INRIA

Project-Team phoenix

Programming Language Technology For Communication Services

Bordeaux - Sud-Ouest

Theme : Distributed Systems and Services

Activity Report

2010

Table of contents

The Phoenix group is located at the INRIA Bordeaux-Sud Ouest center. Phoenix is an INRIA Project-Team joint with University of Bordeaux and CNRS (LaBRI, UMR 5800).

1. Team

Research Scientist

Emilie Balland [INRIA Junior Researcher]

Faculty Member

Charles Consel [Team Leader, Professor, ENSEIRB, HdR]

External Collaborator

Julia Lawall [Associate Professor, University of Copenhagen (DIKU)]

Technical Staff

Benjamin Bertran [Associate Engineer, from December 17, 2007]
Ghislain Deffrasnes [Associate Engineer, from October 1, 2009]

PhD Students

Julien Mercadal [Ministerial scholarship, from October 1, 2006 to April 30 2010, University of Bordeaux and INRIA Research Assistant from May 1, 2010 to October 31, 2010]
Zoé Drey [Research Assistant, from November 1, 2009 to April 31, 2010, INRIA and from May 1, 2010 to July 31, 2010, University of Bordeaux]
Damien Cassou [Ministerial scholarship, from October 1, 2007 to September 30, 2010, University of Bordeaux and Research Assistant, from October 1, 2010, University of Bordeaux]
Julien Bruneau [Thales scholarship, from October 1, 2008]
Henner Jakob [INRIA scholarship, from May 1, 2008]
Hongyu Guan [Region scholarship, from February 9, 2009]
Pengfei Liu [INRIA scholarship, from October 1, 2009]
Stéphanie Gatti [Thales scholarship, from February 1, 2010]
Quentin Enard [Thales scholarship, from February 1, 2010]

Post-Doctoral Fellow

Nicolas Loriant [INRIA scholarship, from October 1, 2009, to September 30, 2010]

Administrative Assistant

Sylvie Embolla [Group Assistant, from September 4, 2006]

2. Overall Objectives

2.1. Overall Objectives

A host of networked devices are populating smart spaces that become prevalent in an increasing number of areas, including supply chain management (e.g., parcel tracking), monitoring (e.g., building surveillance and patient monitoring) and home and building automation (e.g., control of energy consumption). This situation raises a number of challenges (1) safety and security because of the interweaving of these smart spaces in our daily life, (2) productivity because of a high demand of applications matching the wide range of user needs, and (3) abstraction because of the combination of expertise areas involved in smart spaces.

To address these challenges, we develop a software engineering approach that is dedicated to services orchestrating networked devices:

- the specification of robust orchestrating services based on innovative Domain-Specific Languages (DSLs),
- the study of the communication layers underlying these services to improve flexibility and performance,
- the application to concrete areas such as pervasive computing or IP telephony to validate our approach.

2.2. Highlights

1. **DiaSuite release** in October 2010: this release has been the occasion for various demonstrations:
   – during scientific events: the Eigth Annual IEEE International Conference on Pervasive Computing and Communications [14] (PerCom 2010) and the 25th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications [15] (OOPSLA 2010);
   – during popular science events: laying of the foundation stone for the INRIA Bordeaux buildings in September and "fête de la science" in October;
   – during industrial events: INRIA industrial meetings in the topic of "industries du numérique pour la santé" in April, kickoff meetings with Bouygues Telecom in June and September.

2. Initiation of **technology transfer with Bouygues Telecom** in December 2010: this transfer will consist of a porting of DiaSuite on the Bouygues set-top box for domotic applications.

3. Our work in design-driven development received this year a significant international visibility with the **publication in leading conferences** in the domain of object-oriented programming languages and software engineering:
   – Conference on Object-Oriented Programming Systems, Languages, and Applications [18] (OOPSLA 2010),
   – International Conference on Software Engineering [16] (ICSE 2011).

4. **Organization of the workshop in Applications of Cognitive Assistance for Persons with Intellectual Disabilities** in September 2010: this workshop gives concrete expression to a new application area for our research group, opening up a host of research opportunities and collaborations.

3. Scientific Foundations

3.1. Introduction

Our proposed project builds upon results previously obtained by the Compose research group whose aim was to study new approaches to developing adaptable software components in the domain of systems and networking. In this section, we review the accomplishments of Compose, only considering the ones achieved by the current project members, to demonstrate our expertise in the key areas underlying our project, namely:

- Programming language technology: language design and implementation, domain-specific languages, program analysis and program transformation.
- Operating Systems and Networking: design, implementation and optimization.
- Software engineering: software architecture, methodologies, techniques and tools.

By combining expertise in these areas, the research work of the Compose group contributed to demonstrating the usefulness of adaptation methodologies, such as domain-specific languages, and the effectiveness of adaptation tools, such as program specializers. Our work aimed to show how adaptation methodologies and tools could be integrated into the development process of real-size software components. This contribution relied on advances in methodologies to develop adaptable programs, and techniques and tools to adapt these programs to specific usage contexts.

3.2. Adaptation Methodologies

Although industry has long recognized the need to develop adaptable programs, methodologies to develop them are still at the research stage. We have presented preliminary results in this area with a detailed study of the applicability of program specialization to various software architectures [33]. Our latest contributions in this area span from a revolutionary approach based on the definition of programming languages, dedicated to a specific problem family, to a direct exploitation of specialization opportunities generated by a conventional programming methodology.

3.2.1. Domain-Specific languages

DSLs represent a promising approach to modeling a problem family. Yet, this approach currently suffers from the lack of methodology to design and implement DSLs. To address this basic need, we have introduced the Sprint methodology for DSL development [25]. This methodology bridges the gap between semantics-based approaches to developing general-purpose languages and software engineering. Sprint is a complete software development process starting from the identification of the need for a DSL to its efficient implementation. It uses the denotational framework to formalize the basic components of a DSL. The semantic definition is structured so as to stage design decisions and to smoothly integrate implementation concerns.

