

Architectural Stability

Rami Bahsoon

School of Engineering and Applied Science, Computer Sc, Aston University Birmingham, Aston Triangle, B4 7ET, Birmingham, UK

Wolfgang Emmerich

London Software Systems, Dept of Computer Sc., University College London, WC1E 6BT, London, UK

Abstract

Evolution is a key problem in software engineering and exacts huge costs. In managing change and facilitating evolution, considerable emphasis is placed on the architecture of the software system as a key artifact involved. One of the major indicators of the success (failure) of software evolution is the extent to which the software system can endure changes in requirements, while leaving the architecture of the software system intact. We refer to the presence of this “intuitive” phenomenon as architectural stability. Ongoing research on relating requirements to software architectures has considered the architectural stability problem as an open research challenge. This is because the conflict between requirements volatility and architectural stability is a difficult one to handle. Hence, the concept is still far from being understood and many architectural stability related questions are remained unanswered. Reflecting on our extensive research into the problem, we define architectural stability and explore perspectives in handling the problem. We review existing research effort and discuss their limitations. We derive a research agenda and outline research challenges and opportunities. The implications of such contribution need not be overstated: advancing the understanding of the architectural stability, stimulating and possibly motivating future research in architectural stability and related problems.

Keywords

Relating requirements to architectures, architectural stability, requirements evolution, software architectures for maintenance and evolution, economics of software maintenance, software change impact analysis

Paper Type

Research

Correspondence to:

Rami Bahsoon

School of Engineering and Applied Science, Computer Sc, Aston University Birmingham, Aston Triangle, B4 7ET, Birmingham, UK. E-mail: r.bahsoon@aston.ac.uk

Phone: +44 (0) 121 204 3464

fax: +44(0) 121 204 3681

Architectural Stability

Rami Bahsoon

School of Engineering and Applied Science, Computer Sc,
Aston University Birmingham,
Aston Triangle, B4 7ET, Birmingham, UK
r.bahsoon@aston.ac.uk

Wolfgang Emmerich

London Software Systems, Dept. of Computer Science,
University College London,
Gower Street, WC1E 6BT, London, UK
w.emmerich@cs.ucl.ac.uk

Abstract

Evolution is a key problem in software engineering and exacts huge costs. In managing change and facilitating evolution, considerable emphasis is placed on the architecture of the software system as a key artifact involved. One of the major indicators of the success (failure) of software evolution is the extent to which the software system can endure changes in requirements, while leaving the architecture of the software system intact. We refer to the presence of this “intuitive” phenomenon as architectural stability. Ongoing research on relating requirements to software architectures has considered the architectural stability problem as an open research challenge. This is because the conflict between requirements volatility and architectural stability is a difficult one to handle. Hence, the concept is still far from being understood and many architectural stability related questions are remained unanswered. Reflecting on our extensive research into the problem, we define architectural stability and explore perspectives in handling the problem. We review existing research effort and discuss their limitations. We derive a research agenda and outline research challenges and opportunities. The implications of such contribution need not be overstated: advancing the understanding of the architectural stability, stimulating and possibly motivating future research in architectural stability and related problems.

1. Introduction

Software requirements, whether functional or non-functional, are generally volatile; they are likely to change and evolve over time. The change is inevitable as it reflects changes in stakeholders’ needs and the environment in which the software system works. *Software architecture* is the earliest design artifact, which realizes the requirements of the software system. It is the manifestation of the earliest design decisions, which comprise the architectural structure (i.e., components and interfaces), the architectural topology (i.e., the architectural style), the architectural infrastructure (e.g., the middleware), the relationship among them, and their relationship to the other software artifacts (e.g., low-level design) [Bah05a]. One of the major implications of a software architecture is to render particular kinds of changes easy or difficult, thus constraining the software’s evolution possibilities [Jaz02]. A change may “break” the software architecture necessitating changes to the architectural structure (e.g., changes to components and interfaces), architectural topology, or even changes to the underlying architectural infrastructure. It may be expensive and difficult to change the architecture as requirements evolve

[Fin00b]. Conversely, failing to accommodate the change leads ultimately to the degradation of the usefulness of the system. Hence, there is a pressing need for flexible software architectures that tend to be *stable* as the requirements evolve. By a stable architecture, we mean the extent to which a software system can endure changes in requirements, while leaving the architecture of the software system intact. We refer to the presence of this “intuitive” phenomenon as *architectural stability*.

Developing and evolving architectures, which are *stable* in the presence of change and *flexible enough* to be customized and adapted to the changing requirements is one of the key challenges in software engineering [Gar00]. Ongoing research on relating requirements to software architectures has considered the architectural stability problem as an open research challenge [van00; Nus01]. This is because the conflict between requirements volatility and architectural stability is a difficult one to handle [van00]. As a result, many architectural stability related questions are remained unanswered [Nus01]: For example, what software architectures (or architectural styles) are stable in the presence of the changing requirements, and how do we select them? What kinds of changes are systems likely to experience in their lifetime, and how do we manage requirements and architectures (and their development processes) in order to manage the impact of these changes? Meanwhile, industrial evidence reveals situations where high requirements volatility is the norm and much of the promise is leaved to the architecture in accommodating the changes. For example, the number of mergers between companies is increasing and this trend is bound to continue. The different divisions of a newly merged company have to deliver unified services to their customers and this usually demands an integration of their IT systems into the core architecture. The time frame is often so short that building a new system is not an option and therefore existing system components have to be integrated into a distributed system architecture to appear as an integrated computing facility. Secondly, the trend of providing new services or evolving existing services to target new customers, devises and platforms, and distribution settings (e.g., mobility setting) is increasing. For example, moving from a fixed distributed setting to mobility carries critical changes, mainly to non-functionalities, such as changes in availability, security, and scalability requirements. Often the core “fixed” architecture falls short in accommodating the requirements; henceforth, changes to the architecture becomes necessary. Thirdly, it is often the case that components are procured off-the-shelf, rather than built from scratch, in response to changes in requirements and then need to be integrated into the core

architecture. These components often have incompatible requirements on the hardware and operating system platforms they run on. In many software systems, the architecture is the level that has the greatest inertia when external circumstances change and consequently incurs the highest maintenance costs when evolution becomes unavoidable [Coo01]. Hence, a stable architecture which addresses such changes in requirements within limited resources and shorter time-to-market is a significant asset for surviving the business, cutting down maintenance costs and creating value.

Reflecting on our extensive research into the problem, we define architectural stability and explore perspectives in handling the problem. We review existing research effort and discuss their limitations. We derive a research agenda and outline research challenges and opportunities. The implications of such contribution are numerous: exploring the architectural stability problem, advancing the understanding of the said problem, stimulating and possibly motivating future related research. This may result in frameworks with practical implications aimed at developing stable systems in face of changes in requirements.

The paper is further structured as follows. Section 2 looks at architectures and evolution. Section 3 explores perspectives in handling the architectural stability problem. Section 4 outlines research challenges and opportunities. Section 5 concludes.

