature of traffic control in the FTT paradigm creates two obvious single points of failure, namely the FTT master and the communication channel. Thus, applying these protocols in critical applications requires adequate replication schemes.

An initial solution to FTT master replication was developed in the scope of FTT-CAN [10] exploiting specific features of the CAN protocol. Currently, the FT4FTT-Ethernet project is addressing not only the elimination of those single points of failure but, generally, the increase in dependability of Ethernet infrastructures based on a single HaRTES switch by adding adequate fault tolerance mechanisms. An initial solution called Flexible Time-Triggered Replicated Star (FTTRS) replicates both FTT master and switch [11].

The FT4FTT-Ethernet project addresses both transient and permanent faults. In the former case, the project is extending the error detection mechanisms of Ethernet with mechanisms that favor consistent message delivery, roughly corresponding to Atomic Broadcast. This is particularly relevant for protocol critical messages such as the TM. One interesting feature is the simplification achieved in the new Atomic Broadcast protocol developed in FT4FTT-Ethernet when compared with usual approaches, thanks to the global view provided by the HaRTES switch [12].

Concerning permanent hardware faults, and besides the FTT master and switch failures already tolerated by FTTRS, the project also considers node failures by including a replication scheme for node fault tolerance based on restricting nodes failure semantics [13]. This restriction is enforced by the HaRTES switch that prevents the propagation of the most dangerous errors [14]. Moreover, the node replication scheme also supports dynamic changes in the allocation of replicated tasks to nodes [15], opening the way to adaptive fault tolerance consistently with the flexibility of the FTT framework.

Using these mechanisms, the FT4FTT-Ethernet project claims to provide highly-reliable HaRTES-based infrastructures, addressing all relevant system vulnerabilities.

REFERENCES

[1] R. Marau, L. Almeida, P. Pedreiras. Enhancing Real-Time Communication over COTS Ethernet switches. WFCS 2006, IEEE 6th Workshop on Factory Communication Systems, Turin, Italy. June 2006

[2] L. Almeida, P. Pedreiras, J. Ferreira, M. Calha, J. A. Fonseca, R. Marau, V. Silva, E. Martins. On-line QoS adaptation with the Flexible Time-Triggered (FTT) communication paradigm. in Handbook on Real-Time and Embedded Systems, I. Lee, J. Y. Leung, S. Son (eds). CRC Press, ISBN: 9781584886785, July 2007.

[3] L. Almeida, P. Pedreiras, J. Fonseca. The FTT-CAN Protocol: Why and How. IEEE Trans. on Industrial Electronics 49(6):1189-1201, Dec 2002.

[4] P. Pedreiras, P. Gai, L. Almeida, G. Buttazzo. FTT-Ethernet: A Flexible Real-Time Communication Protocol that Supports Dynamic QoS Management on Ethernet-based Systems. IEEE Transactions on Industrial Informatics 1(3):162-172, August 2005.

[5] A. Antunes, P. Pedreiras, L. Almeida, A. Mota. Dynamic Rate and Control Adaptation in Networked Control Systems. INDIN 2007, 5th IEEE Conference on Industrial Informatics, Vienna, Austria, July 2007.

[6] Javier Silvestre, Ricardo Marau, Paulo Pedreiras, Luis Almeida. On-line QoS Management for Multimedia Real-Time Transmission in Industrial Networks, IEEE Transactions on Industrial Electronics, 58(3):1061-1071. DOI: 10.1109/TIE.2010.2049711, March 2011.

[7] Ricardo Marau, Luis Almeida, Mario Sousa, Paulo Pedreiras. A middleware to support dynamic reconfiguration of real-time networks. ETFA 2010 – IEEE Conference on Emerging Technologies for Factory Automation. Bilbao, Spain. 13-16 September, 2010

[8] R. Marau, M. Behnam, Z. Iqbal, P. Silva, L. Almeida, P. Portugal. Controlling multi-switch networks for prompt reconguration. WFCS 2012, 9th Int. Workshop on Factory Communication Systems, May2012.

[9] M. Ashjaei, M. Behnam, L. Almeida, T. Nolte. Performance analysis of master-slave multi-hop switched Ethernet networks.. SIES 2013, 8th IEEE Int. Symposium on Industrial Embedded Systems, July 2013.

[10] Joaquim Ferreira, Luís Almeida, José Alberto Fonseca, Paulo Pedreiras, Ernesto Martins, Guillermo Rodriguez-Navas, Joan Rigo, and Julián Proenza. Combining Operational Flexibility and Dependability in FTT-CAN. IEEE Transactions on Industrial Informatics, 2(2):95–102, 2006

[11] David Gessner, Julián Proenza, Manuel Barranco, and Luís Almeida. Towards a Flexible Time-Triggered Replicated Star for Ethernet. ETFA 2013, 18th IEEE Int Conference on Emerging Technologies and Factory Automation, Cagliari, Italy, September, 2013.

[12] Guillermo Rodriguez-Navas and Julián Proenza. A proposal for Flexible, Real-Time and Consistent Multicast in FTT/HaRTES Switched Ethernet. ETFA 2013, 18th IEEE Int Conference on Emerging Technologies and Factory Automation, Cagliari, Italy, September 2013

[13] Julián Proenza, Manuel Barranco, Joan Llodrà, and Luís Almeida. Using FTT and stars to simplify node replication in CAN-based systems. ETFA 2012, 17th IEEE Int Conference on Emerging Technologies and Factory Automation, Kraków, Poland, September 2012

[14] Alberto Ballesteros, David Gessner, Julián Proenza, Manuel Barranco, and Paulo Pedreiras. Towards Preventing Error Propagation in a Real-Time Ethernet Switch. ETFA 2013, 18th IEEE Int Conf on Emerging Technologies and Factory Automation, Cagliari, Italy, September 2013

[15] Sinisa Derasevic, Julián Proenza, and David Gessner. Towards Dynamic Fault Tolerance on FTT-based Distributed Embedded Systems. ETFA 2013, 18th IEEE Int Conf. on Emerging Technologies and Factory Automation, Cagliari, Italy, September 2013

[16] R. Santos, R. Marau, A. Vieira, P. Pedreiras, A. Oliveira, L. Almeida. "A Synthesizable Ethernet Switch with Enhanced Real-Time Features". IECON'2009 - Porto, Portugal, 3-5 November 2009, pags. 2837-2844