Handbook of Research on Service-Oriented Systems and Non-Functional Properties:

Future Directions

Stephan Reiff-Marganiec
University of Leicester, UK

Marcel Tilly
European Microsoft Innovation Center, Germany

Information Science REFERENCE

Managing Director: Lindsay Johnston
Senior Editorial Director: Heather Probst
Book Production Manager: Sean Woznicki
Development Manager: Joel Gamon
Development Editor: Michael Killian
Acquisitions Editor: Erika Carter
Typesetters: Lisandro Gonzalez, Mackenzie Snader
Print Coordinator: Jamie Snavely
Cover Design: Nick Newcomer

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

Library of Congress Cataloging-in-Publication Data

Handbook of research on service-oriented systems and non-functional properties : future directions / Stephan Reiff-Marganiec and Marcel Tilly, editors.
p. cm.
Includes bibliographical references and index.
ISBN 978-1-61350-432-1 (hardcover) -- ISBN 978-1-61350-433-8 (ebook) -- ISBN 978-1-61350-434-5 (print & perpetual access) 1. Service-oriented architecture (Computer science)--Handbooks, manuals, etc. 2. Business--Data processing--Handbooks, manuals, etc. I. Reiff-Marganiec, Stephan. II. Tilly, Marcel, 1971-
TK5105.5828.H36 2012
004.6’54--dc23
2011031924

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

Chapter 2
Verification of Non-Functional Requirements by Abstract Interpretation

Agostino Cortesi
Università Ca' Foscari, Italy

Francesco Logozzo
Microsoft Research, USA

ABSTRACT

This chapter investigates a formal approach to the verification of non-functional software requirements that are crucial in Service-oriented Systems, like portability, time and space efficiency, and dependability/robustness. The key-idea is the notion of observable, i.e., an abstraction of the concrete semantics when focusing on a behavioral property of interest. By applying an abstract interpretation-based static analysis of the source program, and by a suitable choice of abstract domains, it is possible to design formal and effective tools for non-functional requirements validation.

INTRODUCTION

Effective and efficient management of customer and user requirements is one of the most crucial, but unfortunately also least understood issues (Karlsson, 1997), in particular for Service Oriented Systems. In Service Oriented Architectures the non-functional aspects of services and connections should be defined separately from their functional aspects because different applications use the services and connections in different non-functional contexts. The separation between functional and non-functional aspects improves the reusability of services and connections. It also enables the two different aspects to evolve independently, and improves the ease of understanding application architectures. This contributes to increase the maintainability of applications (Wada, Suzuki, & Oba, 2006 and O'Brien, Merson, & Bass, 2007).

Problems in the non-functional requirements are typically not recognized until late in the development process, where negative impacts are substantial and cost for correction has grown large. Even worse, problems in the requirements may

DOI: 10.4018/978-1-61350-432-1.ch002

go undetected through the development process, resulting in software systems not meeting customers and users expectations, especially when the coordination with other components is an issue. Therefore, methods and frameworks helping software developers to better manage software requirements are of great interest for component based software.

Abstract interpretation (Cousot & Cousot, 1977) is a theory of semantics approximation for computing conservative over-approximations of dynamic properties of programs. It has been successfully applied to infer run-time properties useful for debugging (e.g., type inference (Cousot, 1997 and Kozen, Palsberg, & Schwartzbach, 1994)), code optimization (e.g., compile-time garbage collection (Hughes, 1992)), program transformation (e.g., partial evaluation (Jones, 1997), parallelization (Traub, Culler & Schauser, 1992)), and program correctness proofs (e.g., safety (Halbwachs, 1998), termination (Brauburger, 1997), cryptographic protocol analysis (Monniaux, 2003), proof of absence of run-time errors (Blanchet, Cousot, Cousot, et al., 2003), semantic tattooing/watermarking (Cousot & Cousot, 2004)). As pointed out in (Le Métayer, 1996), there is still a large variety of tasks in the software engineering process that could greatly benefit from techniques akin to static program analysis, because of their firm theoretical foundations and mechanical nature.

In this chapter we investigate the impact of Abstract Interpretation theory in the formalization and automatic verification of Non-Functional Software Requirements, as they seem not adequately covered by most requirements engineering methods ((Kotonya, & Sommerville, 1998), pag. 194). Non functional requirements can be defined as restrictions or constraints on the behavior of a system service (Sommerville, 2000). Different classifications have been proposed in the literature (Boehm, 1976, Davis, 1992, and Deutsch & Willis, 1988)), though their specification may give rise to troubles both in their elicitation and management, and in the validation process.

Let us start from a quite naive question: "*what do we mean when we say that a program is portable on a different architecture?*". In (Ghezzi, Jazayeri & Mandrioli, 2003) a software is said portable if it can run in different environments. It is clear that it is assumed not only that it runs, but that it runs the same way. And it is also clear that if we require that the behavior is exactly the same, portability to different systems (e.g., from a PC to a PDA, or from an OS to another) can almost never be reached. This means that implicit assumptions are obviously made about the properties to be preserved, and about the ones that might be simply disregarded. In other words, portability needs to be parameterized on some specific properties of interest, i.e. it assumes a suitable abstraction of the software behavior. The same holds also for other product non-functional requirements, like space and time efficiency, dependability, robustness, usability, etc. It is clear that, in this context, the main features of abstract interpretation theory, namely modularity, modulability, and effectiveness may then become very valuable.

The main concepts introduced in this chapter can be summarized as follows:

- We extend the usual abstract interpretation notions to the deal with systems, i.e. programs + architectures.
- We show that a significant set of product qualities (non-functional requirements) can be formally expressed in terms of abstraction of the concrete semantics when focusing on a behavioral property of interest. This yields an unifying view of product non-functional requirements.
- We show how existing tools for automatic verification can be re-used in this setting to support requirements validation; their practicality directly depends on the complexity of the abstract domains.

The advantage of this approach with respect to previous attempts of modelling software requirements, e.g., by using Milner's Calculus of Communicating Systems (Halbwachs, 1995) or formal methods like Z (Spivey, 1992) or B (Abrial, 1996 and Abrial, 2003) is twofold: (1) the soundness of the approach is guaranteed by the general abstract interpretation theory, and (2) the automatic validation process can be easily tuned according to the desired granularity of the abstraction.