2. Architecture-Centric Evolution

In Lehman's terminology [FEA], there are two types of systems: these are E-type systems and S-type systems. E-Type systems that are embedded in real world applications and are used by humans for everyday business functions. Examples might be customer service, order entry, payroll, operating systems, databases engines. S-Type systems are executable models of a formal specification. The success of this software is judged by how well it meets the specification. For E-Type systems the "real world" is dynamic and ever changing. As the real world changes the specification changes and the E-Type systems need to adapt to these changes. Hence, E-Type systems tend to be evolvable. For S-Type systems the specification becomes invalid in the presence of change. In this paper, we deal with evolution and architectural stability of E-type systems.

Change is a process that either introduces new requirements into an existing system; modifies the system if the requirements were not correctly implemented; or moves the system into a new operating environment [Ben00]. In software engineering, it has been known that focusing the change on program code leads to loss of structure and maintainability [Ben00]. Upon managing the change of requirements considerable emphasis is thus placed on the architecture of the software system as the key artifact involved [Gar00]. *Architecture-centric evolution* approaches pursue the software architecture as the appropriate level of abstraction for reasoning about, managing and guiding the evolution of complex software systems, and "synchronizing" the software requirements with its detailed design and implementation. A distinctive feature of these approaches is that they explicitly account for the non-functional requirements, the so-called quality attributes. As the quality attributes comprise the most substantial properties of the system, the evolution of such properties can be best reasoned about and managed at the architectural level. For exam-

ple, the current trend is to build distributed systems architectures with middleware technologies such as Java 2 Enterprise Edition (J2EE) and the Common Object Request Broker Architecture (CORBA), resulting in the so-called middleware-induced architectures [DiN99]. Middleware-induced architectures follow an architectural-centric evolution approach, as the emphasis is placed on the induced architecture for simplifying the construction of distributed systems by providing high-level primitives, which shield the application engineers from the distribution complexities, managing systems resources, and implementing low-level details, such as concurrency control, transaction management, and network communication. These primitives are often responsible for realizing many of the non-functional requirements (e.g., scalability, fault tolerance, etc.) in the architecture of the system induced and facilitating their evolution over time. Another example is from product-line architectures. Product lines, a family of products sharing the same architecture, inherently require domain-specific variation and evolution of various products. Due to the higher level of interdependency between the various software artifacts in a product-line, software evolution is too complex to be dealt with at the code level. An essential property of these architectures is that they should be stable over the projected life of the system [Bah05a]. As the focus is on the architecture for "easing" evolution, architecture-centric evolution approaches place considerable emphasis on the *flexibility* of the architecture in responding to change, deriving business value, and minimizing evolution costs [Bah05a, Bah06].

3. Perspectives into Architectural Stability

In subsequent sections, we explore perspectives in looking at the architectural stability problem. We review existing research effort and discuss their limitations.

3.1 Requirements Engineering Perspective

Ongoing research on relating requirements to software architectures has considered the architectural stability problem as an open research challenge and difficult to handle [Nus01; van01; Emm02]. [Nus01] proposed the "Twin Peaks" model, a partial and simplified version of the spiral model. The cornerstone of this model is that a system's requirements and its architecture are developed concurrently; that is, they are "inevitably intertwined" and their development is interleaved. [Nus01] advocated the use of various kinds of patterns – requirements, architectures, and designs- to achieve the model objectives. As far as architectural stability is concerned, Nuseibeh had only exposed a tip of the "iceberg": development processes that embody characteristics of the Twin Peaks are the first steps towards developing architectures that are stable in the face of inevitable changes in requirements. Nuseibeh noted that many architectural stability related questions are difficult and remain unanswered. Examples include: what software architectures (or architectural styles) are stable in the presence of changing requirements, and how do we select them? What kinds of changes are systems likely to experience in their lifetime, and how do we manage requirements and architectures (and their development processes) in order to manage the impact of these changes?

With the motivation of bridging the gaps between requirements and software architectures, [van00] noted that the goal-oriented approach to requirements engineering may support building and evolving software architectures guaranteed to meet its functional and non-functional requirements. As far as the architectural stability problem is concerned:

“Even though streamlined derivation processes may be envisaged for architectural development, things get much more complicated for evolution. For example, the conflict between requirements volatility and architectural stability is a difficult one to handle”. [van00]

In summary, these positions have reflected on open challenges and possible strategies in developing software architectures that need to be stable as requirements evolve. They have highlighted the architectural stability problem from a requirements perspective. Focused research attempts, however, have not followed these lines.

In [Emm02; Bah05b], we reflected on the architectural stability problem with a particular focus on developing distributed software architectures induced by middleware. We advocated adjusting requirements elicitation and management techniques to elicit not just the current non-functional requirements, but also to assess the way in which they will develop over the lifetime of the architecture. These ranges of requirements may then inform the selection of distributed components technology, and subsequently the selection of application server products. Specifically, we considered the architectural stability problem from the distributed components technology in the face of changes in *non-functional requirements*. We argued that addition or changes in functional requirements could be easily addressed in distributed component-based architectures by adding or upgrading the components in the business logic. However, changes in non-functional requirements are more critical; they can stress architecture considerably, leading to architectural “breakdown”. Such a “breakdown” often occurs at the middleware level and is due to the incapability of the middleware to cope with the change in non-functional requirements (e.g., increased load demands). This may drive the architect/developer to consider ad-hoc or proprietary solutions to realize the change, such as modifying the middleware, extending the middleware primitives, implementing additional interfaces, etc. Such solutions could be costly and unacceptable [Bah05b].

3.2. A Value-Driven Design Perspective

An established route to manage the change and facilitate evolution is a universal “design for change” philosophy, where the architecture is conceived and developed such that evolution is possible [Par79]. Parnas’s notion of the “design for change” is based on the recognition that much of the total lifecycle cost of a system is expended in the change and incurred in evolution. A system that is not designed for evolution will incur tremendous costs, which are disproportionate to the benefits. For a system to create value, the cost of a change increment should be proportional to the benefits delivered [Par72]. “Design for change” is thus promoted as a value-maximizing strategy provided one could anticipate changes [Boe00]. The “Design for change” philosophy is believed to be a useful heuristic for

