Remote Device Command and Resource Sharing over the Internet: a New Approach Based on a Distributed Layered Architecture

Francisco José Monaco, Adilson Gonzaga

Abstract—

In addition to the remote access and computer-augmented functionality brought about by the earliest modalities of distance operation, technical advances in the form of telematics have opened up a whole new range of applications, one of which, resource sharing, deserves special attention. Nonetheless, while distance operations through computer networks, and particularly over the Internet, have attracted a great deal of attention in recent years, there is still a noticeable lack of important acquisitions regarding the systemic treatment of essential issues in this field. This paper presents an overview of the current trends in this emerging interdisciplinary area and briefly comments on the fundaments of telematics-supported distance operation. A case study is used to report on an experience involving methodological investigations in this area.

Department of Electrical Engineering, São Carlos School of Engineering, University of São Paulo, Brazil.

Keywords—**Telematics, teleoperation, telecommand, network-based remote control, Web devices, Internet, interactive systems**

I. Introduction

Ever since he invented the most rudimentary tools to dig the ground for edible roots and to search for food inside animal burrows, man has always made use of artificial means to extend his reach, a quest that has become increasingly audacious as technological knowledge continuously advances. Over the last 50 years, exhaustive research has focused on developing the field of distance operations in several application domains. Indeed, the extent to which these operations have been developed and the incongruent terminologies that pervade this inherently multidisciplinary area make them difficult to discuss. However, the underlying technology can be divided into three major classes of distance operation systems.

Mechanical distance operation originated from the first telemanipulators (aka telechirichs, an alternative to the former somewhat etymologically imperfect Greek-Latin world [1] *apud* [2]), direct descendants of the hand-held blacksmith thongs. Telemanipulators were intensively used in the early decades of the 20^{th} century in scientific and industrial environments, particularly in nuclear and military installations, in order to avoid exposure to dangerous radioactive or chemical substances [2]. The main purpose of such equipments, purely mechanical or enhanced by hydraulic and pneumatic mechanisms, was to *manipulate* (handle and move) objects at a distance, specially when direct contact was to be avoided or was impossible due to technical or safety reasons.

The next stage brought *electronic* distance operation, which extended the concept beyond mechanical manipulation to general *device driving*. The technology spread over a large field of applications, including industrial machinery and household appliances. Starting with electric motorized effectors, this technology further benefited from new advances that allowed for long distance communication, which sometimes caused the non-univocally interpreted prefix "tele" to be strictly associated with 'telecommunications', and the term telecommand [3] to signify the transmission of command signals through these systems.

Computer-aided distance operations, which introduced this technology into the age of *automation*, only actually became effective in the 70's. Programmable mechanical devices then led to telerobotics and remote control - used here with the specific meaning associated to the discipline of automation, since it is also the commercial name for devices that are part of modern home appliances. Along with telecommands, remote automation is seen as an important acquisition that has driven the evolution of remotely operated mobile devices, also referred to as teleoperators [2] in some contexts.

II. A New Distance Operation Technology

Today, a new connection that promises revolutionary possibilities is emerging between distance operation and another state-of-the-art technology - *Telematics*.

From the standpoint of techno-social paradigms of distance interaction through information technology, Telematics represents the application of information processing and transmission to overcome general constraints that lessen the effectiveness of user services. Thus, the domain of Telematics involves neither computing nor communication systems but the utility resources that can be developed by bridging the capabilities deriving from their association. From on-line electronic teller machines and information terminals to the myriad interactive user services available through the Internet, telematic applications are tuned to the current trends of an interconnected society.

While both areas have developed independently over the last few decades, with distance operations seeking advances in remote access to general devices - guided by technical, safety and comfort-related needs, and telematics mitigating distances - of a geographical, temporal, cultural or other nature, their union opens up new and exciting possibilities. Indeed, the interconnection of devices through digital networks represents not only on-line availability of their functionality or the simple transmission of remote commands among them. Their significance lies in the fact that, more than channels of communication like so many others, networks constitute computational environments within which complex integrated systems and services can be built.

As access to telematic resources becomes increasingly commonplace in offices, factories and even in the home and with the advent of global networks, with the Internet and its world-wide scope so much in evidence, the idea of sharing general-purpose devices - from intelligent home appliances to data acquisition equipment, high-precision scientific instruments and costly or rare industrial machinery - has become not only a viable purpose but an intensively pursued goal. In fact, the merging of distributed and ubiquitous computing through the distance operation of electronic devices has given rise to a variety of new concepts such as pervasive systems, wearable devices and other current trends.