Applying the Abstract Interpretation theory to the treatment of non-functional software requirements (Cortesi & Logozzo, 2005) can be seen as a contribution towards the achievement of a more challenging objective: to integrate formal analysis by abstract interpretation in the full software development process, from the initial specifications to the ultimate program development (Cousot, 2001).

Chapter Structure: In Section 2, the concrete semantics of a simple imperative language is introduced to instantiate our framework. In Section 3, the core abstract interpretation theory is extended to deal with program and architecture abstractions. In Section 4 we show how to instantiate our framework on a suite of non-functional product requirements. In Section 5 we discuss its use in the Service Oriented scenario. Section 6 concludes the paper.

OPERATIONAL SEMANTICS OF A CORE IMPERATIVE LANGUAGE WITH EXCEPTIONS

In order to better illustrate the approach, we instantiate our framework with a core imperative language with exceptions and a core architecture. The results can be generalized to more complex languages and architectures. We give the syntax, the transition relations and the trace semantics of systems, composed by architectures and a programs.

Syntax

We let an architecture be a tuple $\langle bits, Op, stdio, stdout \rangle$, where $bits$ is the number of bits used to store integer numbers, Op is a set of functions implementing basic arithmetic operations, $stdio$ is the input stream (e.g., the keyboard) and $stdout$ is the output stream (e.g., the screen). The input stream has a method $next$ that returns immediately the next value in the stream, and the output stream has a method add to put a pair $\langle v, c \rangle$, i.e., a value v with a color c. We assume that if an arithmetic error occurs in the application of an operation $op \in Op$ (e.g., an overflow or a division by zero), then the exception $\mathrm{ExcMath}$ is raised.

The syntax of programs is specified by the following grammar:

$$
\begin{aligned}
\mathrm{C} ::= \; & \mathrm{skip} \mid \mathrm{x{=}E} \mid \mathrm{C_1;C_2} \mid \mathrm{if\ (E!{=}0)\ C_1\ else\ C_2} \mid \mathrm{while\ (E!{=}0)\ C} \\
& \mathrm{write(x,col)} \mid \mathrm{throw\ Exc} \mid \mathrm{try\ C_1\ catch(Exc)\ C_2} \\
\mathrm{E} ::= \; & k \mid \mathrm{read} \mid \mathrm{E_1 + E_2} \mid \mathrm{E_1 - E_2} \mid \mathrm{E_1 * E_2} \mid \mathrm{E_1 / E_2}
\end{aligned}
$$

where x and col belong to a given set Var of variables, Exc belongs to a given set $\mathrm{Exceptions}$ of exceptions (including the arithmetic ones) and k is (the *syntactic* representation of) an integer number.

A system is a pair $\langle \mathrm{A,C} \rangle$, where A is an architecture and C is a program.

Semantics

The semantics of a system is described in operational style. We assume that the only available type is that of architecture-representable natural numbers: $\mathsf{N}_{bits} = \{0, \ldots, 2^{bits} - 1\}$. Given the *syntactic* representation k of a number, $\underline{k}$ is the *semantic* correspondent. For instance, $\underline{\mathrm{0xFFFF}} = 65535$ so that $\underline{\mathrm{0xFFFF}}. \notin \mathsf{N}_8$. An environment is a partial map from variables to representable integers: Env= $[\mathrm{Var} \rightarrow \mathsf{N}_{bits}]$. If a variable x is not defined in a state σ, we denote that by $\sigma(\mathrm{x}) = \Omega$. A state is either a command to execute in a given environment, or an environment, or an exception raised within an environment. Formally:

$$\Sigma = \mathrm{C} \times \mathrm{Env} \cup \mathrm{Env} \cup \mathrm{Exceptions} \times \mathrm{Env}\,.$$

The transition relations for expressions and programs are defined by structural induction, and they are depicted in Figure 1. It is worth noting that the transition rules are parameterized by the underlying architecture (e.g., the raising of an overflow exception depends on N_{bits}).

Let Σ^* denote the set of finite traces on Σ, and let $S_0 \subseteq \Sigma$ be a set of initial states. With a slight abuse of notation, we refer to a state as a trace of unitary length. The partial-traces semantics (Cousot & Cousot, 2002) of a system is then

Figure 1. The transition relations for expressions and programs

$$s\langle \mathrm{A},\mathrm{C}\rangle(S_0) = \mathrm{lfp}_{\varnothing}^{\subseteq}\lambda X.\, S_0 \cup \{\sigma_0 \ldots \sigma_n \sigma_{n+1} \mid \sigma_0 \ldots \sigma_n \in X, \sigma_n \longrightarrow \sigma_{n+1}\}.$$

$$\frac{\underline{k} \in \mathsf{N}_{bits}}{\langle \underline{k},\sigma\rangle \xrightarrow{\mathrm{E}} k} \qquad \frac{\underline{k} \notin \mathsf{N}_{bits}}{\langle \underline{k},\sigma\rangle \xrightarrow{\mathrm{E}} \langle \mathrm{ExcMath}\;,\sigma\rangle} \qquad \frac{\mathrm{A}.\mathit{stdio.next} = v}{\langle \mathrm{read},\sigma\rangle \xrightarrow{\mathrm{E}} \langle v,\sigma\rangle}$$

$$\frac{\langle \mathrm{E}_1,\sigma\rangle \xrightarrow{\mathrm{E}} \langle v_1,\sigma\rangle\; \langle \mathrm{E}_2,\sigma\rangle \xrightarrow{\mathrm{E}} \langle v_2,\sigma\rangle\; v_1, v_2 \neq \mathrm{ExcMath}\quad \mathrm{A}.\mathit{op}(v_1,v_2) = v \neq \mathrm{ExcMath}}{\langle \mathrm{E}_1 \mathit{op} \mathrm{E}_2,\sigma\rangle \xrightarrow{\mathrm{E}} \langle v,\sigma\rangle}$$