3.2.2. Declaring adaptation

A less drastic strategy to developing efficient adaptable programs consists of making specific issues of adaptation explicit via a declarative approach. To do so, we enrich Java classes with declarations, named *adaptation classes*, aimed to express adaptive behaviors [22]. As such, this approach allows the programmer to separate the concerns between the basic features of the application and its adaptation aspects. A dedicated compiler automatically generates Java code that implements the adaptive features.

3.2.3. Declaring specialization

When developing components, programmers often hesitate to make them highly generic and configurable. Indeed, genericity and configurability systematically introduce overheads in the resulting component. However, the causes of these overheads are usually well-known by the programmers and their removal could often be automated, if only they could be declared to guide an optimizing tool. The Compose group has worked towards solving this problem.

We introduced a declaration language which enables a component developer to express the configurability of a component. The declarations consist of a collection of specialization scenarios that precisely identify what program constructs are of interest for specialization. The scenarios of a component do not clutter the component code; they are defined aside in a *specialization module* [28], [29], [27], [30].

This work was done in the context of C and declarations were intended to drive our C specializer.

3.2.4. Specializing design patterns

A natural approach to systematically applying program specialization is to exploit opportunities offered by a programming methodology. We have studied a development methodology for object-oriented languages, called design patterns. Design patterns encapsulate knowledge about the design and implementation of highly adaptable software. However, adaptability is obtained at the expense of overheads introduced in the finished program. These overheads can be identified for each design pattern. Our work consisted in using knowledge derived from design patterns to eliminate these overheads in a systematic way. To do so, we analyzed the specialization opportunities provided by specific uses of design patterns, and determined how to eliminate these overheads using program specialization. These opportunities were documented in declarations, called specialization patterns, and were associated with specific design patterns [41]. The specialization of a program composed of design patterns was then driven by the corresponding declarations. This work was presented in the context of Java and uses our Java specializer [40].

3.2.5. Specializing software architectures

The sources of inefficiency in software architectures can be identified in the data and control integration of components, because flexibility is present not only at the design level but also in the implementation. We proposed the use of program specialization in software engineering as a systematic way to improve performance and, in some cases, to reduce program size. We studied several representative, flexible mechanisms found in software architectures: selective broadcast, pattern matching, interpreters, layers and generic libraries. We showed how program specialization can be applied systematically to optimize these mechanisms [32], [33].

3.3. Adaptation in Systems Software

3.3.1. DSLs in Operating Systems

Integrating our adaptation methodologies and tools into the development process of real-size software systems was achieved by proposing a new development process. Specifically, we proposed a new approach to designing and structuring operating systems (OSes) [36]. This approach was based on DSLs and enables rapid development of robust OSes. Such an approach is critically needed in application domains, like appliances, where new products appear at a rapid pace and needs are unpredictable.

3.3.2. Devil - a DSL for device drivers

Our approach to developing systems software applied to the domain of device drivers. Indeed, peripheral devices come out at a frantic pace, and the development of drivers is very intricate and error prone. The Compose group developed a DSL, named Devil (DEvice Interface Language), to solve these problems; it was dedicated to the basic communication with a device. Devil allowed the programmer to easily map device documentation into a formal device description that can be verified and compiled into executable code.

From a software engineering viewpoint, Devil captures domain expertise and systematizes re-use because it offers suitable built-in abstractions [38]. A Devil description formally specifies the access mechanisms, the type and layout of data, as well as behavioral properties involved in operating the device. Once compiled, a Devil description implements an interface to an idealized device and abstracts the hardware intricacies.

From an operating systems viewpoint, Devil can be seen as an *interface definition language* for hardware functionalities. To validate the approach, Devil was put to practice [37]: its expressiveness was demonstrated by the wide variety of devices that have been specified in Devil. No loss in performance was found for the compiled Devil description compared to an equivalent C code.

From a dependable system viewpoint, Devil improves safety by enabling descriptions to be statically checked for consistency and generating stubs including additional run-time checks [39]. Mutation analysis were used to evaluate the improvement in driver robustness offered by Devil. Based on our experiments, Devil specifications were found up to 6 times less prone to errors than writing C code.

Devil was the continuation of a study of graphic display adaptors for a X11 server. We developed a DSL, called GAL (Graphics Adaptor Language), aimed to specify device drivers in this context [44]. Although covering a very restricted domain, this language was a very successful proof of concept.

3.4. Adaptation Tools and Techniques

To further the applicability of our approach, we have strengthened and extended adaptation tools and techniques. We have produced a detailed description of the key program analysis for imperative specialization, namely binding-time analysis [24]. This analysis is at the heart of our program specializer for C, named Tempo [24]. We have examined the importance of the accuracy of these analyses to successfully specialize existing programs. This study was conducted in the context of systems software [34].

Tempo is the only specializer which enables programs to be specialized both at compile time and run time. Yet, specialization is always performed in one stage. As a consequence, this process cannot be factorized even if specialization values become available at multiple stages. We present a realistic and flexible approach to achieving efficient incremental run-time specialization [31]. Rather than developing new techniques, our strategy for incremental run-time specialization reuses existing technology by iterating a specialization process. Our approach has been implemented in Tempo.

While program specialization encodes the result of early computations into a new program, *data specialization* encodes the result of early computations into data structures. Although aiming at the same goal, namely processing early computations, these two forms of specialization have always been studied separately. The Compose group has proposed an extension of Tempo to perform both program and data specialization [23]. We showed how these two strategies can be integrated in a single specializer. Most notably, having both strategies enabled us to assess their benefits, limitations and their combination on a variety of programs.

Interpreters and run-time compilers are increasingly used to cope with heterogeneous architectures, evolving programming languages, and dynamically loaded code. Although solving the same problem, these two strategies are very different. Interpreters are simple to implement but yield poor performance. Run-time compilation yields better performance, but is costly to implement. One approach to reconciling these two strategies is to develop interpreters for simplicity but to use specialization to achieve efficiency. Additionally, a specializer like Tempo can remove the interpretation overhead at compile time as well as at run time. We have conducted experiments to assess the benefits of applying specialization to interpreters [43]. These experiments have involved Bytecode and structured-language interpreters. Our experimental data showed that specialization of structured-language interpreters can yield performance comparable to that of the compiled code of an optimizing compiler.

