

Published in the Proceedings of 3rd International Workshop on Development and Evolution of Software Architectures for Product Families (IWSAPF), Las Palmas de Gran Canaria, Spain, March 2000

Software Connectors and Refinement in Family Architectures

Alexander Egyed Nikunj Mehta Nenad Medvidović

{aegyed, mehta, neno}@sunset.usc.edu

Department of Computer Science
University of Southern California
Los Angeles, CA 90089-0781, USA

Abstract. Product families promote reuse of software artifacts such as architectures, designs and implementations. Product family architectures are difficult to create due to the need to support variations. Traditional approaches emphasize the identification and description of generic components, which makes it difficult to support variations among products. This paper presents an approach to modeling family architectures using generic software connectors that provide bounded ambiguity and support flexible product families. The paper also proposes an approach for transforming a family architecture to a product design through a four-way refinement and evolution process.

1 Introduction

Large, complex systems are often developed in the context of *product families*¹. This enables developers to maximize reuse, accelerate the development process while reducing costs, and deliver products that are generally more reliable. Reuse across product families occurs in terms of architecture, design and implementation. Architectural idioms identify the kinds of building blocks that may be used to compose a system and specify the constraints on the way the composition is done. An explicit focus on common *architectural* idioms has the potential to fundamentally transform the nature of software development, as component *integration* replaces implementation as the predominant development activity. The promise of software architectures is that better software systems can be built in this manner more quickly by modeling their important aspects throughout, and especially early in the development. Coupling the benefits of product family-based and architecture-based development has been recognized as an area with a great potential payoff, as evidenced by a growing number of conferences, workshops, and symposia that focus on this subject [2, 3, 4, 8, 12].

¹ In this paper, we use the following phrases interchangeably: families, application families, product families, product lines, and domain-specific software.

The existing body of research in the area of software architectures for product families is characterized by two major foci:

1. specification of generic, product family architectures (also referred to as *reference* architectures) and their instantiation into application architectures (e.g., [11]); and
2. identification and integration of reusable components that comprise different members of a product family (e.g., [5]).

In this paper, we focus on two additional issues that have not been addressed by existing approaches and that are useful complements to those identified above:

1. the role of software connectors in specifying and ensuring the extra-functional properties of both a product family and individual applications within the family; and
2. refinement of an instantiated product architecture into a design and, eventually, an implementation.

The role of *connectors* in software architectures is to isolate all communication, coordination, and mediation [10]. Connectors do not generally provide domain-specific functionality, but rather enable and streamline interactions among the functional elements (components). Thus, our hypothesis is that certain varying properties of applications within a family (e.g., deployment profile, concurrency, interoperability platform, performance, reliability, security, etc.) can be isolated within connectors. Also, certain types of connectors may occur regularly within a family. Our on-going work on classifying software connectors will serve as a vehicle for exploring these issues.

To enable the *refinement* of an architecture into its implementation, we leverage our work on transforming architecture-level constructs (specified in an architecture description language, or ADL [9]), into design-level constructs (specified in the Unified Modeling Language, or UML [6]), and enabling the refinement of the resulting high-level design in a property-preserving manner [1, 7]. We introduce the notion of *product family design*, analogous at the design level to a product family architecture. A product family design captures recurring design patterns across components in a family. Another hypothesis is that both product family designs and product architectures are needed to enable effective refinement.

The paper is organized as follows. Section 2 identifies the relationships between products and families, and between architectures and designs. Section 3 outlines the role of software connectors in family architectures. Instantiation of a product family architecture and refinement of the resulting product architecture into its implementation is discussed in more detail in Section 4. Section 5 presents an example illustrating the approach. Conclusions and a discussion of open issues round out the paper.

2 Relationship of Products and Families

Software architectures can be described using components, connectors and configurations [9]. Components are units of data store or computation whereas, connectors model the interactions among components. Architectural description identifies the obligations and freedoms of a software system built to that architecture. Obligations allow a high level analysis of system properties while the freedoms allow developers to design and implement the system according to the characteristics and constraints of an underlying infrastructure. Since a product family consists of products with commonalities and differences, it is useful to capture these aspects of the individual products in *family architecture*. Moreover, use of the same architectural elements to describe family architectures and individual product architectures aids in keeping these artifacts consistent and simplifies understanding.