$$\frac{\langle \mathrm{E}_1,\sigma\rangle \xrightarrow{\mathrm{E}} \langle v_1,\sigma\rangle\; \langle \mathrm{E}_2,\sigma\rangle \xrightarrow{\mathrm{E}} \langle v_2,\sigma\rangle\; v_1, v_2 \neq \mathrm{ExcMath}\quad \mathrm{A}.\mathit{op}(v_1,v_2) = \mathrm{ExcMath}}{\langle \mathrm{E}_1 \mathit{op} \mathrm{E}_2,\sigma\rangle \xrightarrow{\mathrm{E}} \langle \mathrm{ExcMath}\;,\sigma\rangle}$$

$$\frac{\langle \mathrm{E}_1,\sigma\rangle \xrightarrow{\mathrm{E}} \langle v_1,\sigma\rangle\; \langle \mathrm{E}_2,\sigma\rangle \xrightarrow{\mathrm{E}} \langle v_2,\sigma\rangle\; (v_1 = \mathrm{ExcMath}\;)\mathit{or}(v_2 = \mathrm{ExcMath}\;)}{\langle \mathrm{E}_1 \mathit{op} \mathrm{E}_2,\sigma\rangle \xrightarrow{\mathrm{E}} \langle \mathrm{ExcMath}\;,\sigma\rangle}$$

$$\frac{}{\langle \mathrm{skip},\sigma\rangle \longrightarrow \sigma} \qquad \frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle v,\sigma\rangle\; v \neq \mathrm{ExcMath}}{\langle \mathrm{x} = \mathrm{E},\sigma\rangle \longrightarrow \sigma[x \mapsto v]} \qquad \frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle \mathrm{ExcMath}\;,\sigma\rangle}{\langle \mathrm{x} = \mathrm{E},\sigma\rangle \longrightarrow \langle \mathrm{ExcMath}\;,\sigma\rangle}$$

$$\frac{\langle \mathrm{C}_1,\sigma\rangle \longrightarrow \sigma'}{\langle \mathrm{C}_1;\mathrm{C}_2,\sigma\rangle \longrightarrow \langle \mathrm{C}_2,\sigma'\rangle} \qquad \frac{\langle \mathrm{C}_1,\sigma\rangle \longrightarrow \langle \mathrm{Exc},\sigma\rangle}{\langle \mathrm{C}_1;\mathrm{C}_2,\sigma\rangle \longrightarrow \langle \mathrm{Exc},\sigma\rangle}$$

$$\frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle \underline{k},\sigma\rangle\; \underline{k} \neq 0}{\langle \mathrm{if}(\mathrm{E}\,!\!= 0)\mathrm{C}_1 \;\mathrm{else}\; \mathrm{C}_2,\sigma\rangle \longrightarrow \langle \mathrm{C}_1,\sigma\rangle} \qquad \frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle 0,\sigma\rangle}{\langle \mathrm{if}(\mathrm{E}\,!\!= 0)\mathrm{C}_1 \;\mathrm{else}\; \mathrm{C}_2,\sigma\rangle \longrightarrow \langle \mathrm{C}_2,\sigma\rangle}$$

$$\frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle \mathrm{ExcMath}\;,\sigma\rangle}{\langle \mathrm{if}(\mathrm{E}\,!\!= 0)\mathrm{C}_1 \;\mathrm{else}\; \mathrm{C}_2,\sigma\rangle \longrightarrow \langle \mathrm{ExcMath}\;,\sigma\rangle}$$

$$\frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle \underline{k},\sigma\rangle\; \underline{k} \neq 0}{\langle \mathrm{while}(\mathrm{E}\,!\!= 0)\;\mathrm{C},\sigma\rangle \longrightarrow \langle \mathrm{C};\mathrm{while}(\mathrm{E}\,!\!= 0)\;\mathrm{C},\sigma\rangle} \qquad \frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle 0,\sigma\rangle}{\langle \mathrm{while}(\mathrm{E}\,!\!= 0)\;\mathrm{C},\sigma\rangle \longrightarrow \sigma}$$

$$\frac{\langle \mathrm{E},\sigma\rangle \xrightarrow{\mathrm{E}} \langle \mathrm{ExcMath}\;,\sigma\rangle}{\langle \mathrm{while}(\mathrm{E}\,!\!= 0)\;\mathrm{C},\sigma\rangle \longrightarrow \langle \mathrm{ExcMath}\;,\sigma\rangle}$$

$$\frac{\mathrm{A}.\mathit{stdout.add}(\sigma(\mathrm{x}),\sigma(\mathrm{col}))}{\langle \mathrm{write}(\mathrm{x},\mathrm{col}),\sigma\rangle \longrightarrow \sigma} \qquad \frac{\mathrm{Exc} \in \mathrm{Exceptions}}{\langle \mathrm{throw\;Exc},\sigma\rangle \longrightarrow \langle \mathrm{Exc},\sigma\rangle}$$

$$\frac{\langle \mathrm{C}_1,\sigma\rangle \longrightarrow \sigma'}{\langle \mathrm{try}\;\mathrm{C}_1\;\mathrm{catch}(\mathrm{Exc})\,\mathrm{C}_2,\sigma\rangle \longrightarrow \sigma'} \qquad \frac{\langle \mathrm{C}_1,\sigma\rangle \longrightarrow \langle \mathrm{Exc},\sigma'\rangle}{\langle \mathrm{try}\;\mathrm{C}_1\;\mathrm{catch}(\mathrm{Exc})\,\mathrm{C}_2,\sigma\rangle \longrightarrow \langle \mathrm{C}_2,\sigma'\rangle}$$

$$\frac{\langle \mathrm{C}_1,\sigma\rangle \longrightarrow \langle \mathrm{Exc}',\sigma'\rangle\; \mathrm{Exc}' \neq \mathrm{Exc}}{\langle \mathrm{try}\;\mathrm{C}_1\;\mathrm{catch}(\mathrm{Exc})\,\mathrm{C}_2,\sigma\rangle \longrightarrow \langle \mathrm{Exc}',\sigma'\rangle}$$

expressed as a least fixpoint over the complete boolean lattice $\langle \Sigma^*, \subseteq \rangle$ as follows.

