

VRDADER: Verification, Performance and Availability of Distributed Applications in Real-world Scenarios

TIC2002-02859

Víctor M. Gulías *
MADS Group - LFCIA
Department of Computer Science
University of Corunna (SPAIN)

Abstract

Nowadays, the development of reliable and efficient distributed systems is a must. In the last years, our group has developed and deployed different solutions to real-world problems using off-the-shelf computer clusters managed by control systems implemented using the functional programming paradigm and, in particular, Erlang/OTP.

The goal of the project is to use the experience in developing distributed functional applications to extract common design/programming functional patterns that can be valuable for new projects. Using these patterns and applying formal techniques such as model checking or theorem proving, distributed functional applications can be improved in terms of reliability, performance and availability.

Among the applications considered, we have used a B2C e-commerce framework, a payment gateway, a distributed filesystem, a risk management information system and a distributed video-on-demand server (VoDKA). One of the refactored systems, the VoDKA server, has caused great impact and it has left the academic environment to become a commercial product, transferring the knowledge gathered by this project to the society.

Keywords: functional programming, distributed computing, cluster computing, design patterns, model checking, theorem provers, performance evaluation.

1 Introduction and Goals

1.1 Project Overview

Nowadays, the development of reliable and efficient distributed systems is a must. For performance reasons, and to be cost-effective, commodity off-the-shelf computer clusters have been suggested a highly flexible and scalable architecture for some problems. As high speed networks and processor have become commodity hardware, affordable and reasonably efficient clusters

*Email: gulias@lfcia.org

are flourishing everywhere. However, the programming of such systems is still a difficult task. For that reason, we believe that a declarative style for the programming of the control system, such as the proposed by functional programming, plays an important role in the development. Moreover, the combined use of design patterns and distributed functional programming has been pointed out as a key factor to quickly produce correct distributed systems running on a cluster of computers with a high degree of adaptability, fault tolerance and scalability.

In the last years, our group has developed and deployed different solutions to real-world problems using off-the-shelf computer clusters managed by control systems developed using the functional programming paradigm and, in particular, Erlang/OTP [1]. The overall goal of the VRDADER project is to use the experience in developing distributed functional applications to extract common design/programming idioms, that is *distributed functional patterns*, that can be valuable for solving recurrent problems in future projects. Using these patterns and applying formal techniques such as model checking or theorem proving, distributed functional applications can be improved in terms of reliability, performance and availability.

At the beginning of the project, the applications considered were a B2C e-commerce framework, a payment gateway and, mainly, a distributed video-on-demand server (VoDKA, *Video-on-Demand Kernel Architecture*, <http://vodka.lfcia.org>). The goal was then to extract recurrent solutions and abstract them out in order to reuse them at different points. Techniques such as model checking or theorem proving can be applied on some of these patterns, improving the quality of the building blocks. Once those patterns have been identified, a refactoring of VoDKA server is proposed to improve the quality of the system.

The resulting great impact of the VoDKA server, due to its reliability and flexibility, has been important to deploy the system in 24x7 real scenarios (cable users, train stations, etc.). This success has motivated the development of more systems using the same techniques such as a distributed file system and the backend of a risk management information system.

1.2 Objectives

To summarize, the concrete goals of the project are:

- To analyze distributed functional applications looking for potential recurrent problems.
- To abstract the solution of each potential recurrent problem.
- To apply different formal techniques on some of the abstractions, in particular, model checking and theorem proving.
- To refactor the VoDKA server to improve its quality in terms of reliability, scalability, flexibility and availability.
- To apply the abstractions to existent and new distributed functional applications.
- To build a catalog of the identified distributed design patterns.

One important point of our proposal is the emphasis that we put on technological transfer to society. We think that distributed functional computing and all the related techniques can be useful for industry and one of our tasks is to demonstrate it. For that reason, we included in our proposal several partners from industry and we should take the appropriate actions to make possible such transfer.

