Guest Editorial Special Section on Communication in Automation—Part I

PAST years have witnessed the ever increasing growth of distributed control systems in industrial and factory automation environments [1-2]. In such distributed environments, intelligent devices work jointly to fulfil specific (and complex) tasks, as a consequence of the need to simultaneously increase productivity and decrease production costs. Such kind of distributed architectures enable a higher degree of flexibility, scalability and availability, that could not be otherwise achieved by means of centralised solutions. Furthermore, the inherent ability to support the seamless integration of devices and applications from different manufacturers makes also possible a noticeable reduction of costs.

On the other hand, all the typical properties of centralised control systems, such as determinism, efficiency, safety and fault-tolerance, just to mention a few, have to be maintained and, possibly, enhanced further. This implies that suitable technologies are needed, that allow devices with a non-negligible degree of processing power to co-operate as they were a single piece of equipment.

In such context, communication certainly plays a critical role at all the levels foresaw by the (still valid) Computer Integrated Manufacturing (CIM) model. While interconnecting equipment at the upper factory levels has never been a real problem, managing the adequate timing and safety behaviour at the lower levels, that is, down to the shop-floor, is certainly a non-trivial task.

The birth of factory communication systems dates back to the ’70s, when the need of replacing the existing point-to-point analogue links led to the use of sophisticated (at the time) digital communication technologies. The available solutions for connecting devices at the field level had however several drawbacks: first, the complexity (and costs) of cabling; second, its reduced reliability level, due to the lack of adequate/advanced diagnostic functions; third, the impossibility to perform remote configuration and management operations, which resulted in higher system deployment and maintenance costs; and, finally, the poor accuracy and limited flexibility of the exchanged information.

Since the very beginning, it was clear that the requirements of complex automated production plants – ranging from the interconnection of decentralised peripherals up to the co-ordination among cell controllers – could only be satisfied by means of digital communications technologies relying on non-trivial communication protocols.

At that time, LANs were becoming quite popular in the office automation world. In particular, Ethernet was gaining widespread acceptance, especially because of its efficient medium access technique that allowed for simple and inexpensive implementations. Unfortunately, such type of solutions was felt unsuitable for the factory automation environments. The main reason was due to the random mechanism used to solve (possible) message collisions, which was based on a probabilistic exponential backoff function. What is worse, half-duplex Ethernet suffered seriously from the congestion phenomenon: when the offered traffic grows, there is an higher number of collisions and, consequently, the number of retransmissions increases, which causes an undesired effect of positive feedback on the network load. Therefore, the likelihood of the transmission delays becoming unpredictably long is not negligible, even for the usual operating conditions.

This led to the definition of a number of so-called “fieldbus networks”, most of which backed up by leading companies operating in the field of automation components and equipment. Those solutions were able to show real-time behaviour and therefore to ensure an adequate degree of determinism (i.e., predictability) [3-12].

Popularity of fieldbus networks increased steadily through the ‘90s, as witnessed by the fact that, at present, most factory automation and process control systems still rely on such solutions. However, while adequate to cope with the communication requirements of control systems at the shop-floor, fieldbuses failed to provide a universal and well-accepted solution for supporting distributed control systems in industrial and factory automation environment. In particular, no single fieldbus succeeded to become the *de facto* standard. This implied that a plethora of different (and incompatible) fieldbus solutions have been (and still are) effectively available to designers and system integrators, which means reduced interoperability, increased costs and slow technological advances.

On the other hand, there was a quite different evolution in the office automation environments, where only one solution emerged as the clear winner, that is, Ethernet. Thanks to the switching technology (also known as full-duplex), there was a steady increase in the transmission speed – from initial 10Mb/s equipment, the bit rate grew to 100Mb/s (fast Ethernet), 1Gb/s (gigabit Ethernet) and, recently, up to 10Gb/s. Such continuous enhancement of the transmission technology, managed to keep the pace with the ever increasing communication requirements, and has provided for the three past decades a rock-solid platform on which to build distributed processing systems. The key factor of its success was that backward compatibility was never lost. This means that we are still able to connect old 10 Mb/s shared Ethernet devices to up-to-date 1 Gb/s Ethernet switches.

Thanks to the availability of multiple priority levels and VLANs, high speed switched LANs were deemed suitable to support traffic with tight timing requirements. Hence, they become quickly appealing for factory automation environments, too. As a logical consequence, a great deal of efforts were spent in designing distributed solutions for such environments that were based on Ethernet. This led to the spring of the so called "industrial Ethernet" networks [13-14].

