

BY ELIAS P. DUARTE JR., RAIMUNDO J. A. MACÊDO,  
ELIANE MARTINS, AND SERGIO RAJSBAUM

# A Tour of Dependable Computing Research in Latin America

COMPUTING TECHNOLOGY HAS become pervasive and with it the expectation for its ready availability when needed, thus basically at all times. Dependability is the set of techniques to build, configure, operate, and manage computer systems to ensure that they are reliable, available, safe, and secure.<sup>1</sup> But alas, faults seem to be inherent to computer systems. Components can simply crash or produce incorrect output due to hardware or software bugs or can be invaded by impostors that orchestrate their behavior. Fault tolerance is the ability to enable a system as a whole to continue operating correctly and with acceptable performance, even if some of its components are faulty.<sup>3</sup>

Fault tolerance is not new; von Neumann himself designed techniques for computers to survive faults.<sup>4</sup>

The premiere conference in the area, the IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), held its 50<sup>th</sup> edition in 2020. In Latin America, the first edition of the Brazilian Symposium on Fault-Tolerant Computers (SCTF) was held in 1985 and took place for 18 consecutive years. In 2003 it evolved into the Latin American Symposium on Dependable Computing, which has since been held in multiple countries. Today, research groups are firmly es-



established, and premier international events have been held in the region, such as DSN in Rio de Janeiro in 2015, the ACM Principles of Distributed Computing (PODC) conference in Mexico in 1998, and the IEEE International Symposium on Reliable Distributed Systems, which has been held twice in Brazil, in Florianópolis in 2004, and in Salvador in 2018, among others.

The first research efforts on dependable computing in Latin America

were focused on aerospace systems. Safety is a dependability property that ensures that a system is capable of avoiding catastrophic consequences. It must be assessed for applications that perform life- or material-critical tasks and is thus a strict requirement of aerospace systems. Established in São Paulo, Brazil, in the early 1960s, the National Institute for Space Research (INPE) has since fostered research in multiple fields including space, climate, and computer

science. At the same time, the aerospace industry took large strides in Brazil, which included the creation of airplane maker Embraer. It is no coincidence that the first SCTF was held at INPE in 1985, chaired by Alderico Rodrigues de Paula Jr., who led the project that developed the fault-tolerant computer launched with the SCD-1 satellite in 1993.<sup>a</sup>

<sup>a</sup> [http://www.inpe.br/scd1/site\\_scd/historico.htm](http://www.inpe.br/scd1/site_scd/historico.htm) (Portuguese)



## Replicating servers over computer networks is one of the building blocks commonly applied to improve the reliability and availability of distributed services.

Since then, research in dependable computing in Latin America (LATAM) has covered many topics, including hardware, software, and communications, from both a practical and theoretical perspective. This article gives a general view of LATAM research in the field and highlights some main results. Please note that we have set up a Web page with an accompanying bibliography.<sup>b</sup>

### **An Overview of Dependable Computing Research in LATAM**

The development of dependable systems requires a combination of fault prevention and fault tolerance techniques to achieve the desired level of reliability, availability, and safety, among other attributes. As such, a major focus of research in LATAM has been the design of hardware and software fault-tolerant architectures to protect systems against faults, either accidental or malicious, complemented by dependability assessment analysis. Code inspection, model-checking, and testing have been widely applied on various LATAM projects, and model-based testing has sometimes been used to generate test cases. Special attention has been devoted to fault injection techniques, not only to dynamically verify the efficacy of fault-tolerant mechanisms, but also to estimate parameters used in analytical models for quantitative evaluation.

On the other hand, architecting dependable systems usually requires the understanding and construction of basic common building blocks or patterns to be reused across multiple application scenarios. Replicating servers over computer networks is one of those building blocks commonly applied to improve the reliability and availability of distributed services—the so-called state machine approach. At the core of this approach are basic problems such as consensus, group communication, group membership, and failure detection. The actual characteristics and behavior of the underlying communication networks, computers, system software, applications, and users impose a myriad of distinct challenges for solving these problems, which have also been exten-

sively addressed in research in LATAM, ranging from theoretical foundations to system engineering and tools.