ABSTRACTING SYSTEMS = PROGRAMS + ARCHITECTURES

Abstract interpretation (Cousot & Cousot, 1977) is a general theory of approximation which formalizes the idea that the semantics of a program can be more or less precise depending on the considered observation level. In this section we revise some basic concepts, and we extend them to deal with composed systems.

In the abstract interpretation terminology, $\langle \Sigma^*, \subseteq \rangle$ is the *concrete domain*, its elements are semantic properties, and the order $\subseteq$ stands for the logical implication. As a consequence, the most precise property about the behavior of a system is the semantics s, called the *concrete* semantics (Cousot, 1999). Set of traces are approximated are represented by suitable abstract elements, which capture interesting properties while disregarding other execution properties that are out of the scope of interest. Abstract properties (or elements) belong to an *abstract domain of observables*, $\bar{\mathrm{D}}$, and they are ordered according to $\leq$, the abstract counterpart for logical implication. In this work we assume that $\langle \bar{\mathrm{D}}, \leq \rangle$ is a complete lattice.

The correspondence between the concrete and the abstract semantic domains is given by a pair of monotonic functions $\langle \alpha, \gamma \rangle$. The function $\alpha \in [\mathsf{P}(\Sigma^*) \rightarrow \bar{\mathrm{D}}]$, called the abstraction function, formalizes the notion of the abstraction ($\mathsf{P}(\Sigma^*)$ denotes the powerset of Σ^*), and $\alpha(T)$ represents the *best* approximation in $\bar{\mathrm{D}}$ of the set of traces T (with respect to the order in $\bar{\mathrm{D}}$). If $\alpha(T) \leq \mathrm{p}$ then p is also a correct, although less precise, abstract approximation of T. On the other hand, the function $\gamma \in [\bar{\mathrm{D}} \rightarrow \mathsf{P}(\Sigma^*)]$, called the concretization function, returns the set of traces that are captured by an abstract property p. The abstraction and concretization functions must satisfy the following property:

$$\forall T \in \mathsf{P}(\Sigma^*). \forall \mathrm{d} \in \bar{\mathrm{D}}.\ \alpha(T) \leq \mathrm{d} \Leftrightarrow T \subseteq \gamma(\mathrm{d}),$$

in such a case, we say that $\langle \alpha, \gamma \rangle$ form a Galois connection between the concrete and the abstract domains. We write is as

$$(\langle \mathsf{P}(\Sigma^*), \subseteq \rangle, \alpha, \gamma, \langle \bar{\mathrm{D}}, \leq \rangle) \qquad (1)$$

The abstract semantics of a system, $\bar{s}$, is defined over an abstract domain that is linked to the concrete domain by a Galois connection. It must satisfy the soundness criterion, (Cousot & Cousot, 1977):

$$\forall S_0 \subseteq \Sigma.\ \alpha(s\langle \mathrm{A,C} \rangle(S_0)) \leq \bar{s}\langle \mathrm{A,C} \rangle(\alpha(S_0)).$$

The soundness criterion above imposes that, when the properties encoded by a given abstract domain are considered, the abstract semantics s $\langle \mathrm{A,C} \rangle$ captures all the behaviors of $\langle \mathrm{A,C} \rangle$. As a consequence, given a specification of a system $\langle \mathrm{A,C} \rangle$ expressed as an abstract property p, if $\bar{s}\langle \mathrm{A,C} \rangle(\alpha(S_0)) \leq \mathrm{p}$, by the soundness criterion and by the transitivity of $\leq$, we have that

$$\alpha(s\langle \mathrm{A,C} \rangle(S_0)) \leq \mathrm{p}.$$

This means that $\langle \mathrm{A,C} \rangle$ respects the specification p.

In the following, we instantiate the abstract domain and p in order to reflect non-functional requirements of systems and we show how well-known static analyses can be re-used in this enhanced context for the automatic verification of such properties.

APPLICATION: NON-FUNCTIONAL REQUIREMENT ANALYSIS

Non-functional software requirements are requirements which are not directly concerned with the specific functions delivered by the system (Sommerville, 2000). They may relate to emergent system properties such as reliability, response time and store occupancy. Alternatively, they may define constraints on the system like the data representation used in system interfaces.

The 'IEEE-Std 830 – 1993' (IEEE, 1988) presents a comprehensive list of non-functional requirements. In the following we will focus on a few of such requirements, namely *portability*, *efficiency*, *robustness* and *usability*. The approach can be extended to cope with other non-functional requirements.

In this section, we show (1) how such requirements admit a rigorous formalization, unlike, e.g., what stated in (Kotonya, & Sommerville, 1998, section 8.2), (2) how, by a suitable choice of abstract domains, existing tools can be re-used to verify such requirements, and (3) the effectiveness of the approach on a public-domain static analyzer (Cousot, 1999).

Portability

Informal Definition

According to (Ghezzi, Jazayeri & Mandrioli, 2003), a software "*is portable if it can run on different environments*". The term *environment* may refer to a hardware platform or a software environment. Analogously, another widespread textbook, (Meyer, 1997), defines portability as "*the ease of transferring software products to various hardware and software environments*". The first observation is that the two definitions implicitly link the requirement to unspecified software metrics. Furthermore, as any natural-based language specifications, they are intrinsically ambiguous. For instance, the word "*run*" can be read as just the possibility of recompiling and executing the software on different system, but also as the request that some behavioral properties of the software are preserved in different platforms.

Formal Definition

We specify portability as a property of the execution of a program that is preserved when it is ported on different architectures. This means that up to a certain property of interest, the behavior of a software is the same on a different architecture. [Portability] Let us consider a program C, an architecture A and a Galois connection ($\langle \mathsf{P}(\Sigma^{*}), \subseteq \rangle, \alpha, \gamma, \langle \bar{\mathrm{D}}, \leq \rangle$). We say that C, developed on A, is portable on the architecture B w.r.t. the observable domain D, if

$$\forall S_0 \subseteq \Sigma.\ \alpha(s\langle \mathrm{B},\mathrm{C}\rangle(S_0)) \leq \alpha(s\langle \mathrm{A},\mathrm{C}\rangle(S_0)).$$

Abstraction

A class property one is interested to keep unchanged among different porting of the software is the behavior w.r.t. arithmetic overflow. For instance, the violation of such a property in porting the control software on a different architecture was at the origin of the Arianne V crash (Lacan, Monfort, Ribal, et al, 1998).