III. From Coke Servers to Web Devices

Like many other interesting creations of Internet technology, the first examples of distance operation over wide-area networks were motivated by the simple initiatives of people attempting to solve trivial problems or perhaps looking for alternative forms of entertainment. Among the earliest experiments were the legendary "coke dispensers", which appeared for the first time in the 70's, when some students at Carnegie Mellon University connected a coke machine in a computer room to the Internet, enabling anyone to read the output of the electronic sensors attached to it and to verify, from a distance, whether the machine contained soda cans and whether they were cold - thereby avoiding frustrating walks through the building [4]. This curious experiment became widely known and was soon reproduced at several universities in the United States and Europe, where alternative versions employed other "unconventional" devices such as the University of Cambridge's famous coffee pot [5], the networked washing machines at the Massachusetts Institute of Technology [6] and other similar projects that achieved the status of classical experiments for engineering students.

Despite their intriguing nature, these early initiatives inspired further development in the field by demonstrating that the idea was feasible and widely accepted. From the first experiments to today's operational systems, the idea has gained increasing consistence and has motivated notable research projects such as the Mercury Project at the University of Southern California [7], the Internet Tele-Robot at the University of Western Australia [8], the Tele-Robotic Telescope at Bradford University [9] and the Remote Microscope of Oak Ridge National Laboratory [10], among others.

The attention that distance operation through the Internet has received is justified by its global coverage and its increasingly accessible technology, not only in terms of cost but also of easy use. While the first coke dispensers were built on standard *finger* and *telnet* Internet services, owing to its advantageous features as a universal network user interface, the Web is considered one of the most promising technologies in this area. This fact is illustrated by the growing number of new *Web devices*, the popular technomania of the

90's, which pop up every day on the Internet: Web cams, lava lamps, weather monitors, etc.

Despite its strong and long-standing appeal in the field of science fiction, distance operation through the Internet in the scientific, education and industrial sectors has only recently come to the attention of companies and research centers, which are seriously interested in its application in telemetry, remote instrumentation, tele-manufacturing [11], distance learning, telemedicine [12] and other promising areas.

IV. Telematics-supported distance operation

Motivated by our growing interest in the association of telematics and distance operation, our current efforts in this sphere focus on an in-depth investigation of its essential technological aspects and are based on a systematic approach that deals with network-based distance operation from the telematic perspective. This, in turn, is the basis for the fourth generation of the technology, which we refer to as *telematics-supported distance operation* (TSDO).

In addition to the remote access and computer-augmented functionality of the early modalities of distance operation, the association of computing and communication produced by telematics opens up a new range of applications, one of which deserves special attention here: *resource sharing*. This is possible owing to the elementary features of the technologies involved, which provide both a multipoint communication channel and the essential computational capabilities needed for multi-user management, e.g., access control, resource arbitration, maintenance of reliability and dependability and other important mechanisms.

Thus, in the light of an integrated view of the fundamental principles of telematics and distance operation, this report describes our progress in the development of methodological investigations on this theme.

V. State-of-Art

While distance operation through computer networks, and particularly over the Internet, has attracted much attention in recent years, this field is still noticeably lacking in important acquisitions regarding the systemic treatment of its essential issues. There is, for instance, no broadly recognized conceptual definition for its specific paradigms, and the formal technical approaches are equally insufficient.

As demonstrated by scientific experience in this area, this situation may be partially attributed to the fact that network-based distance operation is often claimed to be the legitimate competency of this or that technologically specialized area, although each area actually lends it the peculiar features of its own formalism, models, and methodological principles. In short, distance operation is sadly lacking in the characteristic coherency of well-established research lines.

In practice, the lack of consensus regarding terminology, development processes and other resources which are common in well-consolidated disciplines has led to a noticeably poor performance of the available implementations and to these systems' inflexibility in face of the generic configurations of real-world needs - i.e., those inherently associated with distance operation over telematic networks. Thus, the proliferation of incipient implementations and poor-quality applications hinders the progress of this technology by negatively influencing its credibility.