2 Success Level

2.1 Achievements

The achievements so far in the project, according to the the goals presented in section 1.2 are:

- Three different distributed applications have been studied: a video-on-demand server (VoDKA), a clustered payment gateway and an e-commerce framework (SCED). Moreover, two additional applications have been also considered: a distributed filesystem (CheapTB) and the backend of an advanced risk management system (Armistice). There are many recurrent problems found in the case studies.
- Some of the problems found are abstracted as distributed design patterns: generic servers, guarded suspensions, schedulers, data movement, event-based notifications and so on.
- Formal techniques have been applied to some of the problems factored as patterns. For example, model checking has been used to establish global properties for schedulers (for instance, VoDKA scheduler) while a theorem prover (Coq [2]) has been used to prove the correctness of an allocator pattern (for instance, the block allocator of VoDKA's cache).
- VoDKA server has been almost completely rewritten focusing on the identified design patterns. The quality of the system has improved notably in terms of reliability.
- Besides the application to VoDKA server, the patterns have been applied to new system such as the distributed filesystem (CheapTB) or the risk management information system (Armistice).
- A collection of patterns is being gathered. Some of the patterns have been published in conferences and journeys. The whole catalog is being published in the web by the end of the project (<http://vrdader.lfcia.org>).

2.2 Difficulties and Workarounds

In 2003, there was a problem originated by the payment delay (almost one year) of the ERDF part (70% of the total). This problem has been fixed in 2004 and now the project can be executed normally. To avoid this problem, many actions have been accomplished using additional funding for the project (Xunta de Galicia's incentive PGIDIT03PXIC10502PN), related projects (Xunta de Galicia's PGIDT02TIC00101CT), and private funding for developing or exploiting direct or indirect results of the project.

3 Result Indicators

3.1 Formation Activities

During the project, two Ph.D. thesis have been finished:

- Enrique Freire. *Higher-order functions in functional programming. A categorical perspective* (in Spanish). Directed by Jose L. Freire. University of Corunna, 2003.

Even though this thesis was in its final stage at the beginning of the project (and it should not be considered as a result of it), the experience using the theorem prover Coq has been valuable for the project. In its last chapter, Freire comments the problems of using theorem provers, such as Coq, to verify properties of concurrent systems and it provides some useful hints.

- Santiago Jorge. *On the verification of properties of functional programs: from manual proofs to theorem provers* (in Spanish) Directed by Jose L. Freire and Victor M. Gulias. University of Corunna, 2004.

In this thesis, S. Jorge presents how to verify properties of software developed using the functional paradigm. Starting from manual proofs of simple functional programs, the theorem prover Coq is introduced as a tool to assist the developer in the process of checking programs. The technique is applied to two different domains that, however, share the same pattern of behaviour: the buchberger algorithm, in particular all the polynomial operations, and the block allocation mechanism of VoDKA's local caches.

An additional thesis is expected to be finished in 2005: Juan J. Sanchez, former FPU recipient and, recently, teaching assistant at University of Corunna. He works on applying model checking techniques to relevant patterns of VoDKA and, in particular, to the scheduler of the video-on-demand server. In the last two years, he visited twice the IT University of Gothenburg during six months visiting Professor Thomas Arts (external collaborator of the project). The thesis is directed by Victor M. Gulias and Thomas Arts.

There are two additional people working on its Ph.D. thesis (estimate: late 2006) related with the project. Currently, they have predoctoral research grants funded by Xunta de Galicia: Carlos Abalde (joined the project in 2003, now visiting Bristol) and Javier Paris (contracted at the beginning of the project; now applying for joining the project as regular member from 1/10/2004). Both are directed by Victor M. Gulias.

Two additional software engineers will join the project contracted by VRDADER in 2005: Carlos Varela and Laura Castro. They will finish their DEA projects in 2005.

Moreover, many master thesis and DEA projects related to VRDADER were presented in the last two years. They represent preliminar work that we expect to consolidate in the following months.