Arithmetic overflow can be checked by using numerical abstract domains, e.g., (Cousot & Cousot, 1977, Cousot & Halbwachs, 1978 and Miné, 2001). In such domains the range of the values assumed by a variable can be constrained so that it can be checked against the largest representable number in a given architecture.

Example

Let us consider the program C in Algorithm 1(a), and let us consider an architecture A such that $\mathrm{A}.bits = 32$. We can use the Intervals abstract domain (Cousot & Cousot, 1977), and the public-domain static analyzer (Cousot, 1999) to infer that $\overline{s}\langle A,C\rangle(i \mapsto [-\infty,+\infty]) = [1,2^{16}]$, and as 2^{16} is representable on a 32 bit architecture, then program C does not cause any arithmetic overflow. As a consequence, by the soundness of the static analysis (guaranteed by abstract interpretation theory), we can safely infer that the program is portable to any architecture in which 2^{16} is representable (this is not the case in a 16 bits architecture).

Oberve that what is described above is a "low level" portability, assuming implicitly the same interface with the rest of the world (graphical interface, for instance). But we can also define portability as reproducing the same user experience, and that may require different behaviors (for example, to cater for temporary loss of connectivity in mobile systems – of course, this is another example of restriction of domain).

Efficiency

Informal Definition

In the existing literature, efficiency "*refers to how economically the software utilizes the resources of the computer*" (Ghezzi, Jazayeri & Mandrioli,

Algorithm 1. Four programs on which we verify non-functional requirements

```
i=1;
while (2^16-i != 0)
i =i*4;
```

(a) C, a program non portable on 16 bits architectures

```
try
  i=?;
  if (i!=0) c = i/0
  else throw Err
catch (Err)
  c=0;
write(c,255)
```

(b) D, a robust program.

```
i=1;
while (2^16-i != 0)
i =i+2;
```

(c) C', a non efficient program

```
x=?; r=?; g=?; b=?;
if (r+g-1 != 0)
  col = 2^r + 2^g + 2^b
else col=0;
write(x,col)
```

(d) E, a program usable by daltonians

2003), or it is "*the ability of a software system to place as few demands as possible on hardware resources, such as processor time or space occupied*" (Meyer, 1997). Once again, such definitions suffer from the ambiguity of the natural language, e.g., it is not clear if when verifying efficiency requirements the underlying architecture must be considered or not, or if space and time requirements must be considered independently or not.

Formal Definitions

Efficiency can be formally defined as an abstraction of the execution traces of a program. As such behavior depends on the underlying architecture, our definition explicitly mentions the architecture in which the program is executed. Efficiency requirements can be specified by natural numbers, standing, for instance, for the number of processor cycles or the size of the heap. As a consequence our abstract domain will be set of natural numbers with the usual total order, $\langle \mathsf{N}, \leq \rangle$.

We distinguish between efficiency in time and space. The first one corresponds to the length of a trace, i.e. the number of transitions for executing the system, and the second one to the size of the environment, i.e. the maximum quantity of memory allocated during program execution. It is worth noting that the following definitions are well-formed as we consider partial execution traces, i.e., (possible infinite) sets of finite traces. Recall that Ω denotes an uninitialized variable.

For Time Efficiency, let C be a program, A an architecture, $\mathrm{length} \in [\mathsf{P}(\Sigma^*) \to \mathsf{N}]$ be the length of a trace, and $(\langle \mathsf{P}(\Sigma^*), \subseteq \rangle, \alpha_t, \gamma_t, \langle \mathsf{N}, \leq \rangle)$ be a Galois connection where

$$\begin{aligned} \alpha_t &= \lambda T.\ \sup(\{\mathrm{length}(\tau) \mid \tau \in T\}) \\ \gamma_t &= \lambda n.\ \{\tau \in \mathsf{P}(\Sigma^*) \mid \mathrm{length}(\tau) \leq n\}. \end{aligned}$$

We say that the system $\langle A, C \rangle$ respects the time requirement k if

$$\forall S_0 \subseteq \Sigma.\ \alpha_t(s\langle \mathrm{A},\mathrm{C} \rangle(S_0)) \leq k.$$

For Space Efficiency, let C be a program, A an architecture, $\mathrm{size} \in [\mathsf{P}(\Sigma) \to \mathsf{N}]$ be the function defined as

$$\mathrm{size} = \lambda\sigma.\ \#\{\mathrm{x} \in \mathrm{Vars} \mid \sigma(\mathrm{x}) \neq \Omega\},$$

and $(\langle \mathsf{P}(\Sigma^*), \subseteq \rangle, \alpha_s, \gamma_s, \langle \mathsf{N}, \leq \rangle)$ be a Galois connection where

$$\begin{aligned} \alpha_s &= \lambda T.\ \max_{\tau \in T}\{\mathrm{size}(\sigma) \mid \sigma \in \tau\} \\ \gamma_s &= \lambda n.\ \{\tau \in \mathsf{P}(\Sigma^*) \mid \forall \sigma \in \tau.\ \mathrm{size}(\sigma) \leq n\} \end{aligned}$$

We say that the system $\langle \mathrm{A},\mathrm{C} \rangle$ respects the space requirement k if

$$\forall S_0 \subseteq \Sigma.\ \alpha_s(s\langle \mathrm{A},\mathrm{C} \rangle(S_0)) \leq k.$$

Abstractions

In order to automatically verify time requirements, we must find an upper bound to the number of transitions performed during the execution of a system. Once again, we can do it by using a numerical abstract domain. In fact, we can endow a concrete state σ with a (hidden) variable time, to be incremented at each transition (Halbwachs, 1979). Then, the values taken by time will be upper-approximated in the numerical domain, say by $\overline{\mathrm{time}}$, so that the verification boils to check that $\overline{\mathrm{time}} \leq k$. In the same way, the verification of space requirements can be obtained by abstracting a state with the number of variables different from Ω it contains. The approach can be generalized to more complex languages, e.g., a language with recursive functions. In this case, the stack will be approximated by its height.