In the computability and complexity theory realm, handling faults, either by masking or recovering from them, brings additional challenges in algorithm design, especially when systems are distributed, not only in understanding computability limits such as the impossibility of solving certain fault-tolerant problems, but also to proposing, analyzing, and proving solutions for problems such as distributed consensus, renaming, and mutual-exclusion, to cite a few. Thus, a great deal of research in LATAM has been dedicated to the search for appropriate system models, data, and control structures to handle these basic fault-tolerant problems.

At the system and networking management level, other challenges arise. Today systems are widely distributed, concurrent, mobile, and often involve the composition of heterogeneous components, usually requiring autonomous coordination and monitoring. Several pieces of research in LATAM addressed these challenges, proposing new approaches to handle typical problems such as distributed diagnosis, message broadcast, and failure detection, among other problems. Furthermore, theoretical as well as practical approaches to modeling, reasoning, and implementing adaptive and self-adaptive dependable systems in such complex distributed scenarios have been proposed, aiming for autonomous properties such as self-optimization, self-organization, and self-management.

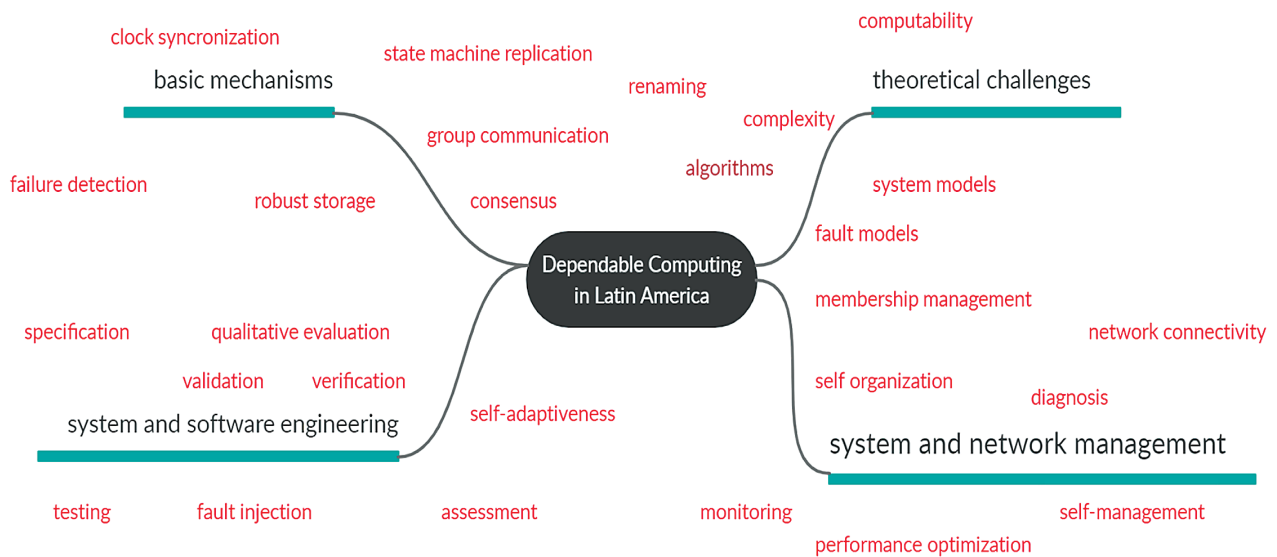
The overall challenges and main topics addressed in dependable computing research in LATAM are shown in the accompanying figure.

### **A Non-Exhaustive Look into Results**

Some of the early dependability projects developed at INPE included robustness testing of satellites, including altitude and orbit control, for which data mining was applied to detect anomalies. Space software was also a major subject, such as the field investigation of errors on space software requirements, done in cooperation with Emilia Villani from the nearby Institute of Aeronautical Tech-

<sup>b</sup> <http://www.inf.ufpr.br/elias/BibDepLA.html>

## Topics of dependable computing research in Latin America.



nology. INPE also fostered projects on integrating model checking and model-based testing for industrial software development. The Brazilian pioneer in this field is Eliane Martins (UNICAMP, Brazil), who has led several projects on dependable computing in cooperation with INPE, in particular with Ana Maria Ambrosio and Fátima Mattiello-Francisco. One of their projects on a Test Environment with Fault Injection by Software (ATIFS)<sup>c</sup> proposed the combination of conformance testing with fault injection. They also worked on the automation of specification-based test case generation for communication systems, and effective testing strategies for the interoperability and robustness of real-time embedded software. Combining model-driven engineering and model-based testing has been recently explored to cope with dynamically evolving systems, with the contribution of Leonardo Montecchi.