Notwithstanding their popularity, *Web devices*, mentioned again for purposes of illustration, can hardly be cited as examples of well-engineered network-based distance operation systems insofar as they are basically improvisations on adapted legacy systems, requiring artful maneuvers in order to overcome inadequacies of the available techniques. Indeed, due to the lack of a methodological basis, even the representative work of acknowledged contributions to the state of the art often corresponds to *ad hoc* artifices rather than to systematically addressed solutions [13]. To name but a few aspects that require further investigation in this respect, there is the questionable trend of converting everything into Web-based applications, often with the sole purpose of keeping up to date with the latest innovations, disregarding the specific requirements of the particular application and the technical features of available implementation resources.

Starting from the teleoperation theory, for instance, the evolution from the early master-slave (irreversible) to bilateral (reversible) control system is recognized as an important step [2]. Web devices, on the other hand, fall into the former category. The fact is not all real-world systems can be suitably modeled under the asymmetric condition implicit in the client-server behavior of elementary Web interactions; indeed, many systems are known to fit better into the peer-to-peer model and distributed organization. A practical example of this situation are the conventional CGI-based Web devices, in which the operator sends commands to the remote device through the Web browser at any time. However, the standard master-synchronized scheme prevents the server at the device end from sending signals to the operator to notify him about external events or from providing full cybernetic behavior.

Another point that has been ignored is the inadequacy of the Web's underlying protocol (basically HTTP, when it is built on top of TCP/IP network), which is conveniently designed for hypermedia document retrieval but not for general distance operation systems. Connection-oriented interaction is quite suitable, for example, when the interval between successive requests is noticeably greater than the time spent for data transfers (e.g., when the user reads the document he has just retrieved before issuing a new request). On the other hand, machine driving, which is the case of distance operation, is often expected to display the opposite behavior, which suggests a connectionless model (possibly on top of a reliable connection channel). An analogous note applies to real-time and dependable systems, which deserve special attention at least until such time as current Web technology imposes certain restrictions - despite the efforts that have recently been dedicated to this matter.

Further reflection on these and other related issues clearly reveals that telematics-supported distance operation is actually neither a particular case of enhanced Web application nor a simple extension of network-based resource control, such as conventional file systems and printer sharing. Instead, it is a special area of telematics in which conceptual and technical principles of remote command must be taken into consideration - a view that requires further development.

By means of a case study, we describe an ongoing project in which the above-mentioned approach has been successfully applied.

VI. Case study: The Distributed Device Network

One of the interesting features of telematic distance operation over computer networks is its distributed nature, which is the reason for the decentralised organisation of the experimental TSDO architecture being developed at LAMI[1].

This project is based on a more general concept, which we refer to as *cooperating networked devices* (CNDs). A CND is a virtual working environment wherein interconnected devices interact cooperatively, requesting and providing services and sharing resources among themselves and with the other elements in the network. Several aspects of this idea have been explored by recent technologies, particularly in industrial control engineering and intelligent residential automation.

The IDN (*interactive-device network*) project is a materialization of the CND concept that provides an experimental interaction model for networked devices and implements the basic features of a multi-user resource sharing system. It has been designed as a telematics-supported distance operation and is based on its specific application features. While the goal of this paper is not to describe IDN internals, we believe that this example is useful to illustrate some essential methodological issues that should be addressed in systematically designed TSDO systems.

[1]Laboratory of Microprocessors, SEL-EESC/USP

A. Comprising Elements

Distance operation systems consist of three basic parts: the *operator*, the command-driven *equipment* and the *control mechanism*. The latter corresponds to the controller interfaces at both the operator and equipment ends and to the communication mechanisms that allow them to exchange control information. In this scheme, generic equipment is often represented by an *interactive system* [14], which, by definition, is a responsive system that produces predictable outputs in response to user inputs. When the *response time* of the interactive system is expected to be limited to a maximum known value, it is said to match the *real-time* criteria [15] (hard real-time if failure to respond within that time implies failure; soft real-time, the contrary). With this design approach, the IDN architecture results from the interconnection of such systems in a network, through which they interact in a distributed organization.

The first limitation we found when experimenting with "traditional" network-based telecommand mechanisms (if one can call such a recent technology traditional) was precisely the asymmetrical behavior of the two distinct entities that form the client-server pair. While this is quite convenient to model systems in which the roles of the user and the interactive system itself are unambiguously determined -as in file retrieval systems, e.g., the Web, this restriction should be avoided in order to cover the wider spectrum of cases that are to be treated by general distance operation.