A family architecture provides generic information common to all the products of the family. This common information may include features present in all systems or a list of possible alternatives that products exhibit. It is easy to represent family architectures in terms of the similarities alone. However, in order to support variations in the individual products, a family architecture needs to describe the architectural elements with a certain amount of ambiguity. The product architecture, on the other hand, identifies specific architectural choices for a single product and thus can be considered as an instantiation of the family architecture. The product architecture is less ambiguous since all architectural elements are already chosen and specified for the sake of completeness. Proceeding in another direction, the family architecture can also be refined to create a set of more detailed family designs that can be used in individual products to obtain the recurring functional and extra-functional properties.

Figure 1 depicts a high-level framework we propose for architectural modeling of product families and their improved refinement and evolution. The white boxes and arrows in the figure denote the traditional way of instantiating product architectures from family architectures, followed by the refinement of product architectures into their designs (and subsequent implementations). To complement this traditional approach, we introduce the concept of *family design*. A family design contains




Figure 1. Design refinement and instantiation using product architecture and family design

design-related information about a product family that the architecture did not (or could not) specify. For instance, a family design could contain different design interpretations of architectural elements (e.g., in the form of design patterns). Merging the product architecture information with family design information can then lead to a product design. Therefore, the product architecture defines at a high level *what* needs to be designed and the family design provides information on *how* to design it. This four-way relationship between architecture, design, family, and product implies that there are at least two alternate but complementary paths of creating designs for the products of a given family.

3 Software Connectors in Family Architecture

Family architectures capture the essential properties and relations of the product family. They describe structural and behavioral freedom and model the functional and extra-functional aspects of a family. Behavioral freedom and functional aspects of the product family are typically captured in components. Our approach also supports the description and analysis of extra-functional properties, coupled with the identification of the structural freedoms through the use of semantically rich connectors.

As discussed above, family architectures need to describe commonalities as well as variations among family members. Commonalities can be captured through elements that are mandatory to all products. The real difficulty lies in modeling variations that have to be supported at the level of a product family. Various approaches have been proposed to describe family architectures including the use of styles, parameters, constraints and service provisioning. However, as discussed in [11], none of these techniques alone adequately addresses the problem of supporting variations in the product family. There is clearly a need for defining family architectures with a certain degree of *bounded ambiguity* in order to support product variations.

Consider the case of a customer service product family that needs to support two product domains, retail banking and telephony. These products require variations in terms of the underlying information, as well as in the interaction of the architectural components. The banking application requires online transactions, whereas the telephony product requires a batch update. A useful family architecture would be able to support the description of both kinds of products.

Software architecture captures the essential structural and behavioral information in the form of components and connectors. Family architectures are useful because they lead to better structured reuse and also because the bulk of the architectural analysis can be performed at the level of an entire family. There is a tradeoff between vagueness of description and the scope of applicability when it comes to specifying the architecture in product families. At the level of a product family, components tend to be vaguely described because family architectures need to support a variety of product features. This vagueness about components reduces our ability to reason about the family architecture. On the other hand, in the product architecture, components are described more precisely [13]. This tradeoff gives rise to reduced

analyzability at the family architecture level and reduced flexibility of the concrete products.

Current techniques for representing components in family architectures, however, tend to be inflexible. Many reuse techniques depend on the availability of interchangeable components that can lead to a component marketplace. However, experience shows that such components can only be achieved through considerable standardization efforts. Standardization tends to be a long process in which decisions are often made at a corporate level rather than industry-wide level. Component centered reuse therefore tends to take longer to adapt and is applicable only to niche domains. On the other hand, the software industry has very quickly embraced component integration frameworks such as DCE RPC [14], COM [15], CORBA[16] and Enterprise Java Beans [17]. This indicates that the industry is more amenable to accepting standards of integrating components than to standards of defining components. We therefore focus our research on the role of software connectors in family architectures.

Software connectors describe the interactions among architectural components and support communication, coordination, conversion and facilitation needs of components [18]. Connectors can be used to describe interactions among components in family architectures. Furthermore, many extra functional properties of a system can be attributed to semantically rich connector mechanisms such as events, distributors and arbitrators. Since connectors can be applied across problem domains, they have a high potential for reusability. Connectors also significantly affect global system properties such as availability, throughput, security and scalability. Various architectural-styles motivated by software connectors have been studied, e.g. pipe and filter [19], real-time data feeds [20], event-driven architecture [21], message-based style [22], middleware-induced styles [23], and push-based systems [11]. Architectural styles are an important mechanism for enabling reuse in family architectures [11, 24], indicating that software connectors have a major role to play in enabling architecture-based reuse.