Besides targeting C, we developed the first program specializer for an object-oriented language. This specializer, named JSpec, processes Java programs [40]. JSpec is constructed from existing tools. Java programs are translated into C using our Java compiler, named Harissa. Then, the resulting C programs are specialized using Tempo. The specialized C program is executed in the Harissa environment. JSpec has been used for various applications and has shown to produce significant speedups [42].

4. Application Domains

4.1. Introduction

After having explored DSLs in isolated domains in the past, we now generalize this experience to attack a larger domain, namely, communication services. Generalizing our work on telephony, we investigated the coordination of *networked entities*, whether or not operated by users. The two main application domains are the pervasive computing systems and the telephony services.

4.2. Pervasive Computing Systems

Pervasive computing systems are being deployed in a rapidly increasing number of areas, including building automation, assisted living, and supply chain management. Regardless of their target area, pervasive computing systems have a typical architectural pattern. They aggregate data from a variety of distributed sources, whether sensing devices or software components, analyze a context to make decisions, and carry out decisions by invoking a range of actuators. Because pervasive computing systems are standing at the crossroads of several domains (*e.g.,* distributed systems, multimedia, and embedded systems), they raise a number of challenges in software development:

- *Heterogeneity.* Pervasive computing systems are made of off-the-shelf entities, that is, hardware and software building blocks. These entities run on specific platforms, feature various interaction models, and provide non-standard interfaces. This heterogeneity tends to percolate in the application code, preventing its portability and reusability, and cluttering it with low-level details.

- *Lack of structuring.* Pervasive computing systems coordinate numerous, interrelated components. A lack of global structuring makes the development and evolution of such systems error-prone: component interactions may be invalid or missing.

- *Combination of technologies.* Pervasive computing systems involve a variety of technological issues, including device intricacies, complex APIs of distributed systems technologies and middleware-specific features. Coping with this range of issues results in code bloated with special cases to glue technologies together.

- Dynamicity. In a pervasive computing system, devices may either become available as they get deployed, or unavailable due to malfunction or network failure. Dealing with these issues explicitly in the implementation can quickly make the code cumbersome.

- Testing. Pervasive computing systems are complicated to test. Doing so requires equipments to be acquired, tested, configured and deployed. Furthermore, some scenarios cannot be tested because of the nature of the situations involved (*e.g.,* fire and smoke). As a result, the programmer must resort to writing specific code to achieve ad hoc testing.

4.3. Telephony Services

IP telephony materializes the convergence between telecommunications and computer networks. This convergence is dramatically changing the face of the telecommunications domain moving from proprietary, closed platforms to distributed systems based on network protocols. In particular, a telephony platform is based on a client-server model and consists of a *signalling server* that implements a particular signalling protocol (*e.g.,* the Session Initiation Protocol [21]). A signalling server is able to perform telephony-related operations that include resources accessible from the computer network, such as Web resources, databases...This evolution brings a host of new functionalities to the domain of telecommunications.

Such a wide spectrum of functionalities enables Telephony to be customized with respect to preferences, trends and expectations of ever-demanding users. These customizations critically rely on a proliferation of telephony services. In fact, introducing new telephony services is facilitated by the open nature of signalling servers, as shown by all kinds of servers in distributed systems. However, in the context of telecommunications, such evolutions should lead service programming to be done by non-expert programmers, as opposed to developers certified by telephony manufacturers. To make this evolution worse, the existing techniques to program server extensions (*e.g.,* Common Gateway Interface [20]) are rather low level, involves crosscutting expertises (*e.g.,* networking, distributed systems, and operating systems) and requires tedious session management. These shortcomings make the programming of telephony services an error-prone process, jeopardizing the robustness of a platform.

5. Software

5.1. DiaSuite: a Development Environment for Pervasive Computing Applications

Participants: Damien Cassou [correspondent], Charles Consel, Benjamin Bertran, Julien Bruneau, Julien Mercadal, Nicolas Loriant, Emilie Balland.

Despite much progress, developing a pervasive computing application remains a challenge because of a lack of conceptual frameworks and supporting tools. This challenge involves coping with heterogeneous devices, overcoming the intricacies of distributed systems technologies, working out an architecture for the application, encoding it in a program, writing specific code to test the application, and finally deploying it.

DIASUITE is a suite of tools covering the development life-cycle of a pervasive computing application:

- *Defining an application area.* First, an expert defines a catalog of entities, whether hardware or software, that are specific to a target area. These entities serve as building blocks to develop applications in this area. They are gathered in a taxonomy definition, written in the taxonomy layer of the DIASPEC language.

- *Architecturing an application.* Given a taxonomy, the architect can design and structure applications. To do so, the DIASPEC language provides an Architecture Description Language (ADL) layer [35]. This layer is dedicated to an architectural pattern commonly used in the pervasive computing domain [26]. Describing the architecture application allows to further model a pervasive computing system, making explicit its functional decomposition.

- *Implementing an application.* We leverage the taxonomy definition and the architecture description to provide dedicated support to both the entity and the application developers. This support takes the form of a Java programming framework, generated by the DIAGEN compiler. The generated programming framework precisely guides the developer with respect to the taxonomy definition and the architecture description. It consists of high-level operations to discover entities and interact with both entities and application components. In doing so, it abstracts away from the underlying distributed technologies, providing further separation of concerns.

- *Testing an application.* DIAGEN generates a simulation support to test pervasive computing applications before their actual deployment. An application is simulated in the DIASIM tool, without requiring any code modification. DIASIM provides an editor to define simulation scenarios and a 2D-renderer to monitor the simulated application. Furthermore, simulated and actual entities can be mixed. This hybrid simulation enables an application to migrate incrementally to an actual environment.

- *Deploying a system.* Finally, the system administrator deploys the pervasive computing system. To this end, a distributed systems technology is selected. We have developed a back-end that currently targets the following technologies: Web Services, RMI, CORBA and SIP. This targeting is transparent for the application code. The variety of these target technologies demonstrates that our development approach separates concerns into well-defined layers.