Several research activities have been carried out to find how a truly deterministic behaviour could be achieved when using Ethernet networks [15-19]. Those activities have been carried out by both Academia and R&D departments of leading manufacturers in the automation field. Even though such a target was often obtained at the price of some modifications to standard Ethernet communication equipment, a much higher degree of interoperability with the existing factory communication backbones became possible.

With the advent of the new millennium, several real-time solutions were defined, and the related devices became available on the market. Therefore, they can be readily embedded both in new projects and in the existing automated production plants. The most part of them features true real-time behaviour and performs noticeably better than traditional fieldbus networks. In some cases, very-low-jitter data exchanges and accurate synchronization are provided as well, in order to support motion control applications. Indeed, industrial Ethernet networks aim at being a far more universal solution for industrial communication applications than the traditional fieldbus networks were.

More recently, besides wired networks, attention has been paid to the adoption of wireless communication technologies in industrial environments [20-25]. At present, the main use of such kind of solutions is to simplify the configuration and diagnostic operations through the reduction of cabling. However, the interest is growing about their adoption at the shop-floor as well, in all those cases where wiring is not possible or, simply, cumbersome. Due to economical reasons, research activities have focused mainly on the ways standard wireless solutions, such as IEEE 802.11 WLANs (in all its different variants, including QoS-enabled 802.11e) and IEEE 802.15.4 LR–WPANs (with the related ZigBee protocol stack), could be adapted to operate in industrial environments. Lately, besides WiFi, wireless sensor networks (WSNs) and Zigbee, also Bluetooth is being considered as a possible solution for interconnecting devices at the shop-floor.

It is worth noting, however, that efforts about industrial networks were spent not only on the lower layers of the communication stack. For instance, an ever growing number of industrial applications relies on standard TCP/UDP/IP communication, for both parameterization and process data exchange. The main advantage of IP-based networks is that, they enable natively geographic connectivity and ensure complete compatibility with the existing application protocols in widespread use over the Internet, such as, for example, HTTP, FTP, SMTP, SNMP and so on.

At an even higher level, we assisted in the past few years at the pervasive adoption in industrial environments of a number of technologies, languages and tools for modelling data and systems, that were borrowed directly from the ICT world, such as for example XML and UML. Together with the concept of device profiles, already in use in automated factory environments for about a decade, they achieve an unprecedented level of interoperability and interchangeability. These new technologies allow designers and system integrators to quickly set-up control systems for new production plants or to quickly reconfigure old ones. At the same time, the concept of middleware became more and more popular to link enterprise applications with production management systems [26-29].

The above technological advances also suffered from some drawbacks. Indeed, open networking technologies also mean that production systems are much more exposed to malicious attacks than in the past, which might jeopardize both production and safety. Consequently, particular care has to be taken in modern plants to prevent threats coming from the outside, e.g., from connections to the Internet. This is why, security in industrial environments is becoming a critical requirement, which should be dealt properly since the design phase.

Besides security, also safety – i.e., avoiding that a malfunction in the system might cause either injuries to human beings or serious damages to the production equipment – and fault-tolerance – to increase both the reliability and availability of the plant – are becoming aspects of utmost importance, in order to have highly-dependable systems. Hence, it is no surprise that, nowadays, industrial communication systems have to support them in a proper way.

As it can be seen from the above discussion, factory automation systems and, in particular, industrial communications are quickly converging toward already available standards and solutions of the ICT world. Although the peculiarities of industrial control systems require that some changes are brought to the existing technologies, nevertheless we should expect that in the near future these synergic actions will lead to steadily increasing performances, enhanced interoperability and reduced costs.

Finally, it is worth noting that the current trend to exploit synergic effects between similar application fields goes well beyond the factory and office worlds. This is why, at present, it is much more meaningful to talk about "Communication in Automation", which embraces all those fields that involve the adoption of digital communication techniques for interconnecting devices and equipment in (virtually) every kind of advanced control system. Besides factory automation and process control environments, remarkable examples of such application areas include building automation, motion control and automotive communication systems, just to mention a few.

This special section on "Communication in Automation" presents some relevant works concerning selected aspects about the topics highlighted above. Obviously, it can not provide a comprehensive overview on the subject. Instead, some insight is provided about the most recent advances in this field. As a consequence, the papers included in this special section cover a quite wide spectrum of topics. The section is split over two issues of the IEEE Transactions on Industrial Informatics. In the present issue, three papers are presented, which deal with performance evaluation of control networks. In particular, they address the analysis of scheduling policies for CAN networks, performance measurements in real-time Ethernet networks, and the theoretical evaluation of CSMA-based networks such as LonTalk. Further papers will appear in the May issue of these Transactions.