It is worth noting that, to a large extent, the user and interactive device functions are interchangeable between the operator and controlled equipment, e.g., when the reactions of both elements are driven by asynchronous events signaled by the corresponding partner (as when the pair is an autonomous machine-machine composition). Therefore, to comply with the generality required by the proposed application, operator-device interaction should be symmetric, i.e., both should be able to send and respond to requests asynchronously at any time. Non-adherence to this

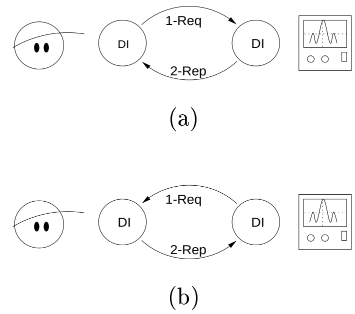

Fig. 1. The symmetric behaviour of the device interfaces.

requirement limits the system's applicability to cases in which the master-slave condition is acceptable.

The above considerations, therefore, suggest the relevance of designing IDN with only one symmetric *device interface* (DI) that can operate in both active and passive configurations. Figure 1(a) shows an operator sending a command request to the controlled equipment. The DI on the left forms with that on the right, in that order, a client-server pair. The opposite is shown in figure 1(b): the equipment sends a non-requested signal to the operator (e.g., an alarm); since the operator is now playing the role of server, it must respond to the client.

Complementary to the DIs, the other component of IDN is the *interface manager* (IM), which is actually a central unit to which all the DIs that

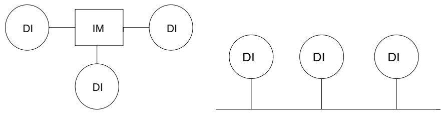

Fig. 2. The real network and the virtual network

participate in the device network must first connect when entering the system, as illustrated in figure 2(a). Once the connection is established, however, the IM becomes transparent, serving only as a channel for the flow of messages exchanged by the DIs, as shown in 2(b), and acting as an omniscient moderator that intervenes in the normal communication process only in response to any abnormal condition (such as requests that cannot be met and unexpected disconnections).

The arrangement is organized similarly to those of modern operating systems. It has been stated that, from the OS point of view, the user is a peripheral (with unpredictable behavior), which means that there is no practical difference between electronic equipment and a human user - both are seen through the I/O system and, therefore, the OS does not distinguish between characters typed by the user on the keyboard and data read from the storage unit. This analogy is in accordance with DIs and IM interaction and, to expand this comparison, the former may be likened to computer device drivers (interfacing with input and output hardware), while the latter is comparable to the operating system kernel (controlling devices and provision for communication among them).

B. The Virtual Working Environment

Once connected to the system, all devices have access to the virtual working environment, where they receive a common name space with transparent address resolution and shared resource management. Up to this point, the differences between command and operation have not been mentioned. Nonetheless, from their literal acceptation, it is recognized that the former consists of a much simpler task than the latter, since issuing control signals does not represent real operation, which, in turn, is only effective if it produces useful work, i.e., if the operator succeeds in controlling the device — by issuing commands — according to his planned action.

The effectiveness of distance operation, in particular, depends on how adequately the system's features support the execution of tasks. Although a more in-depth discussion of such details is beyond the scope of this paper, two aspects are worth commenting on. The first refers to the sequential correlation that characterizes the operation, which expresses how consecutive commands are inter-related; the other is the concurrent correlation that describes how command sequences executed concurrently interfere with each other. When these dependencies have little relevance, as in many Internet client-driven user applications, little or no long-term control is required from the system. In TSDO, however, IDNs rely on the concept of *sessions* to cope with the complexity arising from these relations. A set of parameters is associated with each session, providing information on the system's status and event memory, which is important information for concurrency control.

The second aspect involves the coupling factor associated with the distance operation system, which estimates the level of sensory feedback in a given application. When no extrasensory path is available (e.g., owing to great distances), the operation relies solely on the interaction provided by the system itself. This is the case we call *remote operation*, requiring special treatment, which

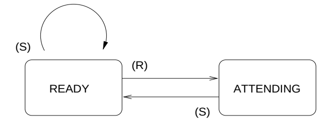

Fig. 3. The two-state DI

is part of the system's requirements and is a feature of the virtual environment that provides for this interactivity. It is often desirable to implement levels of automation and autonomy on behalf of the user so as to artificially overcome natural deficiencies imposed by limited feedback and long response times.