Connectors provide bounded ambiguity that is necessary for supporting variations in family architectures. In order to effectively exploit that ambiguity we have used a taxonomy view of software connectors that describes the connector types, dimensions and their possible values (see Figure 2 for an extract). The ambiguity is contained in the various dimensions along which a connector can be characterized and the range of values that a connector can assume for each dimension. Since there are a finite number of values that can be assigned for each connector dimension, ambiguity involved in defining connectors in an architecture is bounded. Family architectures can be vague about the component interactions and as such can be described using imprecise connectors, i.e. connectors that identify a range of values for connector dimensions. Our taxonomy allows architects to choose the concrete connectors necessary to support interaction among components and to provide the dimensions along which each specific product can choose a different variation of interaction. Many extra-functional properties can thus be evaluated based on connectors as their dimensions of variation are known.

As a solution for the example problem introduced above, the customer service product family architecture would describe the required component interactions in the form of an event connector that allows variations along dimensions shown in Figure

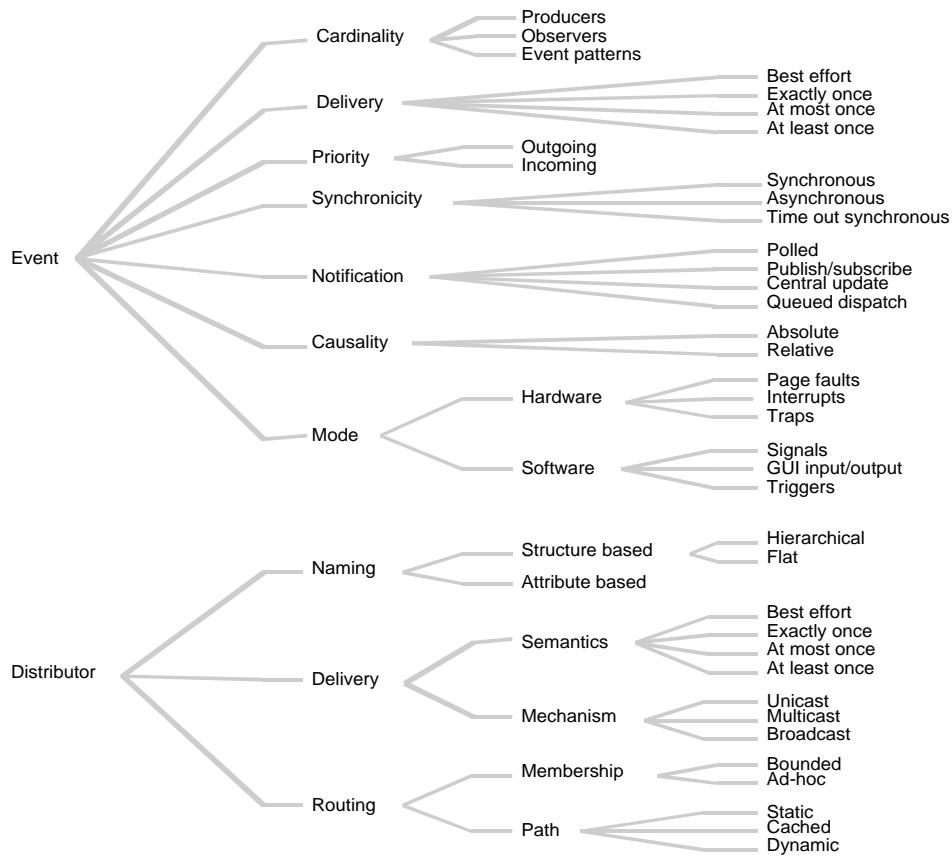


Figure 2. An excerpt from the connector taxonomy showing (from left to right) types, dimensions, subdimensions, and values.

2. It then becomes possible to describe both forms of required interactions - online transactions and batch updates - based on the event dimensions of *notification* and *synchronicity*. It is possible to describe the distribution profile of the interactions using the distributor dimensions *delivery* and *addressing*.

We are currently developing an infrastructure for using and experimenting with connectors for implicit invocation, real time communication and parallel execution. This infrastructure builds upon our previous work with event connectors and adaptors [22].