developing flexible architectures that tend to be stable as requirements evolve. However, the challenge is that there is a general lack of adequate models and methods, which connect this technical engineering philosophy to value creation under given circumstances [Boe00]. [Sul97] note that the important book of Shaw and Garlan on software architecture begins, “As the size and complexity of software systems increase, the design and specification of overall system structure become more significant issues than the choice of algorithms and data structures...” [Sha96]. [Sul97] add, “This statement is true, without a doubt. The problem in the field is that no serious attempt is made to characterize the link between structural decisions and value added”. That is, the traditional focus of software architecture is more on structural and technical perfection than on value. In addressing the architectural stability problem, linking structural decisions to future value becomes more necessary, as presumably evolvable but stable architecture should add value to the system that outweigh what expended in designing for change, as the change materializes [Bah05a]. Furthermore, from an economic perspective, the change in requirements is a source of uncertainty that confronts an architecture during the evolution of the software system. The change places the investment in a particular architecture at risk. Conversely, designing for change incurs upfront costs and may not render future benefits. The benefits are uncertain, for the demand and the nature of the future changes are uncertain. The worthiness of designing or re-engineering an architecture for change should involve a tradeoff between the upfront cost of enabling the change and the future value added by the architecture, if the change materializes. The value added, as a result of enabling the change on a given architecture, is a powerful heuristic which can provide a basis for analyzing: (i) the worthiness of designing for change, (ii) the worthiness of re-engineering the architecture, (iii) the retiring and replacement decisions of the architecture or its associated design artifacts, (iv) the decisions of selecting an architecture, architectural style, middleware, and/or design with desired stability requirements, and/or (v) the success (failure) of evolution [Bah05a]. In ArchOptions [Bah05a; Bah04a; Bah03c; Bah03b], we have taken an economics-driven software engineering perspective [EDS99-05] to evaluate architectural stability using real options theory. We have adopted the view that software design and engineering activity is one of investing valuable resources under uncertainty with the goal of maximizing the value added [Boe00]. In particular, we have viewed evolving software as a value-seeking and value-maximizing activity: software evolution is a process in which software is undergoing an incremental change and seeking value [Bah04a; Bah05a]. We attribute the added value to the *flexibility* of the architecture in enduring changes in requirements. Means for achieving flexibility are typical architectural mechanisms or strategies that are built-in or adapted into the architecture with the objective of facilitating evolution and future growth. This could be in response to changes in functional (e.g., changes in features) or non-functional requirements (e.g., changes in scalability demands). As we are assuming that the added value is attributed to flexibility, arriving at a “more” stable software architecture requires finding an architecture which maximizes the yield in the embedded or the adapted *flexibility* in an architecture relative to the likely changing requirements [Bah05a;

Bah04a; Bah03c]. Optimally, a stable architecture is an architecture that shall add value to the enterprise and the system as the requirements evolve. By valuing the flexibility of an architecture to change, we have aimed at providing the architect/analyst with a useful tool for reasoning about a crucial but previously intangible source of value. This value can be then used for deriving “insights” into architectural stability and investment decisions related to evolving software. To value flexibility, we have contributed to a novel model, ArchOptions, which builds on an analogy with real options theory [Bah05a; Bah04a; Bah03c]. The model examines some critical likely changes in requirements and values the extent to which the architecture is flexible to endure these changes. The model views an investment in an architecture as an upfront investment plus “continual” increments of future investments in likely changes in requirements. We have applied ArchOptions to two architecture-centric evolution problems: assessing the worthiness of re-engineering a “more” stable architecture in face of likely changes in future requirements, where we have taken refactoring as an example of re-engineering [Bah04b]; and informing the selection of a “more” stable middleware-induced software architecture in the face of future changes in non-functional requirements [Bah05b]. Our perspective has provided a compromise through linking technical issues to value creation. The approach has the promise to provide insights and a basis for analyses to support many of the concerns highlighted in previous sections.

3.3. Architectural Evaluation Perspective

Evaluating architectural stability aims at assessing the extent to which the system of a given architecture is evolvable, while leaving the architecture and its associated design decisions unchanged as the requirements change. Approaches to evaluating software architectures for stability can be *retrospective* or *predictive* [Jaz00]. Both approaches start with the assumption that the software architecture’s primary goal is to facilitate the system’s evolution. Retrospective evaluation looks at successive releases of the software system to analyze how smoothly the evolution took place. Predictive evaluation provides insights into the evolution of the software system based on examining a set of *likely* changes and the extent to which the architecture can endure these changes.

Retrospective approaches. Jazayeri [Jaz02] motivated the use of retrospective approaches for evaluating software architectures for stability and looked at the problem from a software evolution perspective. His analyses rely on comparing properties from one release of the software to the next. The intuition is to see if the system’s architectural decisions remained intact throughout the evolution of the system, that is, through successive releases of the software. Jazayeri refers to this “intuitive” phenomenon as architectural stability. Retrospective analysis can be used for empirically evaluating an architecture for stability; calibrating the predictive evaluation results; and predicting trends in the system evolution [Jaz02]. In other words, retrospective analysis can also provide a basis for predictive analysis. For example, previous evolution data of the system may be used to anticipate the resources needed for the next release of the system, or to identify the components most likely that require attention, need restructuring or

replacements, or to decide if it is time to entirely retire the system. In principle, predictive analysis and retrospective analysis should be combined. However, perfect predictive evaluations render retrospective analysis unnecessary [Jaz02].

Jazayeri’s approach uses simple metrics such as software size metrics, coupling metrics, and color visualization to summarize the evolution pattern of the software system across its successive releases. The evaluation assumes that the system already exists and has evolved. This approach is therefore tend to be unpreventive and unsuitable for early evaluation (unless the evolution pattern is used to predict the stability of the next release). The evaluation appears to be expensive (in the absence of dedicated tools), for it requires information to be kept for each release of the software. Such data could be available through configuration management repositories. Yet such data is not commonly maintained, analyzed, or exploited. Moreover, the problem of architectural stability is strategic in essence and not purely technical [Bah06]. Jazayeri addresses the problem from a purely technical perspective. It worth noting that ongoing projects at the S.E.A.L group of Gall at U. of Zurich, have contributed to interesting results on mining software repositories to visualize evolution data. Though the developed infrastructure and analyses have the potential to support retrospective analyses for stability, no explicit results have been reported yet in this direction.

Predictive approaches. Predictive approaches to evaluating architectural stability can be applied during the early stages of the development life cycle to predict threats of the change on the stability of the architecture. Unlike retrospective approaches, predictive approaches are *preventive*; the evaluation aims to understand the impact of the change on the stability of the architecture *if the likely* changes need to be accommodated, so corrective design measures can be taken. Therefore, in predictive approaches, the effort to evaluation is justified as the evaluation is generally cost effective, when compared to retrospective approaches. In ArchOptions[Bah03;Bah05], we examine a set of *likely* changes that are critical to the evaluation. We pursue scenarios as a possible solution to describe these changes. To link the likely future change in requirements to the architecture, we adopt Goal-Oriented Requirements Engineering (GORE) paradigm [Dar93], where the goals are extracted from scenarios [Ant97]. We then predict the extent to which the architecture can endure these changes through constructing and valuing call options for a given change [Bah05a; Bah04a; Bah03c].