C. Basic Interactions

A widely known form of representing real-time systems is through finite-state machines (FSMs) [14], owing to their good representation of how the status of deterministic systems changes upon the occurrence of a known set of events. When described by a FSM, every DI in the system is a two-state machine whose possible states are REDY and ATTENDING, as shown in figure 3.

Two events are possible in the REDY state. One occurs when the DI receives (R) a message: it always treats incoming messages as requests and then evolves into the ATTENDING state, in which it begins to process the requested service. When the DI concludes the task, it sends (S) a reply message, returning to the REDY state. Note that this scheme is analogous to the common operation of traditional network server software, but the same does not apply when an S event occurs in the REDY state. In conventional network applications, a client entity that sends a request changes from the ready to the waiting state while, in DIs, an S event causes no change in the state; instead, this action is comparable to an asynchronous interrupt signal. A useful analogy is a computer device driver, which behaves as a server by accepting requests from the operating system and also triggers hardware interruptions in response to external events in a stateless configuration.

The interface manager stores information about the flow of messages among the system's devices in an internal FSM that contains the whole system status; the FSM, in turn, depends on the states of every single DI on the network. For each DI, a pair in the form of (wf, rt) is maintained in an array, which is updated every time a message is sent by any device. The data field denotes that a message from a given device was sent to the DI identified internally as wf (every connected DI is associated with an internal identification number) and, therefore, it is now *waiting for* a reply from this same device. Likewise, the DI identified as rt is waiting for the response of a previously received message to which it must *reply*.

Based on this status information and on additional configurations provided by the user, the IM may decide how messages should be delivered in order to ensure that every reply message reaches the device that is waiting for it, that busy devices are not overloaded with new requests and that only authorized requests are further processed. This procedure is carried out by a *delivery algorithm*. The implemented prototype uses standard scheduling policies and deadlock prevention strategies to deal with the complexity of distributed resource allocation required in shared access management.

D. The Distributed Layered Approach

Developments in the field of distributed applications are particularly useful in network-based TSDO, informing us about the benefits of system function decentralization among specialized independent units, the flexibility offered by the sepa-

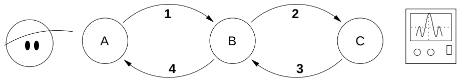

Fig. 4. A layered composition

ration of policies and mechanisms, and about layered functional organizations. The IDN architecture uses all these concepts to simplify system design and to improve its functionality by grouping devices into abstract specialized layers (figure 4).

While many real world distributed systems, especially computer operating systems, are designed to work within closed environments (restricted to the LAN domain, for example), telematics-supported distance operation is often conceived to provide remote command and resource sharing over wide-area networks. In addition to the intrinsic difficulties associated with information transmission over long distances (delay, latency etc.), the distinction between LAN and WAN is not merely a matter of extension, but of the technology associated to them. Interconnecting distinct networks spread over heterogeneous communication channels produces infrastructures with little reliability inasmuch as the user has no control over the features of underlying resources operated by third parties. Therefore, special care must be taken to ensure the system's integrity, since precise assumptions cannot be made about the channel's quality. Likewise, security concerns must be considered and, since it is impracticable to avoid attempted intrusions (as can often occur within LAN setups), authentication and access control mechanisms are crucial.

D.1 Experiences with IDN Project

he IDN is still an ongoing project and no conclusive experimental results have yet been reported, although preliminary tests with an implemented prototype[2] have successfully fulfilled our expectations regarding flexibility requirements. Although the approach founded on the principle of distance operation and telematics has successfully overcome several limitations found in non-systematically designed Internet-based remote commanded experiments in terms of performance and functionality, the pilot system still has to be subjected to a well-defined evaluation plan before conclusive results can be published.

[2]The system is a TCP/IP software suite designed to work through the Internet.

VII. Conclusions

Telematics-supported distance operation is a technology that promises to meet the latest aspirations of the information age. As a result of the ever increasing popularity of the Internet, network-based remote command and resource sharing systems have been investigated in several application domains including distance learning, telemedicine, industrial tele-robotics, remote automation, remote scientific instrumentation, CSCW and tele-manufacturing.