In order to ensure a deterministic behaviour, distributed real-time control systems are required to meet real-time constraints. In many cases, however, other criteria that depend on the specific application have to be satisfied as well. This is the case of *networked control systems* (NCS), which are known to be quite sensitive to the jitter induced by communication delays. Well known scheduling techniques, such as *non-preemptive deadline monotonic* (NP-DM) and *non-preemptive earliest deadline first* (NP-EDF), which can be directly implemented upon popular networks such as the *controller area network* (CAN), despite efficient in terms of bandwidth usage, may show a poor behaviour when other application-dependent performance indices are considered.

The paper "*Fine-Tuning MAC-Level Protocols for Optimized Real-Time QoS*", by Mathieu Grenier and Nicolas Navet, takes into account a class of on-line scheduling policies that schedule frames right at the MAC level, and provides a schedulability analysis that is valid for all the policies in that class. As shown in the paper, the related algorithms can be implemented on COTS components (e.g., CAN controllers). Moreover, they offer a good tradeoff between feasibility on the one hand and, on the other hand, the ability to fulfil satisfactorily other criteria that depend explicitly on the application, such as those concerning jitters on response times.

Besides fieldbuses, Ethernet is currently more and more adopted in industrial automation environments to carry out real-time communications. Thanks to *real-time Ethernet* (RTE) protocols, defined in the IEC 61784-2 standard, new high-performance automation solutions are now available at reasonable prices. In this kind of systems, the communication cycle can be as low as few tens of µs with jitters below one µs. Such tight timings make network testing and debugging a very complex task. Despite most of the existing network & protocol analyzers are able to perform detailed local analysis, they usually can not be employed to carry out distributed measurements on the whole network. Proper characterization of high-performance RTE systems, in fact, requires that transmission delays are precisely measured, which means that measuring instruments have to be synchronized.

The paper "*A Distributed Instrument for Performance Analysis of Real-Time Ethernet Networks*", by Paolo Ferrari, Alessandra Flammini, Daniele Marioli, and Andrea Taroni, introduces a low-cost distributed tool for measuring the timing characteristics of RTE equipment, e.g., end-to-end delays, synchronization and so on. This instrument relies on multiple probes, implemented by means of FPGAs, that allow time measurements to be carried out on different places of the target network, in a simultaneous and synchronized way. The log of measured data is stored on a *monitor station*, implemented on a PC, for further elaboration. Experimental results, obtained from a prototypal implementation of the instrument, show that synchronization accuracy between probes could be as low as 100 ns, which means that very accurate measurements are possible.

A further kind of control networks, mostly adopted in building automation, rely on random access schemes like CSMA. In order to make such communication techniques suitable for control applications, mechanisms based on priorities should be adopted. In this way, an upper bound on transmission latencies for high-priority messages can be ensured. The introduction of priorities has a significant impact on the overall *quality of service* (QoS) obtained by the different connections, and this may lead to some problems. For instance, if the whole network is not adequately dimensioned, messages having low priorities may suffer from excessively long transmission times, or high loss rates. This means, that model-based analysis techniques must be included in tools for the effective capacity planning of networks.

The paper "*On the Analysis of CSMA-Based Control Nets with Priorities and Multicast*", by Peter Buchholz and Andriy Panchenko, presents an approach to analyse the quality of service provided by CSMA-based control networks. Analytical formulas are derived from queueing networks theory that allow several performance indices, like mean throughput, loss rate and response times, to be efficiently and quickly computed, even for systems with several thousands of connected devices. Particular attention is devoted to the analysis of priorities in LonTalk networks, as well as to the effect of timeouts and multicast communication.

A special section like this one relies on the active support of several people. We would like to thank all of them: the authors, for their contributions and their co-operation in promptly replying to the reviewers' comments; the reviewers, for their careful reviews and comprehensive comments, which contributed in a significant way to the quality level of the papers that were published; and, finally, the editors of these Transactions, for their guidance in preparing and finalizing this special section, from the very first steps until its publication.

GIANLUCA CENA, *Guest Editor*
Italian National Research Council, IEIIT-CNR
10129 Torino, Italy (gianluca.cena@polito.it)

FRANCISCO VASQUES, *Guest Editor*
Faculdade de Engenharia da Universidade do Porto
4200-465 Porto, Portugal (vasques@fe.up.pt)

REFERENCES

[1] J. P. Thomesse, Fieldbus Technology in Industrial Automation, *Proceedings of the IEEE*, Vol. 93, No. 6, pp. 1073-1101, June 2005.