A comprehensive survey [Bah03a] of architectural evaluation methods indicates that current approaches to architectural evaluation focus explicitly on construction and only implicitly, if not at all, on the phenomenon of software “evolution”. The survey includes representative methods like ATAM [Kaz98], SAAM [Kaz94], ARID [Kle00], PASA/SPE [Smi99], ABAS [Kle99], and CBAM [Kaz01]. These methods provide frameworks for software architects to evaluate architectural decisions with respect to quality attributes such as performance, security, reliability, and modifiability. Despite the concern with “change” and accommodating changes, none of these methods, addresses stability of an architecture over time. For example, ATAM and SAAM indicate places where the architecture fails to meet its modifiability requirements and in some

cases shows obvious alternative designs that would work better. When used for evaluating modifiability, the input to these methods consists of an enumerated set of stakeholders' scenarios that represent known or likely changes that the system will undergo in the future. These scenarios are prioritized and mapped onto the architecture representation. The activity of mapping indicates problem areas in the architecture: areas where the architecture is overly complex (e.g., if distinct scenarios affect the same component(s)) and areas where changes tend to be problematic (e.g., if a scenario causes changes to a large number of components). The approaches to evaluation involve "thought experiments", modelling, and walking-through scenarios that exemplify requirements, as well as assessment by experts who look for gaps and weaknesses in addressing modifiability based on their experience. However, these methods do not support their prediction with an analytical basis and rigorous models. When methods, such as SAAM and ATAM are used to analyze qualities that are related to change (such as modifiability), they do not predict and measure the capability of the architecture to withstand the change. This renders their predictive effectiveness myopic. Further, these methods have ignored any economic considerations, with CBAM [Kaz01] being the notable exception. The evaluation decisions using these methods tend to be driven by ways that are not connected to, and usually not optimal for value creation. Factors such as flexibility, time to market, cost and risk reduction often have high impact on value creation [Boe00]. Such ignorance is in stark contrast to the objective of architectural evaluation and where cost reduction, risk mitigation, and long-term value creation are among the major drivers behind conducting an evaluation for stability [Bah06]. Such provision is important for it assists the objective assessment of the lifetime costs and benefits of evolving software, and the identification of legacy situations, where a system or component is indispensable but can no longer be evolved to meet changing needs at economic cost [Coo01]. Interested reader may refer to [Bah06], where we have highlighted the requirements for evaluating architectural stability, which address the pitfalls in existing methods.

4. Architecture Stability: Challenges and Opportunities

Rapid technological advances and industrial evidence are showing that the architecture is creating its own maintenance, evolution, and economics problems. Part of the problem stems in (i) the rapid technological advancements where evolution is not limited to a specific domain but extends to "horizontally" cover several domains, (ii) the current practice in engineering requirements, which ignore the above, (iii) and the improper management of the evolution of these requirements and across different design artifacts of the software system. In subsequent sections, we highlight some open issues that future research may consider to address some architectural-centric software evolution problems. Addressing these questions may have a positive implication on understanding architectural stability.

4.1. Architectures Description Languages

Although software evaluation methods are typically human-centred, formal notations for representing and analyzing archi-

tectural designs, generically referred to as Architectures Description Languages (ADLs), have provided new opportunities for architectural analysis [Gar00] and validation. We briefly survey efforts on ADLs as they have implications for supporting the evaluation of software architectures. We explain how ADLs can be used to support the evaluation of software architectures in general and provide some insights on their use to evaluate architectural stability.

ADLs are languages that provide features for modelling a software system's conceptual architecture [Med97]. ADLs provide a concrete syntax and a conceptual framework for characterizing architectures [Gar97]. The conceptual framework typically subsumes the ADL's underlying semantic theory (e.g., CSP, Petri nets, finite state machines). A number of ADLs have been proposed for modelling architectures both within a particular domain and as general-purpose architecture modelling languages [Med97]. Examples are Aesop [Gar95], Darwin [Mag95; Mag96], MetaH [Vest96], C2 [Med96], Rapide [Luc95], Wright [All94], UniCon [Sha95], SADL [Mor95], and ACME [Gar97]. ADLs are often intended to model large, distributed, and concurrent systems. Evaluating the properties of such systems upstream, at the architectural level, can substantially lessen the costs of any errors. The formality of ADL renders them suitable for the manipulation by tools for architectural analysis. In the context of architectural evaluation, the usefulness of an ADL is directly related to the kind of analyses a particular ADL tends to support. The type of analyses and evaluation for which an ADL is well suited depends on its underlying semantic model. We refer to [Med97] to state few examples: Wright is based on CSP; it analyses individual connectors for deadlocks. MetaH and UniCon both support schedulability analysis by specifying non-functional properties, such as criticality and priority. SADL can establish relative correctness of two architectures with respect to a refinement map. Rapide's and C2's event monitoring and filtering tools also facilitate analysis of an architecture. C2 uses critics to establish adherence to style rules and design guidelines. Another aspect of analysis, that supports architectural evaluation, is enforcement of constraints. Parsers and compilers enforce constraints implicit in types, non-functional attributes, component and connector interfaces, and semantic models. Static and dynamic analyses are used. Static analysis verifies that all possible executions of the architecture description conform to the specification. Static analysis helps the developers to understand the changes that need to be made to satisfy the analysed properties. They span approaches such as reachability analysis [Hol91; Val91; God91], symbolic model checking [Bru90; McM93], flow equations, and data-flow analysis [Dwy94]. The applicability of such techniques to architecture descriptions has been demonstrated in [Nau97] using two static analysis tools. These tools are INCA [Cor95] and FLAVERS [Mas91; Dwy94]. Rapide [Luk95] provides a support to simulate the executions of the system. The simulation verifies that the traces of those executions conform to high-level specifications of the desired behavior. [All94] use the static analysis tool FDR [FS92] to prove freedom from deadlock as well as compatibility between the component and connectors in an architecture description. The term dynamic architectures denote that application's architecture evolves during runtime. Examples of analyses sup-

port for dynamic architectures include Darwin [Mag95], which provides a support to the analysis of distributed message-passing systems.

No notable research effort has explored the role of ADLs in supporting evaluating architectural stability. However, ADLs have the *potential* to support such evaluation. For instance comparing properties of ADL specifications for different releases of a software can provide insights on how the change(s) or the likely change(s) tends to threaten the stability of the architecture. This can be achieved by analyzing the parts of newer versions that represent syntactic and semantic changes. Moreover, the analysis can provide insights into possible architectural breakdown upon accommodating the change. For example, the analysis may show how the change may break the architectural topology (e.g., style) and/or the architectural structure (e.g., components, connectors, interfaces ect.). We note that ADLs have potential for performing retrospective evaluation for stability, where the evaluation can be performed at a correspondingly high level of abstraction. Hence, the evaluation may be relatively less expensive as when compared, for example, to the approach taken by [Jaz02], detailed in the previous section.

4.2. Coping with Rapid Technological Advancements and Changes in Domain

Assume that a distributed e-shopping system architecture which relies on a fixed network needs to evolve to support new services, such as the provision of mobile e-shopping. Moving to mobility, the transition may not be straightforward: the original distributed system's architecture may not be respected, for mobility poses its own non-functional requirements for dynamicity that are not prevalent in traditional distributed setting [Cap03]. Examples of these requirements include the need to react to frequent changes in the environment, such as change in location; resource availability; variability of network bandwidth; the support of different communication protocols; losses of connectivity when the host need to be moved; and so forth. These requirements may not be satisfied by the current fixed architecture, the built-in architectural caching mechanisms, and/or the underlying middleware. Replacement of the current architecture may be required.