In our approach, verification of time and space efficiency requirements can be easily combined by considering the reduced product of the two abstract domains (Cousot & Cousot, 1977).

Example

Let us consider the programs C and C' in Algorithm 1, an architecture A, where the multiplication is a primitive operation, and an architecture A' where the multiplication is implemented as a sequence of additions, e.g., $\mathrm{i{=}i{*}4}$ becomes $\mathrm{i{=}i + i;i{=}i + i}$. Using the analyzer described in (Cousot, 1999), we can infer:

$$
\begin{aligned}
\overline{s}\langle \mathrm{A,C}\rangle(\langle \mathrm{i} \mapsto [-\infty,+\infty], \mathrm{time} \mapsto \underline{0}\rangle) &= \langle \mathrm{i} \mapsto [1,2^{16}], \mathrm{time} \mapsto [0,9]\rangle \\
\overline{s}\langle \mathrm{A',C}\rangle(\langle \mathrm{i} \mapsto [-\infty,+\infty], \mathrm{time} \mapsto \underline{0}\rangle) &= \langle \mathrm{i} \mapsto [1,2^{16}], \mathrm{time} \mapsto [0,25]\rangle, \\
\overline{s}\langle \mathrm{A,C'}\rangle(\langle \mathrm{i} \mapsto [-\infty,+\infty], \mathrm{time} \mapsto \underline{0}\rangle) &= \langle \mathrm{i} \mapsto [0,2^{16}], \mathrm{time} \mapsto [0,32769]\rangle.
\end{aligned}
$$

Observe that the results above can be used for comparing different programs on different architectures.

Robustness

Informal Definition

Robustness, or dependability, for (Ghezzi, Jazayeri & Mandrioli, 2003) is "*the ability of a program to behave reasonably, even in circumstances that were not anticipated in the specifications*", for (Meyer, 1997) is "*the ability of software systems to react appropriately to abnormal conditions*", and for (Kotonya, & Sommerville, 1998) is "*the time to restart after failure*". Once again, the three definitions are not rigorous enough: the first definition does not specify what is a reasonable behavior, the second one does not specify what is an abnormal condition, and the latter has implicit the strong assumption that all possible failures are considered.

Formal Definition

A software is robust, if any exception raised during its execution, in any architecture and with any initial state, is caught by some exception handler. We recall that exceptions can be raised either by the architecture, e.g., division-by-zero, or by the software itself. As a consequence, a robust program never terminates in an exceptional state.

Let C be a program, and let

$$(\langle \mathsf{P}(\Sigma^{*}), \subseteq\rangle, \alpha_d, \gamma_d, \langle \mathsf{P}(\Sigma), \subseteq\rangle)$$

be a Galois connection where

$$
\begin{aligned}
\alpha_d &= \lambda T.\ \{\sigma_n \mid \sigma_0 \dots \sigma_n \in T\} \\
\gamma_d &= \lambda S.\ \{\sigma_0 \dots \sigma_{n-1}\sigma_n \mid \forall i \in [0, n-1].\sigma \in \Sigma \wedge \sigma_n \in S\}.
\end{aligned}
$$

We say that a system is robust if for all the architectures A,

$$\forall S_0 \in \mathsf{P}(\Sigma).\ \alpha_d(s\langle \mathrm{A,C}\rangle(S_0)) \cap \mathrm{Exceptions} \times \mathrm{Env} = \varnothing.$$

Abstraction

Robustness can be checked either by considering an abstract domain for inferring the uncaught exceptions (Pessaux & Leroy, 2000), or by considering an abstract domain for reachability analysis (Cousot, 1999). In the first case, a program is robust if the analysis reports that no exception can be raised; in the latter, a program is robust if the analysis reports that the lines of code that may raise an exception (e.g., with a throw statement) are never reached.

Example

Let us consider the program D of Algorithm 1(b). An interval analysis determines that when the true-branch of the if statement is taken, i is different from zero, so that the $\mathrm{MathErr}$ exception cannot be raised. In the other case, the exception Err is raised and then it is also caught. As a consequence, D is robust with respct to the chosen abstraction.

Of course, this is just one possible formalization of Robustness. For instance, the system also might know what to do in case any of the possible exceptions happens, and do so while continuing in an operational state. This can be formalized by abstracting properties of the subtraces that originate from a catch statement.

Usability/Accessibility

Informal Definition

The definition of usability (or accessibility) is probably the most contrived one. The definition in (Ghezzi, Jazayeri & Mandrioli, 2003) says that "*software system is usable [...] if its human users find it easy to use*", whereas (Meyer, 1997) talks about ease of use as "*the ease with which people of various backgrounds [...] can learn to use software*" and (Kotonya, & Sommerville, 1998) defines it in function of other, undefined, basic concepts as "*learnability, satisfaction, memorability*".

Formal Definition

In our setting, usability is a abstraction of the output stream that is preserved when a given property, depending on the particular user, is considered. For instance, an abstraction that considers the colors of the output characters can be used to verify if a system is usable for daltonians. We need some auxiliary definitions. Output streams belong to the set Stdout. Given a state $\sigma \in \Sigma$, the function $\mathrm{out} \in [\Sigma \to \mathrm{Stdout}]$ is such that $\mathrm{out}(\sigma)$ is the output stream in the state σ.

Let C be a program, A an architecture, $(\langle \mathsf{P}(\Sigma^*), \subseteq\rangle, \alpha_\Sigma, \gamma_\Sigma, \langle \mathsf{P}(\mathrm{Stdout}), \subseteq\rangle)$ be a Galois connection where

$$\begin{aligned}\alpha_\Sigma &= \lambda T.\ \{\mathrm{out}(\sigma) \in \Sigma \mid \exists \tau \in T.\ \sigma \in T\}\\ \gamma_\Sigma &= \lambda O.\ \{\tau \in \Sigma^* \mid \forall \sigma \in \tau.\ \exists o \in O.\ \mathrm{out}(\sigma) = o\},\end{aligned}$$

let $(\langle \mathsf{P}(\mathrm{Stdout}), \subseteq\rangle, \alpha, \gamma, \langle \bar{\mathrm{D}}, \leq\rangle)$ be a Galois connection, and let $\mathrm{p} \in \bar{\mathrm{D}}$. We say that the system $\langle \mathrm{A,C}\rangle$ is usable w.r.t. the observable p if

$$\forall S_0.\ \alpha(\alpha_\Sigma(s\langle \mathrm{A,C}\rangle)(S_0)) \leq \mathrm{p}.$$

Abstraction

The definition above can be instantiated to consider the usability of a system for daltonians, i.e., people afflicted by red/green color blindness. In fact, the colors of the output stream can be abstracted in order to collapse together colors indistinguishable by daltonians. As colors are represented by integers in the RGB color system, numerical abstract domains can be used to automatically check properties on colors.