3.2 International Publications

To summarize, the international publications (up to now) are:

1	long paper (25 pages) in the Journal of Functional Programming
2	papers in conferences with publication in IEEE Computer Society Press
1	paper in IFIP/ACM conference
3	papers in ACM conference and workshops
3	papers in conferences with proceedings in Lecture Notes in Computer Science (2 are selections of papers and 1 has a very low acceptance ratio)
2	papers in conference sponsored by IFIP/IEEE Computer Society
1	paper in conference sponsored by Information Processing Society of Japan
1	paper in conference sponsored by IASTED

- V. Gulias, M. Barreiro, J. Freire. *VoDKA: Developing a Video-on-Demand Server using Distributed Functional Programming*. Journal of Functional Programming, special issue on Functional Approaches to High-Performance Parallel Programming (to appear; accepted in August 2004)

This is one of the key objectives pursued by our project. In this 25-page paper, we present a general overview of the video-on-demand server VoDKA, the main case-study of the project, putting emphasis on several distributed design patterns identified (generic server, guarded suspension, scheduler, pipes, and process composition). A performance study suggests that the technology satisfies the requirements of the real scenario for which it is intended and it simplifies the development notably.

- S. Neira, V. Gulias. *An Architecture for Location Dependent Web Based Applications on Mobile Environments*. 2nd International Symposium on Ubiquitous Computing Systems UCS'2004. Organized by the Ubiquitous Computing System Group of Information Processing Society of Japan (IPSJ). Tokio, Japan. November 2004.

An improved and extended version of CSN2004 work. In addition to the architectural pattern to develop location-dependent web-based applications, we propose a real-world scenario for the pattern, a virtual museum, which incorporates video-on-demand services using VoDKA.

- C. Varela, V. Gulias, A. Valderruten, C. Abalde. *Network Management System for a Video-on-Demand Server*. XXX Conferencia Latinoamericana de Informática CLEI'2004. Arequipa (Peru). September 2004 (in Spanish).

In this paper, we introduce VoDKA network management system. This is defined using a variation of the observer pattern to capture events and notifications from different components. In this case, an special agent (monitor) is deployed at each node; this agent is in charge of receiving asynchronous notifications from the rest of concurrent processes within the same node and distributed them to several observer processes registered in the monitor.

Even though the talks are mainly in Spanish and Portuguese in this conference, it has a great tradition and prestige in latinamerica and there is a rigorous selection process (acceptance ratio below 40%). It is sponsored, among others, by IFIP and IEEE Computer Society.

- J. Gulias, V. Gulias, A. Valderruten, C. Abalde. *The Multi-View Pattern for the Generation of Multiple Presentations in an e-Commerce System*. XXX Conferencia Latinoamericana de Informática CLEI'2004. Arequipa (Peru). September 2004 (in Spanish).

In this paper, a design pattern identified in the SCED e-commerce system (one of the VRDADER case studies) is shown. This multiview pattern is an abstraction that makes independent a server from the specific presentations issues related with the client context or its capabilities.

- C. Varela, C. Abalde, L. Castro, J. Gulias. *On Modelling Agent Systems with Erlang*. International Conference of Functional Programming (ICFP'04), 3rd ACM SIGPLAN Erlang Workshop. Snowbird, USA, September 2004.

In this paper, we explore Erlang as a possible implementation language for defining agent systems following the well-known BDI model (beliefs-desires-intentions). In this sense, Erlang can be used as an alternative to other agent languages such as KQML or FIPA, or as a base to translate those languages into an Erlang implementation. Besides the basic architecture, we establish some fundamental mechanisms such as cooperation, reconfiguration and fault-tolerance.

- S. Neira, V. Gulias. *A Transparent Client-Server Platform for Location Dependent Web Based Applications*. The IASTED International Conference on Communication Systems and Networks CSN'2004. Marbella, Spain. September 2004.

In this paper, a new case study of a complex distributed application is presented: a generic framework for location-dependent web-based development. This architectural pattern proposes a proxy in the client that uses a bridge pattern to integrate heterogeneous location devices, such as Bluetooth, IrDA, 802.11, GPS or RFID beacons, and it includes location information in the HTTP headers in order to parametrize server responses.