[2] T. Sauter, Fieldbus Systems: History and Evolution, in *The Industrial Communication Technology Handbook*, R. Zurawski (Ed.), CRC Press, pp. 7-1 – 7-39, 2005.

[3] L. Pinho, F. Vasques, Reliable Real-Time Communication in CAN Networks, *IEEE Transactions on Computers*, Vol. 52, No. 12, Dec. 2003, pp. 1594 - 1607.

[4] S. Cavalieri, Meeting Real-Time Constraints in CAN, *IEEE Transactions on Industrial Informatics*, Vol. 1, No. 2, pp. 124-135, 2005.

[5] T. Nolte, M. Nolin, H.A. Hansson, Real-Time Server-Based Communication With CAN, *IEEE Transactions on Industrial Informatics*, Vol. 1, No. 3, pp. 192-201, 2005.

[6] L. Almeida, E. Tovar, J. Fonseca, F. Vasques, Schedulability Analysis of Real-Time Traffic in WorldFIP Networks: an Integrated Approach, *IEEE Transactions on Industrial Electronics*, Vol. 49, No. 5, Oct. 2002.

[7] E. Tovar, F. Vasques, A. Burns, Communication Response Time in P-NET Networks: Worst-Case Analysis Considering the Actual Token Utilisation, *Journal of Real-Time-Systems*, 22, pp. 229-249, 2002.

[8] E. Tovar, F. Vasques, Real-Time Fieldbus Communications Using Profibus Networks, *IEEE Transactions on Industrial Electronics*, Vol. 46, No. 6, pp 1241-1251, Dec. 1999.

[9] M. Barranco, J. Proenza, G. Rodríguez-Navas, L. Almeida, An active star topology for improving fault confinement in CAN networks, *IEEE Transactions on Industrial Informatics*, Vol. 2, No. 2, pp. 78-85, 2006.

[10] G. Cena, A. Valenzano, On the Properties of the Flexible Time Division Multiple Access Technique, *IEEE Transactions on Industrial Informatics*, Vol. 2, No. 2, pp. 86-94, 2006.

[11] L. Almeida, P. Pedreiras, J. Fonseca, The FTT-CAN protocol: why and how, *IEEE Transactions on Industrial Electronics*, Vol. 49, No. 6, pp. 1189-1201, Dec. 2002.

[12] J. Ferreira, L. Almeida, J. Fonseca, P. Pedreiras, G. Rodríguez-Navas, J. Rigo, J. Proenza, "Combining Operational Flexibility and Dependability in FTT-CAN", *IEEE Transactions on Industrial Informatics*, Vol. 2, No. 2, pp. 95-102, 2006.

[13] J.D. Decotignie, Ethernet-Based Real-Time and Industrial Communications, *Proceedings of the IEEE*, Vol. 93, No. 6, pp. 1102-1117, June 2005.

[14] M. Felser, T. Sauter, Standardization of Industrial Ethernet – the next battlefield?, in *Proc. IEEE 5th International Workshop on Factory Communication Systems (WFCS 2004)*, Wien, Sept. 22-24, 2004, pp. 413-421.

[15] L. Lo Bello, O. Mirabella, Design Issues for Ethernet in Automation, in *Proc IEEE 8th International Conference on Emerging Technologies and Factory Automation (ETFA 2001)*, Vol. 1, pp. 213-221.

[16] L. Lo Bello, G.A. Kaczynski, O. Mirabella, Improving the Real-Time Behavior of Ethernet Networks Using Traffic Smoothing, *IEEE Transactions on Industrial Informatics*, Vol. 1, No. 3, pp. 151-161, 2005.

[17] P. Pedreiras, P. Gai, L. Almeida, G.C. Buttazzo, FTT-Ethernet: A Flexible Real-Time Communication Protocol That Supports Dynamic QoS Management on Ethernet-Based Systems, *IEEE Transactions on Industrial Informatics*, Vol. 1, No. 3, pp. 162-172, 2005.

[18] T. Skeie, S. Johannessen, O. Holmeide, Timeliness of real-time IP communication in switched industrial Ethernet networks, *IEEE Transactions on Industrial Informatics*, Vol. 2, No. 1, pp. 25-39, 2006.