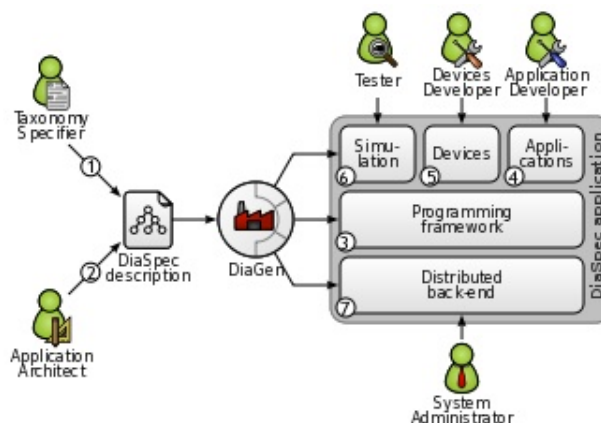This development cycle is summarized in the Figure 1.

Figure 1. DIASUITE *Development Cycle*

See also the web page http://diasuite.inria.fr.

5.1.1. DiaSpec: a Domain-Specific Language for Networked Entities

The core of the DIASUITE development environment is the domain specific language called DIASPEC and its compiler DIAGEN:

- DIASPEC is composed of two layers:

    – The *Taxonomy Layer* allows the declaration of entities that are relevant to the target application area. An entity consists of sensing capabilities, producing data, and actuating capabilities, providing actions. Accordingly, an entity description declares a data source for each one of its sensing capabilities. As well, an actuating capability corresponds to a set of method declarations. An entity declaration also includes attributes, characterizing properties of entity instances. Entity declarations are organized hierarchically allowing entity classes to inherit attributes, sources and actions. A taxonomy allows separation of concerns in that the expert can focus on the concerns of cataloging area-specific entities. The entity developer is concerned about mapping a taxonomical description into an actual entity, and the application developer concentrates on the application logic.

    – The *Architecture Layer* is based on an architectural pattern commonly used in the pervasive computing domain [26]. It consists of context components fueled by sensing entities. These components process gathered data to make them amenable to the application needs. Context data are then passed to controller components that trigger actions on entities. Using an architecture description enables the key components of an application to be identified, allowing their implementation to evolve with the requirements (*e.g.,* varying light management implementations in a controller component to optimize energy consumption).

- DIAGEN is the DIASPEC compiler and runtime, performs both static and runtime verifications over DIASPEC declarations and produces a dedicated programming framework that guides and eases the implementation of components. The generated framework is independent of the underlying distributed technology. As of today, DIAGEN supports multiple targets: Local, RMI, SIP and a simulation target (the Web Services and the Corba targets being currently in development).

5.1.2. DiaSim: a Parametrized Simulator for Pervasive Computing Applications

Pervasive computing applications involve both software and integration concerns. This situation is problematic for testing pervasive computing applications because it requires acquiring, testing and interfacing a variety of software and hardware entities. This process can rapidly become costly and time-consuming when the target environment involves many entities.

To ease the testing of pervasive applications, we are developing a simulator for pervasive computing applications: DIASIM. To cope with widely heterogeneous entities, DIASIM is parameterized with respect to a DIASPEC specification describing a target pervasive computing environment. This description is used to generate with DIAGEN both a programming framework to develop the simulation logic and an emulation layer to execute applications. Furthermore, a simulation renderer is coupled to DIASIM to allow a simulated pervasive system to be visually monitored and debugged. The simulation renderer is illustrated in Figure 2.

5.2. Pantagruel: a Visual Domain-Specific Language for Ubiquitous Computing

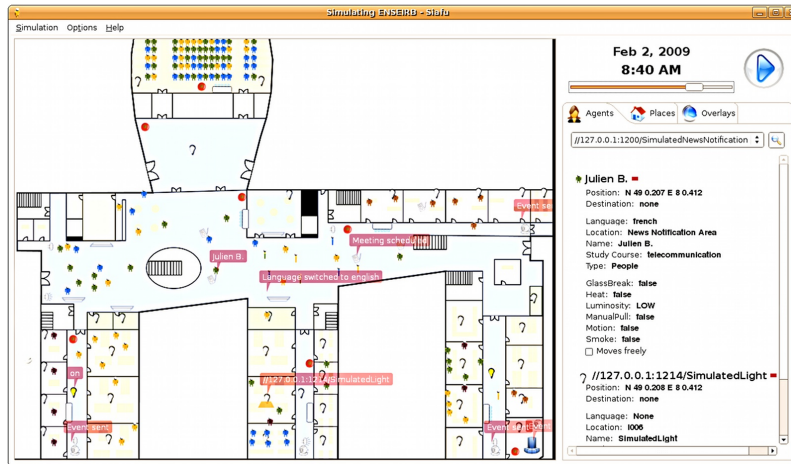**Participants:** Zoé Drey [correspondent], Julien Mercadal, Charles Consel.

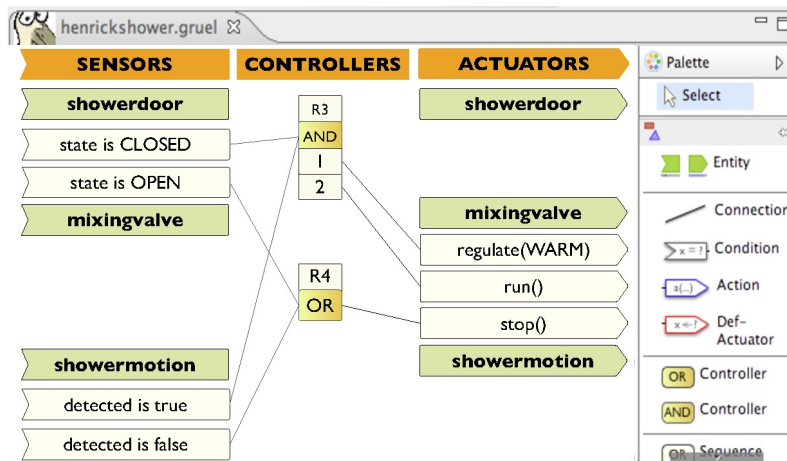*Figure 2. A screenshot of the* DIASIM *simulator*

Figure 3. A screenshot of the Pantagruel graphical editor (2)

Pantagruel aims at easing the description of an orchestration logic between networked entities of a pervasive environment. First, the developer defines a taxonomy of entities that compose the environment, This step provides an abstraction of the entities capabilities and functionalities. Second, the developer defines the orchestration logic in terms of rules. To facilitate its programming, we provide a visual domain-specific language based on the sensor-controller-actuator paradigm. An example of a visual orchestration is given in Figure 3 where a shower automatically runs at the right temperature when someone enters the bathroom and closes the door.