4 Four-Way Refinement and Evolution

Having discussed an approach to exploiting connectors in modeling family architectures and instantiating them into product architectures, we now discuss how to

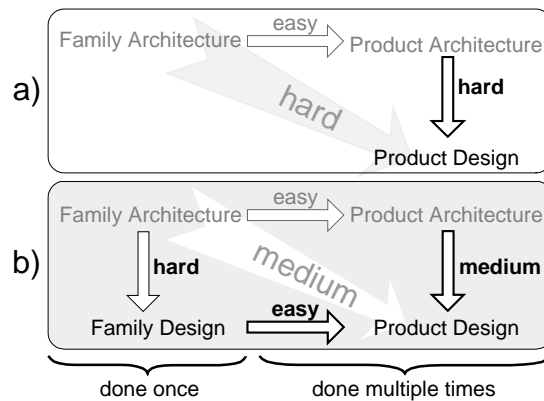


Figure 3. Two refinement approaches

refine the resulting product architectures into individual product designs. A product architecture constitutes an effective milestone [25] for any project since it can be analyzed and simulated to ensure the presence (or absence) of properties of interest. Nevertheless, it is still a difficult task to refine those architectural models into designs and actual implementations.

Refinement involves the creation of lower-level design models (and ultimately source code) and their continuous validation to ensure consistency. Refinement is difficult and it often has to be done manually. This, however, implies that defects may be introduced while refining the product design from its architecture. Thus, we are faced with a major problem: we create the family and product architectures with the understanding that they describe certain desirable properties the end system should exhibit; at the same time, if consistent refinement and evolution cannot be ensured, then there is no guarantee that the final product will indeed exhibit those properties. In other words, inconsistent refinement invalidates the purpose and utility of the family and the product architecture. Meaningful architectural modeling *must* therefore ensure faithful refinement and evolution. This section will discuss an approach to improve the integrity of product models through the automation of refinement.

Automation during Refinement

The traditional way of modeling family architectures involves instantiating a family architecture into a product architecture, followed by refining that product architecture into a product design (Figure 3a – white area). This process may seem simpler in comparison to our proposed approach (Figure 3b – gray area) of using a family design, mainly because our approach additionally requires (1) modeling of a family design and (2) knowledge of how to relate it to the product architecture. However, we believe that these two additional activities ultimately simplify the overall refinement process. Using the traditional “family architecture to product architecture to product design” approach requires an instantiation from family architecture to product architecture and a refinement from the product architecture to a corresponding product design. The instantiation is relatively easy to do compared to the refinement activity, which is complicated by the lack of automation support. Even

if the refinement could be automated, we would still be faced with the possibility of mismatch introduction at a later stage, e.g., when either the design or architecture is altered and those modifications are not properly propagated throughout the family.

Using a family design requires one additional instantiation activity – from a family design to a product design – which is again relatively straightforward. However, by replacing some (hard) refinement activities with (easy) instantiation activities, we achieve the added benefit of having to do less refinement, which, in turn, significantly simplifies doing product designs. Furthermore, we get the benefits of design reuse which complements architectural reuse (enabled through the use of family architecture). On the downside, arriving at a family design is not trivial or easily automated. The major advantage in using a family design is that we need to create it only *once* for each product family. Thus, once the family design is in place, each additional product can be architected and designed much more easily (because of having simplified the refinement of product architectures). On the other hand, without a family design we may avoid the hard initial task of creating and instantiating it, but the task of refining product architecture would be harder. It is not difficult to see the return on investment of doing a one-time difficult task that simplifies a later repetitive task as opposed to avoiding that hard initial task but complicating the repetitive one. Therefore, family design allows us to shift parts of the hard repetitive tasks from product architecture refinement to family architecture refinement.

An added benefit of standard family designs is that they can be used to realize different species of the same connector type. For example, design patterns for central dispatch, publish-subscribe and queued dispatch events can be described in the form of family designs that realize the event connector, and are eventually instantiated in the product designs, based on the specific mode of interaction required. Additional design patterns can describe the other dimensions. This technique requires integration of design patterns in the family design for specific values of the different connector dimensions into a single software product design.

Continuing our previous example, the use of events in the customer support product family architecture as the means of component interaction would leave a lot of flexibility in the design of individual products; the family architecture can support a large number of variations in each product. The product architecture can be used to instantiate an event connector by selecting the dimensions of each connector instance in the family architecture. This is an easy step, as it would involve looking up the taxonomy of connectors and making choices for the dimensions of a connector. Family designs can then provide standard refinements in the form of design patterns of event-based interaction for different platforms and middleware environments. Finally, the product design would select specific design patterns for the target environment and desired product properties of the system.

5 Example: Going from Family Architecture to Product Design

Figure 4 depicts a simple example on how to use generic connectors and family design concepts to generate a product design from a family architecture definition.