Earlier work on fault injection for dependability validation was developed by Martins with Jean Arlat (France). That work included the development of estimators of the coverage of fault-tolerant mechanisms computed using statistical analysis of data obtained with fault injection. Later, Martins also worked on fault injection strategies based

on reflective programming and on patterns. With Regina Moraes (UNICAMP), Henrique Madeira (Portugal), and Marco Vieira (Portugal), they proposed a strategy based on fault injection for risk assessment. A Java framework to specify fault-loads for fault injection campaigns was proposed by Taisy Weber and Sergio Cecchin (UFRGS, Brazil). They also developed the FIONA tool, a fault injector for dependability evaluation of Java-based network applications.

The development of dependable software based on exception handling mechanisms was pioneered by Cecilia Rubira (UNICAMP), who has worked on different aspects of this problem. A comparison of exception handling techniques for object-oriented software with Alessandro Garcia (PUC-Rio, Brazil), Alexander B. Romanovsky, and Jie Xu (U.K.) is highly cited. An architectural approach for effectively representing and analyzing fault-tolerant software systems has been proposed with Rogerio de Lemos (U.K.) that relies on exception handling to tolerate diverse types of faults. Rubira has worked with Andrea Bondavalli (Italy) on the dependability of dynamic software product lines (SPLs) and SPLs for supporting fault-tolerant composite services.

An evaluation of air traffic controller workloads considering manned and unmanned aircraft systems is

one of several projects on the safety of critical cyber physical systems by João Batista Camargo (USP, Brazil). Other projects include anomaly detection in railway and subway systems. As confidence in safety analysis is essential, Camargo has proposed practical analytical approaches to increase confidence in systems based on programmable logic devices, and on safety-critical software.

The TANGRAM-II tool for modeling system performance and availability, which has been widely used in both academia and industry worldwide, was developed by Edmundo de Souza e Silva (UFRJ, Brazil). Among his many contributions are strategies to compute availability and performability measures of repairable computer systems using randomization.

Armando Castañeda and Sergio Rajsbaum (UNAM, Mexico) are well-known for their work on the theoretical foundations of fault-tolerant distributed systems. Together with Achour Mostefaoui and Michel Raynal (France) they investigated the conditions that identify sets of input vectors for which it is possible to solve consensus despite the occurrence of up to  $t$  process crashes, extended later with Roy Friedman (Israel) to the relationship of asynchronous agreement with error-correcting code. In collaboration with Idit Keidar (Israel), Rajsbaum investigated the cost of solv-

<sup>c</sup> <http://www3.inpe.br/atifs/> (Portuguese)



## The development of dependable systems requires a combination of fault prevention and fault tolerance techniques to achieve the desired level of reliability, availability, and safety.

ing consensus in failure-free executions. A recent innovative result is on generalizing linearizability to interval-linearizability, allowing specifications of concurrent problems. Robotics has also been a topic of recent research, including an investigation of robots under the Asynchronous Luminous Robots model operating in look-compute-move rounds in connection with shared-memory wait-free algorithms. Together with Pierre Fraigniaud (France) and Corentin Travers (France) they initiated the study of runtime distributed monitors, which while monitoring the correctness of the underlying system, tolerated failures of the monitors themselves.

Castañeda in his Ph.D. work studied the renaming problem, in which processes start with unique input names from a large space and must choose unique output names taken from a smaller name space. This corrected a result of a paper that had won the 2004 Gödel Prize on the topological structure<sup>2</sup> of asynchronous computability and showed that the solvability of the renaming problem in asynchronous systems depends on whether the number of processes is a power of a prime number. Castañeda won the Best Student Paper Award at PODC 2008, and the paper was included in the ACM list of Notable Computing Books and Articles of 2012. The Babel file system is a software defined, massive, scalable and fault-tolerant distributed storage system developed by Ricardo Marcelín Jiménez (UAM, Mexico). Babel is a middleware for managing replicated databases, based on Paxos, and includes several innovations. Early research in Brazil on dependable distributed systems started perhaps with Rogerio Drummond (UNICAMP), who with Ozalp Babaoglu (then in the U.S.) published seminal work on clock synchronization, as well as the “Streets of Byzantium” paper on implementing reliable broadcasts in distributed systems with broadcast networks.