Pantagruel brings a high-level layer intended to complement existing tools in the activity of safe orchestration logic description, allowing novice-programmers to prototype pervasive applications. The Pantragruel compiler generates code compliant with the DIASUITE toolset. Pantagruel is being completed by tools aimed at verifying safety properties like termination and reachability.

See also the web page http://pantagruel.bordeaux.inria.fr.

6. New Results

6.1. Leveraging Software Architectures to Guide and Verify the Development of Sense/Compute/Control Applications

A software architecture describes the structure of a computing system by specifying software components and their interactions. Mapping a software architecture to an implementation is a well known challenge. A key element of this mapping is the architecture's description of the data and control-flow interactions between components. The characterization of these interactions can be rather abstract or very concrete, providing more or less implementation guidance, programming support, and static verification.

In this work, we have introduced a notion of *behavioral contract* that expresses the set of allowed interactions between components, describing both data and control-flow constraints [16]. This declaration is part of the architecture description, allows generation of extensive programming support, and enables various verifications. We have instantiated our approach in an architecture description language for the domain of Sense/Compute/Control (SCC) applications, and described associated compilation and verification strategies.

The main contributions of this work are the following:

- We have introduced a language for behavioral contracts dedicated to SCC applications.
- We have shown that behavioral contracts can effectively guide the implementation of SCC applications by enabling the generation of highly customized programming frameworks using a dedicated compiler. This approach ensures the conformance between the architecture and the implementation, while facilitating software evolution.
- We have shown that such descriptions are precise enough to verify safety properties such as information flow reachability or behavioral invariants.
- Based on an implementation of behavioral contracts in an ADL targeting SCC applications, we haved assessed the benefit of behavioral contracts at a conceptual level and in terms of metrics on the resulting code.

6.2. A Domain-Specific Approach to Architecturing Error Handling in Pervasive Computing

The challenging nature of error handling constantly escalates as a growing number of environments consists of networked devices and software components. In these environments, errors cover a uniquely large spectrum of situations related to each layer ranging from hardware to distributed platforms, to software components. Handling errors becomes a daunting task for programmers, whose outcome is unpredictable. Scaling up error handling requires to raise the level of abstraction beyond the code level and the try-catch construct, approaching error handling at the software architecture level.

We have proposed a novel approach that relies on an Architecture Description Language (ADL), which is extended with error-handling declarations [18].

The main contributions of this work are the following:

- We have proposed a novel approach that raises the level of abstraction of error handling from programming to architecturing. Our approach allows reasoning, and programming is driven by this extended form of software descriptions.

- We have extended a domain-specific ADL with declarations dedicated error handling. These architecture-level declarations provide a separation between functional and error-handling concerns. Furthermore, error handling is made specific by decomposing it into application and system compensation strategies.

- Architecture descriptions are processed by a compiler that generates dedicated programming frameworks in Java. We have extended this compiler to produce additional programming support for signaling, propagating and treating errors that originate as Java exceptions. This support makes the programming of error handling more rigorous and systematic.

- We have used our approach to develop a variety of dependable applications in areas including home/ building automation and healthcare. Our largest case study is a system for managing a *13 500*-square meter building, amounting for more than *3 000* LOC.

6.3. SIP as a Universal Communication Bus: A Methodology and an Experimental Study

In this work, we have proposed a methodology and a programming support that use the SIP protocol as a universal communication bus in pervasive computing environments [13]. In doing so, our work enables homogeneous communications between heterogeneous distributed entities. We have also presented a classification of a wide variety of entities in terms of features, capabilities and network connectors. Based on this classification, a methodology and a programming support are described for connecting entities on the SIP communication bus. This work has been validated by applications using the SIP communication bus to coordinate widely varying entities, including serial-based sensors (RS232, 1-Wire), ZigBee devices, X10 devices, PDA, native SIP entities, and software components.

The main contributions of this work are the following:

- A classification of a wide variety of entities that facilitates their integration in the SIP communication bus.

- A methodology and programming support that make each class of entities SIP compliant.

- An experimental study that validates SIP as a communication bus for pervasive computing environments. This study comprises numerous entities with vastly varying features and capabilities.

7. Contracts and Grants with Industry

7.1. Designing and developing simulation capabilities for network-centric systems – Industrial Fellowship (CIFRE / Thales)

Participants: Charles Consel, Julien Bruneau.

The goal of this project is to provide simulation capabilities for testing network-centric systems. To achieve this goal, a formal description of the component behavior of such system must be defined. Hybrid testing (combining virtual and real) of components, data and scenarios, as well as observability tools for component-tools will be studied in this project.

Models, DSLs and protocols for modeling a network-centric system will be designed and developed during this project. A dedicated framework for the simulation must also be provided. Finally, the simulation of a system must allow to qualify the functional logic of this system.

7.2. Integrating non-functional properties in a Design Language and its execution environment – Industrial Fellowship (CIFRE / Thales)

Participants: Charles Consel, Emilie Balland, Stéphanie Gatti, Quentin Enard.

The goal of this project is to add non-functional properties in the Diaspec language and in the Diagen generator. More especially, these non-functional properties are considered on three different levels:

- *The component level.* The non-functional properties define temporal, physical and software constraints restrictive for a component.

- *The component coupling level.* The component coupling level. The non-functional properties define the dependency between the components as well as the Quality of Service provided and required by each component of the environment.

- *The software architecture level.* The software architecture level. The non-functional properties describe the resources that must be allocated to a component (memory, processing capacity). They also define the necessary resources for a component to interact with other components (network QoS).

This work will be illustrated and validated with a concrete application in the avionics domain.

7.3. SmartImmo: Towards intelligent and environmentally-friendly buildings (french competitiveness pole)

Participants: Charles Consel, Benjamin Bertran, Ghislain Deffrasnes.