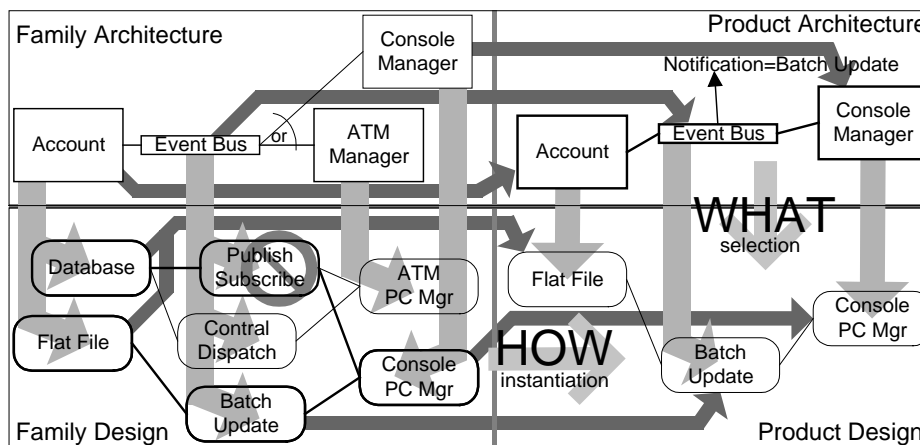


Figure 4. Example of instantiation, refinement and traceability links

The figure depicts a simple accounting family architecture (upper left) that supports access to accounting information via an event-based connector. This particular family architecture allows two types of interfaces, one for ATM machines and one for terminal consoles, but only one at a time. The family design (lower left) depicts possible realizations of above architectural elements. Note that there are realizations for both architectural components and connectors. Furthermore, we need to be able to deal with incomplete family design specifications including missing links (e.g., missing glue code) and missing realizations for some architectural elements. For instance, *Flat File* is a realization of *Account*; however, it does not work together with any realization of *Event Bus* (*publish-subscribe*, *control dispatch*, or *batch update*).

In our example, we decided to instantiate a product architecture that consists of *Account*, *Event Bus*, and *Console Manager* (upper right). Using this product architecture as a reference and the family design as a resource database, we can now design and build the product using a predefined set of realizations. For instance, we could use the *Console PC Mgr* and combine it with either a *Publish-Subscribe* bus and a *Database*, or a *Batch-Update* bus and a *Flat File*. When we specified the product architecture, we also specified some attributes the architectural connectors should demonstrate. For instance, we pre-selected the *Event* bus to be of the *Batch Update* style. With this additional information, we can now automatically select a possible product design from the family design that would be compatible with the product architecture (lower right). The product architecture supplies information on *what* to design, while the family design provides a details on *how* to design a product.

6 Conclusions

This paper identified and addressed two significant challenges in product family development: modeling family architectures via generic connectors and supporting automatic architectural refinement via family designs. Our approach involves the use

of a taxonomy of connectors to model the bounded ambiguity in family architectures. We do not claim that connectors are more important than components for enabling family architectural descriptions; however we have found that, in some respect, connectors are significantly more flexible and reusable than components.

To enable automatic refinement and evolution, we introduced the concept of family design. Family designs provide a set of realizations of architectural components and connectors (e.g., in the form of design patterns). They simplify refinement by providing an additional path from a family architecture to a product design. We believe that combining a product architecture and a family design provides simplified and more precise refinement.

To date we have developed a suite of tools that allows automated mapping between architecture and design as well as their consistency checking. We have also studied the role of complex connectors in simplifying component integration and generating designs [1] and implementations [26]. The techniques used in this paper extend our previous work in the area of product line architectures [24]. This work is still in progress and it will evolve in several directions, including refining of our taxonomy of connectors, providing automated support for creating family designs, and resolving mismatches among architectural and design views at the level of a product family.

Acknowledgements

This research is sponsored by the Defense Advanced Research Projects Agency, and Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement numbers F30602-94-C-0195 and F30602-99-C-0174, as well as by the Affiliates of the USC Center for Software Engineering. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency, Air Force Research Laboratory or the U.S. Government.

The authors are thankful to Sandeep Phadke and Marija Rakić who provided useful feedback while discussing the topics presented in this paper.