It was a Brazilian researcher, Joni Fraga (UFSC, Brazil), who coined the term “intrusion tolerance” in 1985. The term has been adopted worldwide and became truly popular years later with the growth of the Internet and related security problems. In coopera-

tion with his previous Ph.D. students Alysson Bessani (Portugal), Eduardo Alchieri (UnB, Brazil), and the late Lau Cheuk-Lung, they proposed relevant algorithms and tools for distributed systems that tolerate Byzantine faults. DepSpace is a system to improve the dependability of tuple spaces. Connectivity requirements for solving Byzantine consensus with unknown participants was studied with Fabiola Greve (UFBA, Brazil).

The well-known BFT-Smart system for state machine replication was proposed by Alchieri together with Bessani. With Fernando Pedone (Switzerland) they have investigated strategies to boost the concurrency of parallel state machine replication.

Raimundo Macêdo (UFBA, Brazil) proposed the concept of causal blocks to represent group message ordering. Based on this concept, with Paul Ezhilchelvan (England) and Santosh Shrivastava (England), they proposed the early and well-cited fault-tolerant, general-purpose Newtop protocol for partitionable and overlapping process groups, as well as pioneering mechanisms for flow control in group communication. The investigation of mobile process groups with virtual synchrony was an innovative contribution with Flavio Assis-Silva (Brazil). Distinct prominent aspects of consensus have been investigated: such as mobile systems with Michel Hurfin (France) and Nadjib Badache (Algeria); general agreement framework with Hurfin, Raynal, and Frederic Tronel (France); and adaptive message patterns with Hurfin, Mostefaoui, and Raynal. Macêdo addressed adaptive failure detection in asynchronous systems and the use of neural networks and Simple Network Management Protocol for adaptive detectors with his supervised student Fabio Ramon (Brazil).

Macêdo proposed alternative hybrid distributed system models that welded both synchronous and asynchronous assumptions under the same framework, initially with his Ph.D. student Sérgio Gorender (Brazil), and later with Raynal and Gorender. Group communication and simulation tools in this hybrid model have been developed in the Ph.D. thesis of Allan Freitas (Brazil). Macêdo also proposed the partitioned-syn-

chronous model with Gorender, and some problems have been addressed under this model like failure detection, mutual-exclusion, and with his supervised students Marcos Ramos, Anne Blagojevic, and Wellington Silva, Byzantine failures.

Self-manageable distributed system protocols inspired by the feedback control theory, where protocol's objectives can be modified and controlled at runtime, were also innovative contributions by Macêdo and his supervised students. The Ph.D. theses of Alirio Sá (Brazil), Freitas, and Sandro Andrade (Brazil) have applied these principles to QoS-based self-configuring failure detectors, self-manageable group communication, adaptive Byzantine replication, and self-adaptive software architecture design. Finally, an initial effort to secure IoT-based cyber-physical human systems against collaborative attacks has been undertaken with Sathish Kumar, Bharat Bhargava, and Ganapathy Mani (U.S.).

Multiple other groups have worked on diverse aspects of dependable distributed systems. Improving the precision of failure detectors using time series to predict communication delays was proposed by Ingrid Jansch-Porto (UFRGS, Brazil) with Raul Ceretta (UFSM, Brazil). Ceretta has led multiple projects on security and has had a partnership with the Universidad de Paraguay for more than 10 years. The Impact Failure Detector, which takes into account both process relevance and confidence in the system to assess the state of monitored processes, was proposed by Cláudio Geyer (UFRGS, Brazil) with Luciana Arantes (France) and Anubis Rossetto (IFSUL, Brazil). In cooperation with Pedone, Fernando Dotti (PUCRS, Brazil) has worked on several aspects of parallel state machine replication, also in cooperation with Odorico Mendizabal (UFSC). Among the several relevant results of this fruitful cooperation is the Byzantine fault-tolerant atomic multicast strategy proposed with Bessani. The dependability of streaming systems has been investigated by Andrey Britto (UFCG), with Christof Fetzer (Germany). Britto's main focus is on security. Total order broadcast and consensus have been investigated by Luiz Buzato (UNICAMP). Together with Islene

Garcia (UNICAMP) they have proposed relevant strategies for checkpoint and rollback.