The challenge is thus to cope with the co-evolution of both the architecture and the non-functional requirements as we change domains. This poses challenges in understanding the evolution trends of non-functional requirements; designing architectures, which are aware of how these requirements will change over the projected lifetime of the software system and tend to evolve through the different domains. From an economics perspective, such is necessary to reduce the future "switching cost", which could hinder the success of evolution. In this perspective, engineering requirements and designing architectures need to be treated as value-maximizing activities in which we can maximize the net benefits (or real options) by minimizing the future "switching costs" while transiting across different domains. This necessitates amending the current practice of engineering requirements and brings a need for methods and techniques, which explicitly model the domain, the "vertical" evolution of the software system within the do-

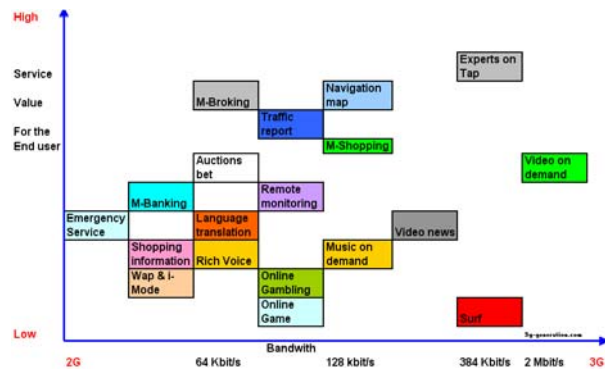
main itself and how the domain is likely to change over the projected lifetime of the software system. Again, goal-oriented requirements engineering(e.g.,[Dar93]) could be a promising starting point to "horizontally" capture the evolution across various domains and "vertically" across the domain itself. The problem of selecting an architecture, which tend to be stable as the "vertical" and the "horizontal" requirements evolve, become a multi-optimization design problem, where the selected architecture must maximize the value added relative to the "vertical" and the "horizontal" changes. The modeling could be then complemented by valuation frameworks which have the promise for answering questions of interest such as which architectural styles and middlewares, have the promise to reduce the switching costs and could prevail over the life time of the software system? This we believe is a practical need for engineering requirements to support stable software architectures. Managing the change is a process which involves recognizing the change through continued requirements elicitation, requirements evaluation of risk, and evaluation of systems in their operational environments [Nus00]. Identifying and documenting possible future changes is important in order to manage software evolution [Leh98] and evaluate architectural choices [Nus00].

In software engineering, the use of technology roadmapping, for example, is left unexplored in predicting and eliciting change in requirements. Technology roadmapping is an effective technology planning tool which help identifying product needs, map them into technology alternatives, and develop project plans to ensure that the required technologies will be available when needed [Sch99]. Technology roadmapping, as a practice, emerged from industry as a practical method of planning for new technology and product requirements. According to [Sch99], a roadmap is not a prediction of future breakthroughs in the technology, but rather an articulation of requirements to support future technical needs. A roadmap assumes a given future and provides a framework toward realizing it. Often, a roadmap is part of the business and/or the product strategy towards growth and evolution.

Figure 1 is a product roadmapping of Company x, a mobile service provider. Figure 1 shows how the mobile services are said to evolve as we transit from 2G to 3G networking. As the bandwidth is improved, an emerging number of content-based services, ranging from voice, multi-media, data, and location-based services might be possible. This, in turn, will translate into future requirements (functional and non-functional), which need to be planned in advance so it can be accommodated by the architecture responsible for delivering the services. Note that many of the likely changes in the requirements may be derived from the roadmapping process, rather than the roadmap itself. As an example, M-banking is a service, which allows customers to check bank balances, view statements, and carry bank transactions using mobile phones. A distributed architecture of a banking system, which envisions providing such a service as the bandwidth is improved, may need to anticipate changes due to mobility like changes in security requirements, load, availability, etc. The architect may then need to anticipate relevant change scenarios and ways of accommodating them on the architecture.

E.g., M-banking availability:

(Requirements) Loss of connectivity is the norm in mobility. The M-banking service shall be available 99% of the time,
 (Architecture) New caching mechanisms are then required.



Source: <http://www.3g-generation.com/>

Figure 1. Company’s x technology road mapping showing the evolution of its mobile services as it moves from 2G to 3G and its value to the end user

4.3. Architectural Stability and Middleware

Recent research effort (e.g., [Jaz95; Gal97; Sul97; Ore98; Din99; Met00; Den04]) on the relation between software architectures and middleware has been motivated by pragmatic needs. The effort has revolved on issues such as investigating the compliancy of architectural styles with middleware; capabilities that the middleware and the architecture can bring when “coupled” to understand quality attributes of the system such as performance; mapping between middleware and software architectures; and semantics and syntactical issues related to the mapping process. As it has been noted in several occasions [Emm00b; Emm02], research on software architectures has over-emphasized functionality and not sufficiently addressed how global properties and non-functional requirements are achieved in an architecture, where these requirements cannot be attributed to individual components or connectors. Though ongoing research on the “coupling” of middleware and architectures could have an impact on understanding the relation between architectures and non-functional requirements, their contributions to such understanding is still insufficient. As far as the architectural stability problem is concerned, no effort has been devoted for understanding the evolution of non-functional requirements in relation to both the architecture and the middleware, when coupled. In [bah05b], we have claimed that as the non-functional requirements evolve, the “coupling” between the middleware and architecture becomes the focal point for understanding the *stability* of the distributed software system architecture in the face of change. We have hypothesised that the choice of a stable distributed software architecture depends on the choice of the underlying middleware and its *flexibility* in responding to future changes in non-functional requirements. Drawing on a case study that adequately represents a medium-size component-based distributed architecture: we have reported on how a likely future change in scalability could impact the architectural structure of two

versions, each induced with a distinct middleware: one with CORBA and the other with J2EE. We have devised an option-based model to value the flexibility of the induced-architectures and to guide the selection. Both versions have exhibited similar styles (i.e., three-tier), yet they have differed in the way they cope with the change in scalability. The difference was not due to the architectural style, but due to the primitives that are built in the middleware to facilitate scaling the software system. The governing factor, hence, appears to be to a large extent dependent on the flexibility of the middleware (e.g., through its built-in primitives) in supporting the change. Our preliminary observations suggest that the style by itself is not revealing to the analysis of architectural stability with respect to changes in non-functional requirements. Though this observation reveals a trend that agrees with the state-of-practice, confirming the validity of these observations are still subject to some systematic empirical studies. These studies may need to consider other non-functional requirements, their concurrent evolution, and their corresponding change impact on different architectural styles and middleware, which worth future research.