- T. Arts, C. Benac, J. Sanchez. *Translating Erlang to mCRL - Make Industrial Code Available for Research Tools*. 4th International Conference on Application of Concurrency to System Design (ACSD'2004). IEEE CS Press, Canada. June 2004

In this paper, we present a method and a set of tools oriented towards the formal verification of properties of concurrent programs developed in Erlang. Taking as working examples a model of the VoDKA scheduler and a model of the resource manager of Ericsson AXD301 switch, we propose a transformation from Erlang code to a process algebra (mCRL) and then we use the tools and formal methods available for this process algebra in order to prove program properties.

- V. Gulias, A. Valderruten, C. Abalde. *Building Functional Patterns for Implementing Distributed Applications*. IFIP/ACM LANC'2003. La Paz, Bolivia, October 2003

In this paper, three distributed design patterns, identified in VRDADER case studies, are presented and commented. The patterns are structured similarly to other pattern collections (for example, GoF patterns [3]). These patterns are proposed as building blocks for developing large distributed applications. An implementation skeleton is also shown, putting emphasis on distributed functional programming (Erlang, as an example) as an interesting development technique for such systems.

- J. Sanchez, C. Abalde. *Extending the VoDKA Architecture to Improve Resource Modeling*. Principles, Logics and Implementation of High-Level Programming Languages (PLI'03), 2nd ACM SIGPLAN Erlang Workshop. Uppsala, Sweden, August 2003

In order to model resource control, VoDKA defines a delegation scheme based on the chain of responsibility pattern. This solution has become inappropriate for some deployments that need a finer grain of control at each node. Besides this pattern, the paper shows a refinement of it designed to overcome this problem. The refactoring of VoDKA is inspired on the modeling of the negotiation process using Tomas Arts' ETOMCRL tool.

- D. Cabrero, C. Abalde, C. Varela, L. Castro. *ARMISTICE: An Experience Developing Management Software with Erlang* Principles, Logics and Implementation of High-Level Programming Languages (PLI'03), 2nd ACM SIGPLAN Erlang Workshop. Uppsala, Sweden, August 2003

This paper summarizes the experience developing a risk management information system using Erlang. The project, with 100% private funding from one of VRDADER industrial partners *Alfa21 Outsourcing*, has been motivated by the success of VoDKA development. The server side of the system has been developed using the same techniques and patterns used in VoDKA and other case studies. In addition, it provides new potential patterns to be exploited in the future.

- J. Sanchez, T. Arts. *VoDKAV Tool: Model Cheking for Extracting Global Scheduler Properties from Local Restrictions*. 3rd International Conference on Application of Concurrency to System Design (ACSD2003). IEEE CS Press, Portugal. June 2003

In this paper, some results derived from the use of model checking techniques to an abstraction of VoDKA scheduler is shown. Using a customized version of the ETOMCRL tool, a simple abstraction of the source code of the system is first translated into an intermediate process algebra and, later, the state space of a given configuration is generated. From this state space, some global properties of the system can be extracted.

- C. Abalde, V. Gulias, J. Freire, J. Sanchez, J. Garcia-Tizon. *Development of a Scalable, Fault Tolerant and Low Cost Cluster-based e-Payment System with a Distributed Functional Kernel*. 9th International Conference on Computer Aided System EUROCAST'2003. LNCS 2809, Las Palmas de Gran Canaria, February 2003

This work summarizes the experience developing a clustered payment gateway system. Special attention is paid to the coordination behaviour shown by system actors. Some design patterns, related with availability and scalability issues, are sketched. This is one of the case studies of the VRDADER project.

- J. Freire, A. Blanco, V. Gulias, E. Freire. *The Reduction Lemma in Coq*. 9th International Conference on Computer Aided System EUROCAST'2003. LNCS 2809, Las Palmas de Gran Canaria, February 2003

This paper presents the verification of the reduction lemma using the theorem prover Coq. It defines the methodology to follow when using this provers. One of our goals at this point is to use theorem provers to verify properties of relevant components of the system.