The SmartImmo project gathers research groups in pervasive systems and french companies working in the building construction, installation, and management. This project led by Orange Labs aims to make a building able to "communicate" with its occupants and to be environmentally-friendly (*e.g.,* automatic temperature adjusting).

The main objectives of this project are to design a M2M (Machine-To-Machine) box for the heterogeneous equipment communication and to build several services on top of this platform. This project is funded by the SCS (Secured Communicating Solutions), a french pole of competitiveness.

8. Other Grants and Activities

8.1. International Collaborations

We have been exchanging visits and publishing articles with the following collaborators.

- Julia Lawall, DIKU, University of Copenhagen (Denmark, Copenhagen).
- Walid Taha, Rice University (US, Houston).

8.2. Visits and Invited Researchers

The Phoenix group has been visited by:

- Anne-Françoise Le Meur (Associate Professor at the University of Lille, member of the ADAM INRIA project team) from February 15, 2010 to February 16, 2010;

- Julia Lawall (Associate Professor at the University of Copenhagen, Denmark) from March 21, 2010 to March 23, 2010;

- Xavier Blanc (former Associate Professor at the Pierre and Marie Curie University, Paris) on the 18th of March 2010;

- Vicente Sanchez-Leighton (consultant Hyptique, Paris) from September 13, 2010 to September 14, 2010;

- Arne Svensk (University of Lund, Sweden) from September 13, 2010 to September 14, 2010;

- Stefan Parry Carmien (consultant Fatronik, San Sebastian, Spain) from September 13, 2010 to September 14, 2010;

- Yves Lachapelle and Dany Lussier-Desrochers (University of Québec, Trois-Rivières, Canada) from September 13, 2010 to September 14, 2010.

9. Dissemination

9.1. Animation of the scientific community

Charles Consel has been involved in the following events as:

- Program Committee member of
  - ICWS 2010 (IEEE International Conference on Web Services),
  - ICMT 2010 (International Conference on Model Transformation),
  - NFM 2010 (Second NASA Formal Methods Symposium),
  - PSIEtA 2010 (Workshop on Programming Support Innovations for Emerging Distributed Applications, satellite event of SPLASH 2010),
  - SLE 2010 (International Conference on Software Language Engineering),
  - GPCE 2010 (International Conference on Generative Programming and Component Engineering),

- Guest Editor for the Annals of Telecommunications, Springer;

- Member of the scientific commitee on "GDR génie de la programmation du logiciel" (CNRS);

- Member of the steering committee of the International Conference on Generative Programming and Component Engineering (GPCE);

- Member of the IFIP WG 2.11 on Program Generation;

- Member of the INRIA working group on research and perspectives in the domain "Réseaux, systèmes et services, calcul distribué".

- President of the selection committee for the associate professor position at ENSEIRB-IPB.

Charles Consel has participated in the following thesis defense committees:

- Jesper Andersen, University of Copenhagen, Denmark, February,

- Dima Aladidi, University of Concordia, Canada, January,

- Tien Dung Cao, University of Bordeaux, December (president of the PhD thesis defense committee).

Emilie Balland has been involved as a Program Committee member in:

- WASDeTT 2010 (3d workshop on Academic Software Development Tools and Techniques, satellite of ASE 2010);
- LDTA 2010 (10th Workshop on Language Descriptions, Tools and Applications, satellite of ETAPS 2010).

9.2. Participation in Popular Science Events

Participation of the Phoenix INRIA project team in the following events:

- "1000 chercheurs" national event in Paris,
- "Unité ou Café" talk at the INRIA Bordeaux Sud-Ouest research center,
- "Visage des Sciences" regional event in Aquitaine,
- Demonstrations of DiaSuite during the "Fête de la science" national event.

9.3. Teaching

Charles Consel has been teaching Master level courses on:

- Domain-Specific Languages and Program Analysis;
- Telephony over IP (related protocols, the SIP protocol, existing programming interfaces). Students are also offered practical labs on various industrial-strength telephony platforms. These labs are supervised by Benjamin Bertran and Julien Bruneau.

Charles Consel and Damien Cassou have been teaching a course on Architecture Description Languages.

Emilie Balland has been teaching a Master level course on Software Development guided by modeling and verification.

9.4. Presentations and Invitations

Charles Consel gave a number of invited presentations:

- at Concordia University (Montreal, Canada) in January 2010,
- at the University of Waterloo for the ninth meeting of the IFIP WG 2.11 in December 2010.

Emilie Balland gave a number of invited presentations:

- at the University of Halmstad (Sweden) in July 2010,
- at the annual seminar of the Pareo research group in November 2010,
- at the University of Waterloo for the ninth meeting of the IFIP WG 2.11 in December 2010.

9.5. PhD Thesis

One student of the Phoenix group obtained his PhD in 2010:

- Zoé Drey, "Vers une méthodologie dédiée à l'orchestration d'entités communicantes" [11].

10. Bibliography

Major publications by the team in recent years

[1] D. CASSOU, E. BALLAND, C. CONSEL, J. LAWALL. *Leveraging Software Architectures to Guide and Verify the Development of Sense/Compute/Control Applications*, in "ICSE'11: Proceedings of the 33rd International Conference on Software Engineering", Honolulu United States, ACM, 2011, to appear, http://hal.inria.fr/inria-00537789.

[2] C. CONSEL. *From A Program Family To A Domain-Specific Language*, Lecture Notes in Computer Science, State-of-the-Art Survey, Springer-Verlag, 2004, n⁰ 3016, p. 19–29, http://phoenix.labri.fr/publications/papers/dagstuhl-consel.pdf.

[3] C. CONSEL, J. LAWALL, A.-F. LE MEUR. *A Tour of Tempo: A Program Specializer for the C Language*, in "Science of Computer Programming", 2004, http://phoenix.labri.fr/publications/papers/tour-tempo.ps.gz.

[4] C. CONSEL, L. RÉVEILLÈRE. *A Programmable Client-Server Model: Robust Extensibility via DSLs*, in "Proceedings of the 18th IEEE International Conference on Automated Software Engineering (ASE 2003)", Montréal, Canada, IEEE Computer Society Press, November 2003, p. 70–79, http://phoenix.labri.fr/publications/papers/Consel-Reveillere_ase03.pdf.