4.4. Traceability of Requirements to Architectures

An important outcome of the initial development of the software system is the knowledge that the development team acquires: the knowledge of the application domain, user requirements, role of the application in the business process, solutions and algorithms, data formats, strength and weakness of the architecture, and operating environment [Ben00]. This knowledge is acknowledged to be crucial prerequisite for evolution. In particular, both the architectures and the team knowledge make the evolution possible [Ben00]. These to a great extent allow the team to make changes in the software without damaging the architectural integrity. Once one or the other aspect disappears, the system is no longer evolvable and enters the stage of servicing (also referred to as maturity by Lehman) [Ben00]. At the servicing stage, only small tactical changes would be possible. For the business, the software is likely to be no longer a core product and the cost-benefit of the change becomes marginal. According [Ben00], there is a positive feedback between the loss of software architecture coherence and the loss of software knowledge. Less coherent architectures requires more extensive knowledge in order to evolve the system of the given architecture. However, if the knowledge necessary for evolution is lost, the changes in the software will lead to faster deterioration of the architecture. Very often on software projects, the loss of knowledge is triggered by loss in key personnel and the project slips into the servicing stage. Hence, planning for evolution and stable software architectures urges the need for traceability techniques, which traces requirements and their evolution back and forth into the architecture and aid in “preserving” the team knowledge.

[Dav93] defines traceability as “the ability to describe and follow (track) the lifetime of an artifact, in both a forward and a backward direction, i.e., from its origin to development and vice versa” [Dav93]. [Got95] have preserved the spirit of Davis’s definition of traceability. They, however, have scoped the definition on tracing a requirement through its “life”. The

requirements life covers periods of a requirement origin, development and specification, deployment, use, and on-going refinement. They have defined *requirements traceability* as “the ability to describe and follow the life of a requirement in both a forwards and backwards direction (i.e., from its origins, through its development and specification, to its subsequent deployment and use, and through periods of on-going refinement and iteration in any of these phases)”. [Got95] have particularly discussed the importance of tracing requirements back to their source. These sources might be people, other requirements, documents, or standards.

Traceability is important for modeling dependencies among software objects and for managing the change across software artifacts. Traceability information records the dependencies between requirements and the sources of these requirements, dependencies between requirements themselves, and dependencies between requirements and the system implementation. Advances in software-development environments and repository technology have enabled software engineers to trace the change in software using traceability techniques. These techniques span a variety of approaches ranging from cross-referencing schemes (e.g., cross-referencing schemes, based on some form of tagging, numbering, indexing, traceability matrices, and matrix sequences), through document-centered techniques (e.g., Templates, hypertext, and integration documents), to more elaborate structure-centered techniques (e.g., assumption-based truth maintenance networks, constraint networks, axiomatic, key phrase, and/or relational dependencies).

We define *requirement to architecture traceability* as the ability to describe the “life” of a requirement through the requirements engineering phase to the architecture phase in both forwards and backwards. Forwards demonstrates which (and how) architectural element(s) satisfy an individual requirement in the requirements specification. Backwards demonstrates which requirement(s) in the requirements specification an individual architectural element relate to and satisfy. Current architectural practices, however, do not provide support for traceability from the requirements specification to the architectural description. Maintaining traceability “links” is necessary for managing the change, the co-evolution of both the requirements and the architecture, confining the change, understanding the change impact on both the structure and the other requirements, providing a support for automated reasoning about a change at a high level of abstraction. Further, such traceability “links” make it easier to preserve the acquired knowledge of the team through guided documentation. This may then minimize the impact of personnel losses, and may allow the enterprise to make changes in the software system without damaging the architectural integrity and making the software system unevolvable.

4.5. Architectural Change Impact Analysis

Although change impact analysis techniques are widely used at lower levels of abstractions (e.g., code level) and on a relatively abstract levels (e.g., classes in O.O. paradigm), little effort has been done on the architectural level (i.e., architectural impact analysis). Formal notations for representing and analyzing architectural designs such as ADLs have provided

new opportunities for architectural analyses [Gar00]. Examples of such analyses includes system consistency checking [All94; Luc95], and conformance to constraints imposed by an architectural style [Abo93].

Notable effort using dependency analysis on the architectural level includes the “chaining” technique suggested by [Sta97]. The technique is analogous in concept and application to program slicing. In chaining, dependence relationships that exist in an architectural specification are referred to as links. Links connect elements of the specification that are directly related. The links produce a chain of dependencies that can be followed during analysis. The technique focuses the analysis on components and their interconnections. A component may have a set of input and output ports (which correspond to the component’s interface). These ports may have been connected to one another to form a particular architectural configuration. Communication between components is accomplished by sending events to the component’s ports. [Staf97] supports the approach with an analysis tool, Aladdin. Aladdin accepts an architectural specification as input. A variety of computations can be then performed. The computations include unconnected component identification, change impact analysis (i.e., which components will be affected by an architectural change), and event dependence analysis (i.e., which components can send the following event to this port). These computations start at a particular component and/or port. Forward and/or backward chaining are then performed to discover related components. Forward and backward chaining is analogous in concept to forward and backward walk in the data-flow slicing. The applicability of this technique is demonstrated on small scale architectures and could be extended to address current architectural development paradigms. For example, how such a concept could be refined to perform what-if analysis on large-scale software architectures such as product-line or model-driven architectures? For product-line architectures, this is necessary for reasoning about how the change could impact the commonality, variability, and their interdependence. These techniques could be then complemented by analysis tools which could facilitate automated reasoning and provide a basis for what-if analyses to manage the change across instances of the core architecture. Understanding how the change could then ripple across different products might be feasible. For model-driven architectures, for example, this could help in reasoning about how the change could affect the Platform Independent Model (PIM) and ripple to affect the Platform Specific Models (PSM). These techniques could be complemented by automated reasoning to manage evolution. When combined with traceability links, the combination could provide a comprehensive framework for managing the change and guiding evolution.

4.6. Empirical Studies

A key benefit of adopting an architecture-centric approach to manage evolution is driven by the objective of reducing future evolution costs, while attaining a net benefit and embedding options [Bah05a]. Though this is the motivation behind many architectural-centric evolution approaches, such as product-line architectures and model-driven architectures, little -if no- documented empirical evidence is

available on the extent to which the architecture has succeeded or failed in attaining its objectives. In particular, the architectural stability problem is just a hint on the fact that the architecture is also creating its own problems. This brings a need for systematic empirical studies to analyze real life cases, which lead to substantial “break” in the architecture of the software system upon accommodating changes. The “breakage” could be attributed to the nature of the change, personnel, the architectural style, the adopted middleware, and so forth. Lessons to be learned from these studies may have positive implications on the way we engineer our future requirements, design architectures to meet these changing requirements, and have better understanding on how we can control risks associated with the change and its impact. The main objective is to learn from the state-of-practice to improve the state-of-the-art.

5. Conclusion

Reflecting on our research into the problem, we have defined architectural stability and explored perspectives in handling the problem. We have reviewed existing research effort, have discussed their limitations, and have outlined research challenges and opportunities. The implications of such contribution need not be overstated: advancing the understanding of the architectural stability, stimulating and possibly motivating future research in architectural stability and related problems.