Example

Let us consider the program E in Algorithm 1(d), an architecture where the input stream is a sequence of 0/1 digits, and colors are represented as in RGB schema using 3 bits, i.e. colors range between 0 (black) and 7 (white). Using the static analyzer of (Cousot, 1999) instantiated with the Intervals abstract domain, and refined with trace partitioning (Handjieva & Tzolovski, 1998), one infers that

$$\begin{aligned}\bar{s}\langle \mathrm{A,E}\rangle(\ &\langle\!\langle \mathrm{x} \mapsto [0,1], \mathrm{r,g,b} \mapsto [0,1]\rangle\!\rangle)\\ &= (\langle\!\langle \mathrm{x} \mapsto [0,1], \mathrm{r,g,b} \mapsto [0,1], \mathrm{col} \mapsto [0,1] \cup [6,7]\rangle\!\rangle),\end{aligned}$$

so that as col is always in the set of the colors distinguishable by daltonians (i.e. $\{$ black, blue, yellow, white$\}$), E respects the usability specification.

Other Non-Functional Requirements

We showed how four typical non-functional requirements can be encapsulated in our framework. This approach based on preservation of a property up to a given observation, can be easily generalized to other product non-functional requirements. For instance, *upgrade* means that when a new program N, replaces a program O on a given architecture A, then the observed behavior is preserved: $\alpha(s\langle A,N\rangle) \leq \alpha(s\langle A,O\rangle)$. Similarly, if *compatibility* is a property specified by an abstract element c, then we say that two programs P and P' are compatible w.r.t. c if $\alpha(s\langle A,P\rangle) \leq c$ and $\alpha(s\langle A,P'\rangle) \leq c$.

Non Functional Requirements in Service Oriented Scenarios

In the previous sections, we discussed how Abstract Interpretation theory can be used as a provably sound way to model non functional software requirements and to support their automatic validation. One may argue that the we dealt mainly with a one computer-one program view with statically decidable semantics, which in today's world of dynamic, complex and distributed systems seems a limited domain. In Service Oriented Scenarios, systems interact by asking/providing services, that are expressed in terms of functional and non functional requirements that should satisfy suitable Service Levels Agreements. The adeguacy of the functionalities with respect to the client requirements can be formally verified through type matching techniques.

On the other hand, the approach we advocate in this chapter can be used to verify the adeguacy of non-functional properties of Service Oriented Applications. In fact, one can think of expression non-functional properties in a suitable specification language, and use abstract interpretation-based static analysis tools to verify the conformance of the implementation or of the model with the specification. This is for instance the key idea behind the design and the development of a language agnostic abstract interpretation-based static contract analyzer and checker for .NET, whose static checker can be downloaded as part of the Code Contracts in DevLabs (http://msdn.microsoft.com/en-us/devlabs/dd491992.aspx).

CONCLUSION

Recent very encouraging experiences show that abstract interpretation-based static program analysis can be made efficient and precise enough to formally verify a class of properties for a family of large programs with few or no false alarms, also in case of critical embedded systems (Blanchet, Cousot, Cousot, et al., 2003). We strongly believe that also the treatment of non functional requirements can well fit in this picture. For instance, recent works on Security analysis through Abstract Interpretation (Zanioli & Cortesi, 2011), focussing on information leakage detection, may be seen as another instance of the framework presented in this chapter.

As already mentioned, the key issue (which deserves to be investigated end experimentally validated) is finding appropriate formalizations of the properties to be "observed", by re-using or designing new numerical domains (e.g., this would be the case when approaching requirements like Effectiveness, Availability, Response time, etc.) or logical domains keeping track of dependencies (e.g., for Modifiability of Privacy), or categorical symbolic domains (e.g., for Testability or Supportability).

REFERENCES

Abrial, J.-R. (1996). *The B-Book: Assigning programs to meanings*. Cambridge University Press. doi:10.1017/CBO9780511624162

Abrial, J.-R. (2003). B#: Toward a synthesis between Z and B. In Bert, D., Bowen, J. P., King, S., & Waldén, M. (Eds.), *ZB'2003 - Formal Specification and Development in Z and B, Lecture Notes in Computer Science* (pp. 168–177). Turku, Finland: Springer. doi:10.1007/3-540-44880-2_12

Blanchet, B., Cousot, P., Cousot, R., Feret, J., Mauborgne, M., & Miné, A.…. Rival, X. (2003). A static analyzer for large safety-critical software. *Proceedings of the 2003 ACM Conference on Programming Language Design and Implementation (PLDI'03)*, (pp. 196-207). ACM Press.

Boehm, B. W. (1976). Software engineering. *IEEE Transactions on Computers*, 1266–41.

Brauburger, J. (1997). Automatic termination analysis for partial functions using polynomial orderings. In P. Van Hentenryck (Ed.), *Proc.4 th Int. Symp.SAS '97, Paris, Lecture Notes in Computer Science*: *Vol. 1302,* (pp. 330--344).

Cortesi, A., & Logozzo, F. (2005). Abstract interpretation-based verification of non functional requirements. *Proceedings of the 7th International Conference on Coordination Models and Languages, LNCS 3654, (COORD'05)*, (pp. 49-62). Berlin, Germany: Springer-Verlag.

Cousot, P. (1997). Types as abstract interpretations, invited paper. *24th ACM Symposium on Principles of Programming Languages (POPL '97)*, (pp. 316-331). ACM Press.

Cousot, P. (1999). The calculational design of a generic abstract interpreter. In Broy, M., & Steinbrüggen, R. (Eds.), *Calculational system design. NATO ASI, Series F*. Amsterdam, The Netherlands: IOS Press.