- V. Gulias, C. Abalde, J. Sanchez. *Lambda Goes to Hollywood*. 5th International Symposium on Practical Aspects of Declarative Languages PADL'03. LNCS 2562. New Orleans, USA Enero 2003

Presentation of VoDKA as a succesful example of a declarative program addressing a demanding real-world problem. This is the main case study of the VRDADER project. The paper sketches some of the abstractions used in the development of VoDKA and how that can be implemented using higher-order functions.

3.3 Technological Transfer to Society

A key aspect of the project is to make possible that the explored research can be useful to industry. The main collaborations with the companies that supported our proposal has been:

- **R Cable y Comunicaciones de Galicia S.A.**. It has started a commercial deployment of VoDKA. During 2003 and 2004, the server was in use by selected final users for testing.
- **Ericsson AB**. Invitation (twice) of Bjarne Dacker (Ericsson contact person) to Victor M. Gulias to be a member of ACM Erlang Workshop program commitee. VoDKA experience is well-known in Ericsson. The final catalog of patterns is quite interesting for people in charge of OTP (Erlang's Open Telecom Platform).

- **Alfa21 Outsourcing S.L.**. Commitment to commercialize products developed using distributed functional technology. First success stories deploying VoDKA server at different locations: pilot deployment for R Cable y Comunicaciones de Galicia, integration with the AVE (Spanish high-speed train) information system in the Guadalajara-Zaragoza line and pilot experience in CaixaGalicia's (large bank in Galicia) OficinaON (an office with about 30 video-information points). In 2003, the company has started a new project with our group (100% private funding) to develop a risk management information system (ARMISTICE) for a Galician multinational company using VRDADER approach.

To make possible the technological transfer, two additional actions have been carried out:

- A **spinoff company**, *LambdaStream Servicios Interactivos S.L.*, has been created at the end of 2003 to industrialize VoDKA-related products and technology. The company is devoted entirely to R&D related with reliable services and, in particular, digital video. The spinoff uses larger partner companies such as Alfa21 (R, CaixaGalicia) or Indra S.A. (GIF, Amena) to commercialize its products. Currently, four former researchers of our group are now contracted in the company, working in new projects that apply their acquired knowledge. LOU contracts with the laboratory maintain the flow of information in both directions and they become an additional source of fundings for our research.
- The interest of the techniques explored has originated a **patent** request: *Method and system of distributed storage for multimedia contents oriented toward streaming* (in Spanish) request number P200302837 (3/12/2003). There is a firm offering of exploiting the patent.

3.4 Collaboration with other Groups

In 2003, as a result of the start of VRDADER and due to the dissemination conducted primarily by our external collaborators T. Arts and L. Fredlund, a consortium was created to prepare an European project proposal (*Formally-based tool support for Erlang development*). The project consortium consists of the major internationally recognised Erlang researchers in Europe. Within the academic partners, we found University of Kent (J. Derrick), Brunel University (M. Harrman), Chalmers University (J. Hughs), **University of Corunna (V. Gullias)**, **IT University in Gothenburg (T. Arts)** and **SICS (L. Fredlund)**. The consortium also includes companies such as Cesarini Consulting, Ericsson, Mobile Arts and **R Cable y Comunicaciones de Galicia**. The goal of the project, as in VRDADER, is summarized as *making reliable products cheaper and cheap products more reliable*.

The project was presented to the VI Framework Programme, as a Specific Targeted Research Project (STRP). Unfortunately, the first attempt was not accepted by the commission. Nevertheless, the consortium is working to refine the proposal for future calls.

References

- [1] J.Armstrong, M.Williams, R.Virding. *Concurrent Programming in Erlang*. Prentice-Hall, 1993.
- [2] G.Huet, G.Kahn, C.Paulin-Mohring. *The Coq Proof Assistant, a tutorial*. INRIA 2004.
- [3] E.Gamma, R.Helm, R.Johnson, J.Vlissides. *Design Patterns*. Addison-Wesley, 1996.