6. References

- [Abo93] Abowd, G., Allen, R., and Garlan, D.: Using Style to Understand Descriptions of Software Architecture. In: Proc. of Foundations of Software Engineering, ACM Press (1993) 9-20
- [All94] Allen, R., and Garlan, D.: Formalizing Architectural Connection. In: Proc. of the 14th Int. Conf. on Software Engineering, ACM Press(1994) 71-80
- [Ant96] Antón, A.I.: Goal-based Requirements Analysis. In: Proc. 2nd IEEE Int. Conf. on Requirements Engineering, IEEE CS(1996) 136-144
- [Bah03a] Bahsoon, R. and Emmerich, W.: Evaluating Software Architectures: Development, Stability, and Evolution. In: Proc. of IEEE/ACS Computer Systems and Applications, IEEE CS Press (2003a) 47-57
- [Bah03b] Bahsoon, R. and Emmerich, W.: ArchOptions: A Real Options-Based Model for Predicting the Stability of Software Architecture. In: Proc. of the 5th Workshop on Economics-Driven Software Engineering Research, with the 25th Int. Conf. on Software Engineering, IEEE CS (2003b) 35-40
- [Bah03c] Bahsoon, R.: Evaluating Software Architectures for Stability: A Real Options Approach. In: Proc. of the Doctoral Symposium of the 25th Int. Conference on Software Engineering, IEEE CS Press (2003)
- [Bah04a] Bahsoon, R. and Emmerich, W.: Evaluating Architectural Stability with Real Options Theory. In: Proc. of the 20th IEEE Int. Conf. on Software Maintenance, IEEE CS Press (2004a) 443-447
- [Bah04b] Bahsoon, R. and Emmerich, W.: Applying ArchOptions to Value the Payoff of Refactoring. In: Proc. of the Sixth Workshop on Economics-Driven Software Engineering Research, with the 26th Int. Conf. on Software Engineering, IEE Press (2004b) 66-70
- [Bah05a] Bahsoon, R. : Evaluating Architectural Stability with Real Options Theory, PhD thesis, U. of London, UK (2005)
- [Bah05b] Bahsoon, R., Emmerich, W., and Macke, J.: Using ArchOptions to Select Stable Middleware-Induced Architectures. In: IEE Proceedings Software, Special issue on Relating Requirements to Architectures, IEE Press 152(4) (2005) 176-186
- [Bah06] Bahsoon, R. and Emmerich, W.: Requirements for Evaluating Architectural Stability. In: Proc. of Int. Software Stability workshop, in Conj with IEEE/ACS Computer Systems and Applications, IEEE CS Press (2006)
- [Ben00] Bennet, K. and Rajilich, V.: Software Maintenance and Evolution: A Roadmap. In: A. Finkelstein (ed.): The Future of Software Engineering. ACM Press (2000) 73-90
- [Boe00] Boehm, B., and Sullivan, K. J.: Software Economics: A Roadmap. In: A. Finkelstein (ed.): The Future of Software Engineering. ACM Press (2000) 320-343
- [Bur90] Burch, J., Clarke, E., McMillan, E., Dill, D., and Hwang, L.: Symbolic Model Checking: 10²⁰ States and Beyond. In: Proc. of the Fifth Annual IEEE Symposium on Logic in Computer Science. IEEE CS (1990) 428-439
- [Cap03] Capra, L.: Reflective Mobile Middleware for Context-Aware Applications. PhD Thesis. University of London, UK (2003)
- [Cle00] Clements, P.: Active Reviews for Intermediate Designs. Technical Report (CMU/SEI-2000-TN-009), Software Engineering Institute, (2000)
- [Coo01] Cook, S., Ji, H., and Harrison, R.: Dynamic and Static Views of Software Evolution. In: Int. Conf. on Software Maintenance, Florence, Italy. IEEE CS (2001) 592-601
- [Corb97] Corbett, J., and Avrunin, G.: Using Integer Programming to Verify General Safety and Liveness Properties. Formal Methods in System design, 6(2)(1997) 97-123.
- [Dar93] Dardenne, A., van Lamsweerde A., and Fickas, S.: Goal-Directed Requirements Acquisition, Science of Computer Programming, 20(1-2) (1993) 3-50
- [Dav93] Davis, A: Software Requirements: Objects, Functions and States. Englewood Cliffs, New Jersey: Prentice-Hall (1993)
- [Den04] Denaro, G., Polini A., and Emmerich W.: Performance Testing of Distributed Component Architectures. In: S. Beydeda and V. Gruhn (eds.), Building Quality into COTS Components - Testing and Debugging, Springer (2004) 294-314
- [DiN99] Di Nitto, E., and Rosenblum, D.: Exploiting ADLs to Specify Architectural Styles Induced by Middleware Infrastructures. In: Proceedings of the 21st International Conference on Software Engineering, ACM Press (1999) 13-22
- [Dwy94] Dwyer, M. and Clarke, L.: Dataflow Analysis for Verifying Properties of Concurrent Programs. In: Proc. of 2nd ACM SIGSOFT Symp. on Foundations of Software Engineering, ACM Press (1994) 62-75
- [EDS99-05] EDSER 1-7: Proceedings of the Workshops on Economics-Driven Software Engineering Research: In conj. with the 21st through 27th International Conference on Software Engineering (1999 - 2005)
- [Emm00a] Emmerich, W.: Engineering Distributed Objects. John Wiley & Sons, Chichester, UK (2000a)
- [Emm00b] Emmerich, W.: Software Engineering and Middleware: A Road Map. In: A. Finkelstein (ed.), Future of Software Engineering, ACM Press (2000b) 117-129
- [Emm02] Emmerich, W.: Distributed Component Technologies and their Software Engineering Implications. In: Proc. of the 24th Int. Conf. on Software Engineering, Orlando, Florida, ACM Press (2002) 537-546
- [FEA] Lehman, M.M.: Feedback, Evolution and Software Technology, FEAST 1-2. <http://www-dse.doc.ic.ac.uk/~mml/feast/>
- [Fin00a] Finkelstein, A., and Kramer, J.: Future of Software Engineering. In: A. Finkelstein (ed.): The Future of Software Engineering, ACM Press (2000) 5-21
- [Fin00b] Finkelstein, A.: Architectural Stability. <http://www.cs.ucl.ac.uk/staff/a.finkelstein/talks.html> (2000)
- [FS92] Formal Systems (Europe) Ltd.: Failures Divergence Refinement: User Manual and Tutorial (1992)
- [Gal97] Gall, H., Jazayeri, M., Klösch, R., and Trausmuth, G.: The Architectural Style of Component Programming. COMPSAC, IEEE CS Press (1997) 18-27
- [Gar00] Garlan, D.: Software Architecture: A Roadmap. In: A. Finkelstein (ed.): The Future of Software Engineering, ACM Press (2000) 91-101