Cousot, P. (2001). Abstract interpretation based formal methods and future challenges, invited paper. In Wilhelm, R. (Ed.), *Informatics - 10 years back, 10 years ahead, Lecture Notes in Computer Science* (pp. 138–156). Springer-Verlag.

Cousot, P., & Cousot, R. (1977). Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. *4th ACM Symposium on Principles of Programming Languages (POPL '77)*, (pp. 238-252). ACM Press.

Cousot, P., & Cousot, R. (2002). Systematic design of program transformation frameworks by abstract interpretation. *29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '02)*, (pp. 178-190). New York, NY: ACM Press.

Cousot, P., & Cousot, R. (2004). An abstract interpretation-based framework for software watermarking. *Conference Record of the Thirtyfirst Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Venice, Italy, (pp. 173--185). New York, NY: ACM Press.

Cousot, P., & Halbwachs, N. (1978). Automatic discovery of linear restraints among variables of a program. *5th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '78)*, (pp. 84-97). ACM Press.

Davis, A. (1992). *Software Requirements: Objects, functions and states*. Prentice Hall.

Deutsch, M. S., & Willis, R. R. (1988). *Software quality engineering*. Prentice-Hall.

Ghezzi, C., Jazayeri, M., & Mandrioli, D. (2003). *Foundamentals of software engineering* (2nd ed.). Prentice Hall.

Halbwachs, N. (1979). Determination Automatique de Relations Lineaires Verifees par les Variables d'un Programme. *These de 3eme cycle d'informatique, Universite scientifique et medicale de Grenoble*.

Halbwachs, N. (1995). Non-functional requirements in the software development process. *Software Quality*, *5*(4), 285–294.

Halbwachs, N. (1998). About synchronous programming and abstract interpretation. *Science of Computer Programming*, *31*(1), 75–89. doi:10.1016/S0167-6423(96)00041-X

Handjieva, M., & Tzolovski, S. (1998). Refining static analyses by trace-based partitioning using control flow. *Proceedings of the Static Analysis Symposium (SAS '98), Lectures Notes in Computer Science,* (pp. 200-215). Springer-Verlag.

Hughes, S. (1992). Compile-time garbage collection for higher-order functional languages. *Journal of Logic and Computation*, *2*(4), 483–509. doi:10.1093/logcom/2.4.483

IEEE. (1988). *Recommended practice for software requirement specification*. IEEE Press.

Jones, N. D. (1997). Combining abstract interpretation and partial evaluation. In P. Van Hentenryck, (Ed.), *Static analysis, Proc. 4th Int. Symp SAS '97, Paris, Lecture Notes in Computer Science*, (pp. 396-405). Springer-Verlag.

Karlsson, J. (1997). Managing software requirements using quality function deployment. *Software Quality Control*, *6*(4), 311–326. doi:10.1023/A:1018580522999

Kotonya, G., & Sommerville, I. (1998). *Requirements engineering - Processes and techniques*. Wiley.

Kozen, D., Palsberg, J., & Schwartzbach, M. I. (1994). Efficient inference of partial types. *Journal of Computer and System Sciences*, *49*(2), 306–324. doi:10.1016/S0022-0000(05)80051-0

Lacan, P., Monfort, J. N., Ribal, L. V. Q., Deutsch, A., & Gonthier, A. (1998). The software reliability verification process: The Ariane 5 example. *Proceedings DASIA 98 – Data Systems In Aerospace,* Athens, GR. ESA Publications.

Le Métayer, D. (1996). Program analysis for software engineering: New applications, new requirements, new tools. *ACM Computing Surveys*, *28*(4es), 167. doi:10.1145/242224.242435

Meyer, B. (1997). *Object-oriented software construction. Professional technical reference* (2nd ed.). Prentice Hall.

Miné, A. (2001). *The octagon abstract domain. AST 2001 in WCRE 2001* (pp. 310–319). IEEE CS Press.

Monniaux, D. (2003). Abstracting cryptographic protocols with tree automata. *Science of Computer Programming*, *47*(2-3), 177–202. doi:10.1016/S0167-6423(02)00132-6

O'Brien, L., Merson, P., & Bass, L. (2007). Quality attributes for service-oriented architectures. *Proceedings of the International Workshop on Systems Development in SOA Environments (May 20 - 26, 2007)*, Washington, DC: IEEE Computer Society.

Pessaux, F., & Leroy, X. (2000). Type-based analysis of uncaught exceptions. *TOPLAS*, *22*(2), 340–377. doi:10.1145/349214.349230

Sommerville, I. (2000). *Software engineering* (6th ed.). Addison Wesley.

Spivey, J. M. (1992). *The Z notation*. Prentice Hall.

Springer Verlag.

Traub, K. R., Culler, D. E., & Schauser, K. E. (1992). Global analysis for partitioning non-strict programs into sequential threads. *ACM LISP Pointers*, *5*(1), 324–334. doi:10.1145/141478.141568

Wada, H., Suzuki, J., & Oba, K. (2006). Modeling non-functional aspects in service oriented architecture. *SCC '06: Proceedings of the IEEE International Conference on Services Computing*, (pp. 222-229). Washington, DC: IEEE Computer Society.

Zanioli, M., & Cortesi, A. (2011). Information leakage analysis by abstract interpretation. Proceedings of the 37th International Conference on Current Trends in Theory and Practice of Computer Science. *Lecture Notes in Computer Science*, *6543*, 545–557. doi:10.1007/978-3-642-18381-2_45

KEY TERMS AND DEFINITIONS

Abstract Interpretation: A theory of semantics approximation for computing conservative over-approximations of dynamic properties of programs.

Abstract System: An abstraction of an architecture and of programs running on it.

Portability: Set of properties of the execution of a program that are preserved when it is ported on different architectures.

Space Efficiency: Maximum quantity of memory allocated during program execution.

Software Robustness: A software is robust, if any exception raised during its execution, in any architecture and with any initial state, is caught by some exception handler.

Software Usability: Set of properties of the program's user interface that are preserved during program execution.

ENDNOTE

This chapter is an extended and revised version of (Cortesi & Logozzo, 2005).