Dynamic distributed systems with unknown participants have been investigated by Greve, including failure detection and eventual leader election in evolving mobile networks in cooperation with Arantes. A consensus algorithm for systems with unknown participants in shared memory was developed with Catia Khouri (IFBA, Brazil) and Sébastien Tixeuil (France). Also, with Tixeuil, Greve has investigated the knowledge connectivity versus synchrony requirements for consensus in unknown networks. A solution to the group priority inversion problem in the timed asynchronous model was proposed by Greve in cooperation with Francisco "Fubica" Brasileiro (UFCG, Brazil), Emmanuelle Anceaume (France), and Hurfin. Earlier, Greve and Brasileiro, with Mostefaoui and Raynal, in a seminal work proposed consensus in one communication step. Brasileiro, with Livia Sampaio (UFCG), proposed an adaptive process ordering module to improve the performance of adaptive indulgent consensus protocols. The implementation of fail-silent nodes for distributed systems was an early work of Brasileiro with Neil Spears (U.K.).

Working on the frontier of distributed systems and networking, Elias P. Duarte Jr. (UFPR, Brazil) has proposed with Luis Bona (UFPR) the VCube virtual topology for distributed systems. VCube is a hypercube when all processes are correct, but as processes fail and recover the structure reorganizes itself, keeping several logarithmic properties. The topology was first introduced in the context of distributed diagnosis with Takashi Nanya (Japan). Multiple distributed algorithms have been proposed for the VCube, including reliable broadcast and distributed mutual exclusion with Luiz A. Rodrigues (UNIOESTE, Brazil) and Arantes, and a publish-subscribe algorithm, with Pierre Sens.

Duarte has also worked on the diagnosis of dynamic partitionable general topology networks and comparison-based diagnosis. Results include a survey covering 30 years of research in this field, as well as a nearly optimal algorithm for general topologies. The search for a fault-tolerant

routing strategy for the Internet led to the development of connectivity numbers, which are centrality metrics that reflect network node connectivity, proposed with Jaime Cohen (UEPG, Brazil). The two also proposed parallel algorithms for cut trees, a very relevant combinatorial data structure. Recent work includes a strategy to improve the dependability of cloud systems based on replicating the cloud manager itself. With Rogerio Turchetti (UFSM) and Edson Camargo (UTFPR), they have investigated the dependability of programmable/virtualized networks, including the usage of network function virtualization technologies to implement in-network distributed services, including failure detectors, reliable and ordered broadcast, and consensus.

## Conclusion

To conclude, we must remark that the article is far from exhaustive, not only in terms of results of the research groups mentioned, but also because there are many other groups that could not be discussed due to space restrictions. In particular, security is a dependability attribute and there is a very large number of groups working in that field. Most of the research groups presented here are at universities that are forming a good number of young researchers who are very enthusiastic about dependable computing research. The future looks bright! 

## References

1. Avizienis, A. et al. Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable and Secure Computing* 1.1 (2004), 11–33.
2. Herlihy, M., Kozlov D.N. and Rajsbaum S. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, 2013.
3. Pradhan, D.K. (Ed.). *Fault-Tolerant Computer System Design*. Prentice-Hall, 1996.
4. von Neumann, J. Probabilistic logics and synthesis of reliable organisms from unreliable components. *Automata Studies*. C. Shannon, J. McCarthy, J. (Eds.). Princeton University Press, Princeton, NJ, 1956, 43–98.

**Elias P. Duarte Jr.** is a professor in the Department of Informatics at Federal University of Paraná, Curitiba, Brazil.

**Raimundo J. A. Macêdo** is a professor in the Computer Science Department at Federal University of Bahia, Salvador, Brazil.

**Eliane Martins** is a collaborating professor in the Institute of Computing at the State University of Campinas, Campinas, Brazil.

**Sergio Rajsbaum** is a professor at the Instituto de Matemáticas at the Universidad Nacional Autónoma de México, in Mexico City, México.