[19] F. Carreiro, R. Moraes, J. A. Fonseca, F. Vasques, Real-Time Communication in Unconstrained Shared Ethernet Networks: The Virtual Token-Passing Approach, in *Proc. IEEE 10th International Conference on Emerging Technologies and Factory Automation (ETFA 2005)*, Catania, Italy, 2005.

[20] A. Willig, K. Matheus, A. Wolisz, Wireless Technology in Industrial Networks, *Proceedings of the IEEE*, Vol. 93, No. 6, pp. 1130-1150, June 2005.

[21] G. Cena, I. Cibrario Bertolotti, A. Valenzano, C. Zunino, Evaluation of Response Times in Industrial WLANs, *IEEE Transactions on Industrial Informatics*, Vol. 3, No. 3, pp. 191-201, Aug. 2007.

[22] F. De Pellegrini, D. Miorandi, S. Vitturi, A. Zanella, On the use of wireless networks at low level of factory automation systems, *IEEE Transactions on Industrial Informatics*, Vol. 2, No. 2, pp. 129-143, May 2006.

[23] A. Willig, M. Kubisch, C. Hoene, A. Wolisz, Measurements of a wireless link in an industrial environment using an IEEE 802.11-compliant physical layer, *IEEE Transactions on Industrial Electronics*, Vol. 49, No. 6, pp. 1265-1282, Dec. 2002.

[24] S. Lee, K. C. Lee, M. H. Lee, F. Harashima, Integration of mobile vehicles for automated material handling using Profibus and IEEE 802.11 networks, *IEEE Transaction on Industrial Electronics*, Vol. 49, No. 3, pp. 693-701, June 2002.

[25] C. Koulamas, S. Koubias, G. Papadopoulos, Using Cut-Through forwarding to retain the Real-Time Properties of Profibus over Hybrid Wired/Wireless Architectures, *IEEE Transactions on Industrial Electronics*, Vol. 51, No. 6, pp. 1208-1217, Dec. 2004.

[26] J. Peschke, A. Lüder, A. Klostermeyer, A. Bratoukhine, T. Sauter, Distributed Automation: PABADIS vs. HMS, *IEEE Transactions on Industrial Informatics*, Vol. 1, No. 1, pp. 31-38, 2005.

[27] F. Jammes, H. Smit, Service-Oriented Paradigms in Industrial Automation, *IEEE Transactions on Industrial Informatics*, Vol. 1, No. 1, pp. 62-70, 2005.

[28] J.L.M. Lastra, M. Delamer, Semantic web services in factory automation: fundamental insights and research roadmap, *IEEE Transactions on Industrial Informatics*, Vol. 2, No. 1, pp. 1-11, 2006.

[29] A. Kalogeras, J. Gialellis, C. Alexacos, M. Georgoudakis, S. Koubias, Vertical Integration of Enterprise Industrial Systems Utilizing Web Services, *IEEE Transactions on Industrial Informatics*, Vol. 2, No. 2, pp. 120-128, 2006.

Gianluca Cena was born in Torino, Italy, in 1966. He received the Laurea degree in electronic engineering and the Ph.D. degree in computer engineering from the Politecnico di Torino, Italy, in 1991 and 1996, respectively. In 1995, he became Assistant Professor at the Department of Computer Engineering of the Politecnico di Torino and, in 2001, he joined the Italian National Research Council (CNR) as a Senior Researcher. Since 2005, he has been Director of Research with the Istituto di Elettronica e di Ingegneria dell'Informazione e delle Telecomunicazioni (IEIIT-CNR), where he is engaged in research activities concerning communications in manufacturing and automotive environments.

He is the author of many technical papers in the area of computer communications. His current research interests include industrial communication systems, wired and wireless protocols and real-time networks. He is also active in the international scientific community working on these subjects, and served as Program Co-Chairman for the 2006 and 2008 editions of the IEEE Workshop on Factory Communication Systems.

Francisco Vasques (M'00) was born in Porto, Portugal, in 1964. He received the "Licenciatura" degree in electrical engineering from the University of Porto, Portugal, in 1987 and both the MSc and PhD degrees in computer science from LAAS-CNRS, Toulouse, France, in 1992 and 1996, respectively. Since 2004, he is Associate Professor of the University of Porto. His current research interests include real-time communication systems, factory communications, fault-tolerant systems, and real-time system architectures. He served as Program Co-Chairman for the 2000, 2004 and 2006 editions of the IEEE Workshop on Factory Communication Systems. He is author or coauthor of more than 75 technical papers in the area of real-time systems and factory communications. He is Associate Editor of the IEEE Transactions on Industrial Informatics since 2007.