[5] C. CONSEL, L. RÉVEILLÈRE. *A DSL Paradigm for Domains of Services: A Study of Communication Services*, Lecture Notes in Computer Science, State-of-the-Art Survey, Springer-Verlag, 2004, n⁰ 3016, p. 165–179, http://phoenix.labri.fr/publications/papers/dagstuhl04_consel_reveillere.pdf.

[6] A.-F. LE MEUR, J. LAWALL, C. CONSEL. *Specialization Scenarios: A Pragmatic Approach to Declaring Program Specialization*, in "Higher-Order and Symbolic Computation", 2004, vol. 17, n⁰ 1, p. 47–92, http://phoenix.labri.fr/publications/papers/spec-scenarios-hosc2003.ps.gz.

[7] D. MCNAMEE, J. WALPOLE, C. PU, C. COWAN, C. KRASIC, A. GOEL, P. WAGLE, C. CONSEL, G. MULLER, R. MARLET. *Specialization tools and techniques for systematic optimization of system software*, in "ACM Transactions on Computer Systems", May 2001, vol. 19, n⁰ 2, p. 217–251, http://phoenix.labri.fr/publications/papers/tocs01-namee.pdf.

[8] J. MERCADAL, Q. ENARD, C. CONSEL, N. LORIANT. *A Domain-Specific Approach to Architecturing Error Handling in Pervasive Computing*, in "OOPSLA'10: Proceedings of the 25th Annual ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications", États-Unis Reno, October 2010, http://hal.inria.fr/inria-00486930/en.

[9] F. MÉRILLON, L. RÉVEILLÈRE, C. CONSEL, R. MARLET, G. MULLER. *Devil: An IDL for Hardware Programming*, in "Proceedings of the Fourth Symposium on Operating Systems Design and Implementation", San Diego, California, October 2000, p. 17–30, http://phoenix.labri.fr/publications/papers/osdi00-merillon.pdf.

[10] S. THIBAULT, C. CONSEL, G. MULLER. *Safe and Efficient Active Network Programming*, in "17th IEEE Symposium on Reliable Distributed Systems", West Lafayette, IN, October 1998, p. 135–143, http://phoenix.labri.fr/publications/papers/srds98-thibault.ps.gz.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] Z. DREY. *Vers une méthodologie d'ediée à l'orchestration d'entités communicantes*, Université Sciences et Technologies - Bordeaux, September 2010.

International Peer-Reviewed Conference/Proceedings

[12] E. BALLAND, C. CONSEL. *Open Platforms: New Challenges for Software Engineering*, in "PSIEtA'10: Proceedings of the International Workshop on Programming Support Innovations for Emerging Distributed Applications", ACM Digital Library, November 2010, to appear, http://hal.inria.fr/inria-00533721/en.

[13] B. BERTRAN, C. CONSEL, W. JOUVE, H. GUAN, P. KADIONIK. *SIP as a Universal Communication Bus: A Methodology and an Experimental Study*, in "ICC'10: Proceedings of the IEEE International Conference on Communications", Afrique Du Sud Cape Town, May 2010 [*DOI :* 10.1109/ICC.2010.5502591], http://hal.inria.fr/inria-00453548/en.

[14] D. CASSOU, J. BRUNEAU, C. CONSEL. *A Tool Suite to Prototype Pervasive Computing Applications (Demo)*, in "PERCOM'10: Proceedings of the 8th IEEE Conference on Pervasive Computing and Communications", Allemagne Mannheim, IEEE Computer Society Press, 2010, p. 1–3, http://hal.inria.fr/inria-00484067/en.

[15] D. CASSOU, J. BRUNEAU, J. MERCADAL, Q. ENARD, E. BALLAND, N. LORIANT, C. CONSEL. *Towards a Tool-based Development Methodology for Sense/Compute/Control Applications (Poster)*, in "SPLASH'10: Proceedings of the 1st International Conference on Systems, Programming, Languages, and Applications: Software for Humanity", États-Unis Reno/Tahoe, ACM, 2010, p. 1–2, poster, http://hal.inria.fr/inria-00510378/en.

[16] D. CASSOU, E. BALLAND, C. CONSEL, J. LAWALL. *Leveraging Software Architectures to Guide and Verify the Development of Sense/Compute/Control Applications*, in "ICSE'11: Proceedings of the 33rd International Conference on Software Engineering", Honolulu United States, ACM, 2011, to appear, http://hal.inria.fr/inria-00537789.

[17] Z. DREY, C. CONSEL. *A Visual, Open-Ended Approach to Prototyping Ubiquitous Computing Applications*, in "PERCOM'10: Proceedings of the 8th IEEE Conference on Pervasive Computing and Communications", Allemagne Mannheim, March 2010, http://hal.inria.fr/inria-00484083/en.

[18] J. MERCADAL, Q. ENARD, C. CONSEL, N. LORIANT. *A Domain-Specific Approach to Architecturing Error Handling in Pervasive Computing*, in "OOPSLA'10: Proceedings of the 25th Annual ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications", États-Unis Reno, October 2010, http://hal.inria.fr/inria-00486930/en.

[19] H. WAN, Z. DREY, Z. YOU, L. LIU. *Formal Modeling and Verification of Services Managements for Pervasive Computing Environment*, in "ICSSSM'10: Proceedings of the 7th International Conference on Service Systems and Service Management", Japon Tokyo, June 2010, http://hal.inria.fr/inria-00516020/en.

References in notes

[20] *CGI: The Common Gateway Interface*, 1993, http://www.w3.org/CGI/.

[21] *Session Initiation Protocol (SIP)*, March 2001, Request for Comments 3261.

[22] P. BOINOT, R. MARLET, J. NOYÉ, G. MULLER, C. CONSEL. *A Declarative Approach for Designing and Developing Adaptive Components*, in "Proceedings of the 15th IEEE International Conference on Automated Software Engineering (ASE 2000)", Grenoble, France, IEEE Computer Society Press, September 2000, p. 111–119.