- [Gar94] Garlan, D., Allen, R., and Ockerbloom: Exploiting Style in Architectural Design Environments. In: Proceedings of SIGSOFT'94, Foundations of Software Engineering, ACM Press(1994)175-188
- [Gar95] Garlan, D., Monroe, R. and Wile, D.: ACME: An Architectural Interconnection Language. Technical Report (CMU-CS-95-219) (1995)
- [God91] Godefroid, P., and Wolper, P.: Using Partial Orders for the Efficient Verification of Deadlock Freedom and Safety Properties. In: Proceedings of the Third Workshop on Computer Aided Verification, Lecture Notes in Computer Sc, Springer (1991)417-428
- [Got95] Gotel, O., and Finkelstein, A.: Contribution Structures. In: 3rd Proc. of the Requirements Engineering Symposium, York, UK, IEEE CS Press (1995) 169-178
- [Hol91] Holzman, G.: Design and Validation of Computer Protocol, Prentice Hall Software Series (1991)
- [Jaz02] Jazayeri, M.: On Architectural Stability and Evolution. Lecture Notes in Computer Science, Springer Verlag, (2002) 13-23
- [Jaz95] Jazayeri, M.: Component Programming - a Fresh Look at Software Components, In: 5th European Software Engineering Conference. Lecture Notes in Computer Sc, Springer (1995) 457-478
- [Kaz01] Kazman, R., Asundi, J., and Klein, M.: Quantifying the Costs and Benefits of Architectural Decisions. In: Proc. of 23rd Int. Conf. on Software Engineering, IEEE CS Press (2001) 297-306
- [Kaz94] Kazman, R., Abowd, G., Bass, and L., Webb, M.: SAAM: A Method for Analyzing the Properties of Software Architectures. In: Proc. of 16th Int. Conf. on Software Engineering, IEEE CS (1994) 81-90
- [Kaz98] Kazman, R., Klein, M., Barbacci, M., Lipson, H., Longstaff, T., and Carrière, S.J.: The Architecture Tradeoff Analysis Method. In: Proc. of 4th Int. Conf. on Engineering of Complex Computer Systems IEEE CS Press (1998) 68-78
- [Kle99] Klein, M., and Kazman, R.: Attribute-Based Architectural Styles. CMU/SEI-99-TR-22, Software Engineering Institute, (1999)
- [Leh00] Lehman, M.M., Kahen, G., and Ramil, J.F.: Replacement Decisions for Evolving Software. In: Proc. of the Second Workshop on Economics-Driven Software Engineering Research (2000)
- [Leh98] Lehman, M.M.: The Future of Software – Managing Evolution. IEEE Software (Jan. 1998)
- [Luc95] Luckham, D.C., and Vera, J.: An Event-Based Architecture Definition Language. IEEE Trans. on Software Engineering, 29(9) (1995) 717-734
- [Mag95] Magee, J., Dulay, D., Eisenbach, N., and Kramer, J.: Specifying Distributed Software Architecture. In: Proc. of Fifth European Software Engineering Conf., Lecture Notes in Computer Sc, Springer (1995) 137-153
- [Mas91] Masticola, S., and Ryder, B.: A Model of ADA Programs for Static Deadlock Detection in Polynomial Time. In: Proc. of the Workshop on Parallel and Distributed Debugging (1991) 97-107
- [Med03] Medvidovic N., Dashofy E., and Taylor R.: On the Role of Middleware in Architecture-based Software Development. International Journal of Software Engineering and Knowledge Engineering, 13(4) (2003) 229-306
- [Med97] Medvidovic, N., and Taylor, R.: A Framework for Classifying and Comparing Architecture Description Languages. In: Proc. of 6th. European Software Engineering Conf., with the Fifth ACM SIGSOFT Symp. on the Foundations of Software Engineering, ACM Press (1997)60-76
- [Meh00] Mehta, N., Medvidovic, N., and Phadke, S.: Towards a Taxonomy of Software Connectors. In: Proceedings of the 22nd International Conference on Software Engineering, ACM Press (2000) 178-187
- [Mor95] Moriconi, M., Qian, X., and Riemenschneider, R.: Correct Architecture Refinement. IEEE Trans. on Software Engineering, 21(4) (1995) 356-372
- [Nau97] Naumovich, G., Avrunin, G.S, Clarke, L.A., and Osterweil, L.J.: Applying Static Analysis to Software Architectures. UM-CS-1997-008, University of Massachusetts, Amherst (1997)
- [Nus00] Nuseibeh, B., and Easterbrook, S.: Requirements Engineering: A Roadmap. In: A. Finkelstein (ed.): The Future of Software Engineering, ACM Press (2000) 35-46
- [Nus01] Nuseibeh, B.: Weaving the Software Development Process between Requirements and Architectures. In: Proc. of the First Int. workshop from Software Requirements to Architectures, Toronto, Canada (2001)
- [Orei98] Oreizy, P., Medvidovic, N., Taylor, R., and D. Rosenblum, D.: Software Architecture and Component Technologies: Bridging the Gap. In Digest of the OMG-DARPA-MCC Workshop on Compositional Software Architectures, Monterey, CA (1998)
- [Par72] Parnas, D.L.: On the Criteria to Be Used in Decomposing Systems into Modules, Communications of the Association of Computing Machinery, 15(12) (1972)1053-58
- [Par79] Parnas, D.L.: Designing Software for Ease of Extension and Contraction, IEEE Transaction on Software Engineering, 5 (2) (1979)
- [Sch99] Schaller, R.R.: Technology Roadmaps: Implications for Innovation, Strategy, and Policy, The institute of Public Policy, George Mason University Fairfax, VA (1999)
- [Sch99] Schaller, R.R.: Technology Roadmaps: Implications for Innovation, Strategy, and Policy, The institute of Public Policy, George Mason University Fairfax, VA (1999)
- [Sha95] Shaw, M., DeLine, R., Klein, D., Ross, T., and Young, D.: Abstractions for Software Architecture and Tools to Support them. IEEE Transactions on Software Engineering, 21(4) (1995) 314-335
- [Smi99] Smith, C., and Woodside, M.: System Performance Evaluation: Methodologies and Applications. CRC Press (1999)
- [Staf01] Stafford, J. A., and Wolf, A. W.: Architecture-Level Dependence Analysis for Software System. International Journal of Software Engineering and Knowledge Engineering, 11(4) (2001) 431-453
- [Sul97] Sullivan, K. J., Socha, J., and Marchukov, M.: Using Formal Methods to Reason about Architectural Standards. In: Proceedings of the 19th Int. Conf. on Software Engineering, MA, ACM Press (1997) 503-513
- [Sul99] Sullivan, K. J.: Chalasani, P., Jha, S., and Sazawal, V.: Software Design as an Investment Activity: A Real Options Perspective. Real Options and Business Strategy: Applications to Decision-Making. In: Trigeorgis L. (ed.) Risk Books (1999) 215-260
- [Val91] Valmari, A.: A Stubborn Attack on State Explosion. In: E. M. Clarke and R. Kurshan(editors), Computer-Aided Verification 90. American Mathematical Society, Providence RI. Number 3 in DIMACS Series in Discrete Mathematics and Theoretical Computer Science, (1991) 25-41
- [van00] van Lamsweerde, A.: Requirements Engineering in the Year 00: A Research perspective. In: Proc. 22nd Int. Conf. on Software Engineering, (2000) ACM Press 5-19
- [Ves96] Vestal, S.: MetaH Programmer's Manual, Version 1.09, Technical Report, Honeywell Technology Center (1996)