References

1. M. Abi-Antoun and N. Medvidovic. Enabling the Refinement of a Software Architecture into a Design. In Proceedings of The Second International Conference on The Unified Modeling Language (UML'99), Fort Collins, CO, October 1999.
2. ARES. Proceedings of the International Workshop on Development and Evolution of Software Architectures for Product Families, Las Navas del Marqués, Ávila, Spain, November 1996. <http://hpv17.infosys.tuwien.ac.at/Projects/ARES/public/AWS/>
3. ARES II. F. van der Linden, editor. Proceedings of the Second International Workshop on Development and Evolution of Software Architectures for Product Families, Las Palmas de Gran Canaria, Spain, February 1998.
4. ARES III. The Third International Workshop on Development and Evolution of Software Architectures for Product Families, Las Palmas de Gran Canaria, Spain, February 2000.

5. D. Batory, L. Coglianesi, S. Shafer, and W. Tracz. The ADAGE Avionics Reference Architecture. In Proceedings of AIAA Computing in Aerospace 10, San Antonio, 1995.
6. G. Booch, I. Jacobson, and J. Rumbaugh. The Unified Modeling Language User Guide. Addison-Wesley, 1998.
7. A. Egyed and N. Medvidovic. A Formal Approach to Heterogeneous Software Modeling. Alexander Egyed and Nenad Medvidovic, to appear in Proceedings of Foundational Aspects of Software Engineering, Berlin, Germany, 2000.
8. R. Hayes-Roth and W. Tracz. DSSA Tool Requirements for Key Process Functions. ADAGE Technical Report, ADAGE-IBM-93-13B, October 1994.
9. N. Medvidovic and R.N. Taylor. A Classification and Comparison Framework for Software Architecture Description Languages. Accepted for publication in IEEE Transactions on Software Engineering, 2000. (To appear)
10. D. E. Perry. Software Architecture and its Relevance to Software Engineering, Invited Talk. Second International Conference on Coordination Models and Languages (COORD '97), Berlin, Germany, September 1997.
11. D. E. Perry. Generic Descriptions for Product Line Architectures. In Proceedings of the Second International Workshop on Development and Evolution of Software Architectures for Product Families (ARES II), Las Palmas de Gran Canaria, Spain, February 1998.
12. The First Software Product Line Conference, August 28-31, 2000, Denver, Colorado, USA. <http://www.sei.cmu.edu/plp/conf/SPLC.html>
13. D. Batory and S. O'Malley. The Design and Implementation of Hierarchical Software Systems with Reusable Components. ACM Transactions on Software Engineering and Methodology, 1(4), October 1992, pp. 355-398.
14. The Open Group, <http://www.opengroup.org>
15. Microsoft Corp. <http://www.microsoft.com/com>
16. Object Management Group, <http://www.omg.org>
17. Sun Microsystems. <http://java.sun.com/j2ee>
18. N. Mehta, N. Medvidovic and S. Phadke, Towards a Taxonomy of Software Connectors, Technical Report, Center for Software Engineering, University of Southern California, USC-CSE-99-529, 1999.
19. M. Shaw and D. Garlan. Software Architecture: Perspectives on an Emerging Discipline, Prentice-Hall, Upper Saddle River, NJ, 1996.
20. N. Roodyn and W. Emmerich. An Architectural Style for Multiple Real-Time Data Feeds. 21st International Conference on Software Engineering (ICSE '99), Los Angeles, CA, May 1999.
21. A. Carzaniga, E. Di Nitto, D. S. Rosenbloom and A. L. Wolf. Issues in Supporting Event-based Architectural Styles. 3rd International Software Architecture Workshop (ISAW3), Orlando FL, 1998.
22. R. N. Taylor, N. Medvidovic, K. M. Anderson, E. J. Whitehead and J. E. Robbins. A component- and message-based architectural style for GUI software. IEEE Transactions on Software Engineering, 1996, 22(6), pp. 390-406.
23. E. Di Nitto and D. Rosenbloom. Exploiting ADLs to Specify Architectural Styles Induced by Middleware Infrastructures. 21st International Conference on Software Engineering (ICSE '99), Los Angeles, CA, May 1999.
24. N. Medvidovic and R. N. Taylor. Exploiting architectural style to develop a family of applications. In IEE Proceedings Software Engineering, Vol. 144 No 5-6, October 1997.
25. B. Boehm, Anchoring the Software Process, IEEE Software, July 1996.
26. N. Medvidovic, D. S. Rosenblum, and R. N. Taylor. A Language and Environment for Architecture-Based Software Development and Evolution. In Proceedings of the 21st International Conference on Software Engineering (ICSE'99), pp. 44-53, Los Angeles, CA, May 16-22, 1999.