[23] S. CHIROKOFF, C. CONSEL, R. MARLET. *Combining Program and Data Specialization*, in "Higher-Order and Symbolic Computation", December 1999, vol. 12, n^o^ 4, p. 309–335.

[24] C. CONSEL, J. LAWALL, A.-F. LE MEUR. *A Tour of Tempo: A Program Specializer for the C Language*, in "Science of Computer Programming", 2004.

[25] C. CONSEL, R. MARLET. *Architecturing software using a methodology for language development*, in "Proceedings of the 10th International Symposium on Programming Language Implementation and Logic Programming", Pisa, Italy, C. PALAMIDESSI, H. GLASER, K. MEINKE (editors), Lecture Notes in Computer Science, September 1998, vol. 1490, p. 170–194.

[26] A. K. DEY, G. D. ABOWD, D. SALBER. *A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications*, in "Human-Computer Interaction", 2001, vol. 16, n^o^ 2, p. 97–166.

[27] A.-F. LE MEUR, C. CONSEL, B. ESCRIG. *An Environment for Building Customizable Software Components*, in "IFIP/ACM Conference on Component Deployment", Berlin, Germany, June 2002, p. 1–14.

[28] A.-F. LE MEUR, C. CONSEL. *Generic Software Component Configuration Via Partial Evaluation*, in "SPLC'2000 Workshop – Product Line Architecture", Denver, Colorado, August 2000.

[29] A.-F. LE MEUR, J. LAWALL, C. CONSEL. *Towards Bridging the Gap Between Programming Languages and Partial Evaluation*, in "ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation", Portland, OR, USA, ACM Press, January 2002, p. 9–18.

[30] A.-F. LE MEUR, J. LAWALL, C. CONSEL. *Specialization Scenarios: A Pragmatic Approach to Declaring Program Specialization*, in "Higher-Order and Symbolic Computation", 2004, vol. 17, n^o^ 1, p. 47–92.

[31] R. MARLET, C. CONSEL, P. BOINOT. *Efficient Incremental Run-Time Specialization for Free*, in "Proceedings of the ACM SIGPLAN'99 Conference on Programming Language Design and Implementation (PLDI'99)", Atlanta, GA, USA, May 1999, p. 281–292.

[32] R. MARLET, S. THIBAULT, C. CONSEL. *Mapping Software Architectures to Efficient Implementations via Partial Evaluation*, in "Conference on Automated Software Engineering", Lake Tahoe, NV, USA, IEEE Computer Society, November 1997, p. 183–192.

[33] R. MARLET, S. THIBAULT, C. CONSEL. *Efficient Implementations of Software Architectures via Partial Evaluation*, in "Journal of Automated Software Engineering", October 1999, vol. 6, n^o^ 4, p. 411–440.

[34] D. MCNAMEE, J. WALPOLE, C. PU, C. COWAN, C. KRASIC, A. GOEL, P. WAGLE, C. CONSEL, G. MULLER, R. MARLET. *Specialization tools and techniques for systematic optimization of system software*, in "ACM Transactions on Computer Systems", May 2001, vol. 19, n^o^ 2, p. 217–251.

[35] N. MEDVIDOVIC, R. N. TAYLOR. *A Classification and Comparison Framework for Software Architecture Description Languages*, in "IEEE Transactions on Software Engineering", 2000, vol. 26, n^o^ 1, p. 70–93, http://dx.doi.org/10.1109/32.825767.

[36] G. MULLER, C. CONSEL, R. MARLET, L. BARRETO, F. MÉRILLON, L. RÉVEILLÈRE. *Towards Robust OSes for Appliances: A New Approach Based on Domain-Specific Languages*, in "Proceedings of the ACM SIGOPS European Workshop 2000 (EW2000)", Kolding, Denmark, ACM Press, September 2000, p. 19-24.

[37] F. MÉRILLON, L. RÉVEILLÈRE, C. CONSEL, R. MARLET, G. MULLER. *Devil: An IDL for Hardware Programming*, in "4th Symposium on Operating Systems Design and Implementation (OSDI 2000)", San Diego, California, October 2000, p. 17–30.

[38] L. RÉVEILLÈRE, F. MÉRILLON, C. CONSEL, R. MARLET, G. MULLER. *A DSL Approach to Improve Productivity and Safety in Device Drivers Development*, in "Proceedings of the 15th IEEE International Conference on Automated Software Engineering (ASE 2000)", Grenoble, France, IEEE Computer Society Press, September 2000, p. 101–109.

[39] L. RÉVEILLÈRE, G. MULLER. *Improving Driver Robustness: an Evaluation of the Devil Approach*, in "The International Conference on Dependable Systems and Networks", Göteborg, Sweden, IEEE Computer Society, July 2001, p. 131–140.

[40] U. SCHULTZ, J. LAWALL, C. CONSEL, G. MULLER. *Towards Automatic Specialization of Java Programs*, in "Proceedings of the European Conference on Object-oriented Programming (ECOOP'99)", Lisbon, Portugal, Lecture Notes in Computer Science, June 1999, vol. 1628, p. 367–390.

[41] U. SCHULTZ, J. LAWALL, C. CONSEL. *Specialization Patterns*, in "Proceedings of the 15th IEEE International Conference on Automated Software Engineering (ASE 2000)", Grenoble, France, IEEE Computer Society Press, September 2000, p. 197–208.

[42] U. SCHULTZ, J. LAWALL, C. CONSEL. *Automatic Program Specialization for Java*, in "ACM Transactions on Programming Languages and Systems", 2003, vol. 25, n^o 4, p. 452–499.

[43] S. THIBAULT, C. CONSEL, J. LAWALL, R. MARLET, G. MULLER. *Static and Dynamic Program Compilation by Interpreter Specialization*, in "Higher-Order and Symbolic Computation", September 2000, vol. 13, n^o 3, p. 161–178.

[44] S. THIBAULT, R. MARLET, C. CONSEL. *Domain-Specific Languages: from Design to Implementation – Application to Video Device Drivers Generation*, in "IEEE Transactions on Software Engineering", May 1999, vol. 25, n^o 3, p. 363–377.