The association of the two elementary capabilities of telematic systems, computing and communication are the building blocks for effective worldwide resource sharing systems - a concept that has already inspired several research projects. These projects range from simple Internet-based telemetry to autonomous intelligent device networks and advanced multi-institutional endeavors involving the creation of large virtual environments, in which specialists can carry out their collaborative tasks irrespectively of geographic constraints, interacting with people, databases and shared instruments as though they worked in the same physical location [16].

In this paper, we have presented an overview of the current trends in the emergent field of remote command through computer networks and have briefly commented on the foundations of telematics-supported distance operation. Using a case study to introduce the IDN project, we have also reported on our experiences with methodological investigations in this area.

Our observations of today's scenario suggest that the somewhat inadequate combination of excessive enthusiasm for the latest developments in network technology and the lack of a systemic treatment of the multidisciplinary issues involved in this rapidly evolving field is responsible for the proliferation of incipient real-world systems. The purpose of our contribution is to share the results of our investigative efforts with other researchers in this and correlated areas so that a convergence of cooperative action in this field can be achieved.

Acknowledgements

The authors are grateful to the São Paulo State Foundation for Research Support (*Fundação de Amparo à Pesquisa do Estado de São Paulo - FAPESP)* for its financial support of this project.

References

[1] M. W. Thring. Telechiric mining. In *2nd R.O.MAN.SY*, page 459, 1976.

[2] J. Vertut and P. Coiffet. *Teleoperation and Robotics*, volume 3A of *Robot Technology*. The Anchor Press Ltd, Lodon, 1985.

[3] National Telecommunications The Institute for Telecommunication Sciences and Information Administration. Federal standard 1037c, 7 Aug. 1996. http://ntia.its.bldrdoc.gov/fs-1037/dir-036_5347.htm.

[4] CMU. CMU coke machine. Computer Science Department, Carenegie-Mellon University, 1998. http://www.cs.cmu.edu/ coke/.

[5] M Johnson. Cambridge university's internet tojan room coffee machine. Cambridge University, Trojan Room, 1994. http://www.cl.cam.ac.uk/coffee/coffee.html.

[6] Jacob B. Schwartz. Welcome to random hall laundry. at Massachussets Institute of Tecnology Web Pages, 1997. http://spleen.mit.edu/laundry.

[7] K. Goldberg, M. Mascha, S. Genter, N. Rothenberg, C. Sutter, and J. Wiegley. Beyound the web: Excavating the real world via mosaic - the Mercury project. In *Second International WWW Conference*, Chicago, USA, Oct. 1994.

[8] K. Taylor and J. Trevelyan. A telerobot on the world wide web. In *Robots for Australian Industries Conference*, pages 108–120, Melbourn, Australia, Jul 1995. Australian Robot Association.

[9] M.J. Cox and J.E.F. Baruch. Robotic telescopes: An interactive exhibit on the world wide web. In *2nd International Conference of the World Wide Web*, Chicago, USA, Oct 1994.

[10] ORNL. ORLN's telepresence connects researchers with remote microscopes. Oak Ridge National Laboratory, Press Release, 1996. http://www.ornl.gov/Press_Release/archive/960627-01.html.

[11] F. J. Monaco. Toward distributed cooperative telemanufacturing over the internet. In *proceedings of 15th International Conference on CAD/CAM, Robotics & Factories of the Future*, volume 1, pages MT7–17, guas de Lindoia, Brazil, Aug 1999.

[12] J. E. Borelli, F. J. Monaco, E. A. Neves, and A. Gonzaga. An application of telediagnosis in the investigation of pathological condition and remote monitoring of individuals with disabilities - tissue teparation of leg trofic ulcers. In *Iberdiscape 2000, III Congresso Iberoamericano de de Comunicacin Alternativa y Aumentativa, I Congresso Iberoamericano de Tecnologas de Apoyo para la Discapacidad*, pages 435–437, 19-20 Oct 2000.

[13] F. Monaco and A Gonzaga. Web-based system paradigms and the conceptual design of network applications. In *I Workshop on Web Engineering, 21th International Conference on Sofware Engineering*, Los Angeles, USA, 16-22 May 1999.

[14] P. Laplante. *Real-Time Systems Design and Analysis.* IEEE Press, 1993.

[15] S. Levi and A. K. Agrawala. *Real-Time System Design.* McGraw-Hill Book Company, 1990.

[16] V.G. Cerf et al. *National Collaboratories: Applying Information Technology for Scientifc Research.* National Academic Press, Washington DC, USA, 1993.