# GENERIC ARCHITECTURE DESCRIPTIONS
# FOR PRODUCT LINES

## D E Perry

**Rapporteur:** Dr Paul Ezhilchelvan

# Generic Architecture Descriptions for Product Lines

Dewayne E. Perry

Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974 USA
dep@research.bell-labs.com, www.bell-labs/com/usr/dep/

## 1 Introduction

Two of the fundamental needs in defining an architecture for a product line are

- to be able to generalize or abstract from the individual products to capture the important aspects of the product line and
- to be able to instantiate an individual product architecture from the product line architecture.

In other words, having a product line implies having a generic architecture from which the individual product architectures can be derived in some manner.

There are a number of different ways in which one might go about defining the product line architecture so that this desired level of genericity can be achieved. Five possible ways of doing this are

- use a software architecture style,
- use an under–constrained architecture description,
- define a variance–free architecture,
- use parametric descriptions with varying binding times, and
- use a service oriented description for selective provisioning.

In the end, I think you will need all of these for a systematic and complete generic product line architecture. I will discuss each of these in turn and delineate their strengths and weaknesses.

## 2 A Style as a Generic Architecture

There is a certain intuitive appeal in using a product line specific architectural style as the generic architecture for a product line. It would capture the essential characteristics of the product line while ignoring the variations and leave them to be supplied as needed in the actual product architecture. These essential characteristics would encompass the necessary components that each instance must have, the basic minimum interactions that each instance must have and the basic constraints on these components and interactions.

The utility of a style description is that it represents the minimalist approach to software architecture in general and product line architecture in specific. Only the critical aspects of the product line need to be considered in the architectural

specification. One primary advantage is that new products can be added to the line with ease as long as they conform to the basic product line stylistic constraints. This provides a wide degree of latitude in the the various products and what they provide relative to the core essence of the product line.

One of the negative side effects of this approach is the amount of work needed to refine the product line style into a particular product architecture. With the intent of a style as capturing only the essential architectural aspects of the product line, those aspects must be extended and added to in order to create in individual product architecture. As such the product architecture must be analyzed for conformity to the product line architecture.

As a result of this lack of completeness other aspects of architectural based development suffer as well. For example, analysis of the product line architecture will, of necessity, be less comprehensive. Project planning will be similarly less comprehensive at the product line level and the majority of planning work will be delayed until after a complete product architecture has been extended from the core style.

Further care must be taken in evolving the product line's architectural style so as not to invalidate existing product architectures. With each change to the product line style, the individually derived product architectures will have to be re-analyzed to ensure that the product architectures remain conforming to the style.

On the whole there are better uses of styles for product line architectures than defining the generic product line architecture itself. For example, one could define a set of styles defining such things as initialization, fault recovery, etc that all the various components in the architecture must adhere to.

## 3  An Under–Constrained Architecture as a Generic Architecture

The difference between an architectural style and an under–constrained architecture is a subtle one. The difference is fundamentally the difference in the completeness of the architectural description. A style is meant to focus on certainly critical features and isolate them from non–essential and non–stylistic features. There is no requirement for completeness of an architectural description in any way.

With an under–constrained architecture the idea is to capture the product line as completely as possible but in such a way that the variations are not ruled out by overly constraining the architecture. The variance is within the confines of the architectural constraints, not within the aspects that have not been defined.

This approach goes a long way towards solving the weaknesses of the stylistic approach in terms of analysis and planning at the product line level. Further it is much easier to create a product architecture from the product line architecture. However, it is still not a simple matter to produce the product architecture from the product line architecture (it is still primarily a creative process as with styles

but one which has marrower bounds) and one will have to analyze the product architecture to ensure its conformance to the product line constraints.

This approach seems to be an appropriate one to use if the primary difference among the products is something like performance and in which the functionality is primarily the same. On the negative side, extending the product line is a significantly more constraining task. Unless you evolve the product line architecture, the new products must be definable within the current constraints. In evolving the under-constrained product line architecture, care must be taken in its expansion not to inadvertently nullify current products as constituents of the line through the addition of further components or constraints. Constraint relaxation, of course, does not cause such a problem.

## 4  A Variance–Free Architecture as a Generic Architecture

Again the differences between this and the preceding ones are subtle. Here the architecture is not under-constrained. It is instead a fully described architecture but one in which the variances among the products are not considered to architecturally important – that is, the product differences are an issue of design and implementation, not an issue of architecture.

This approach is useful when your product line spans a significant range of options with respect to a particular aspect. One such example is that of whether the system is centralized or distributed. If the products range from simple centralized systems through to complex multi-processor and distributed processor systems. then this characteristic of the system might well be one that you want to bury in the infrastructure and not have as an important architectural issue. In this case, you might want to have a distribution independent architecture. Distribution then becomes an implementation or even a administrative issue, but not an architectural issue.

What is interesting in this case is that there is a significant implication for the implementation to support this kind of variance independence. To make the architecture independent of issues of distribution implies a class of architectural components which will support that independence.

Another example might be platform independence. Here again, there is an implication about what the structure of part of the architecture must be in order to bury the actual platform specific aspects in the design and implementation rather than have them visible at the architectural level.

There is a significant appeal in this approach. Analysis and planning can be done at the product line architecture level. If the right product characteristics are made independent of the architecture, then new products can be derived from the product line architecture with relative ease merely by providing the appropriate implementation specific components in the design and coding phase in such a way that they conform to the product line architecture. The individual product architecture is the product line architecture; there is no derivation

involved. Evolution of the product line architecture implies evolution of the product architectures.

Because of the identity of the product and product lines architectures, issues of analysis and planning at the product line level apply to the product level.

The downside of this approach is that it may not be possible to isolate all the variations in this way. Certain properties such as distribution, fault-tolerance, etc may be amenable to this, but differing functionality may not be.

Another negative aspect is the standard specification problem of talking about what is not there.

## 5 A Parametric Architecture as a Generic Architecture

A standard approach for generalizing is that of parametric abstraction. The parameterized component is then applicable across a wide range of arguments (in programming languages defined typically by types). The limits of applicability depend on the constraints that are checked on those arguments. That partly depends on the type system and what is allowed as a first class parameter types. For example, in Ada generics, the range of types usable as parameters is larger than for functions and procedures. In macro languages there are typically no constraints at all. But then there is no guaranteed substitution safety either.

The utility of this approach is the same as for packages and operations: the architecture specification defines a family of possible instantiations and for which the properties of the product line can be ensured for the various instantiations. The variations required for each possible product in the line are well-defined and known. Moreover, the instantiation of a specific product architecture is a well-understood technology and the instance can be derived automatically from the argumented product line description.

Here again, analysis and planning are doable at the product line rather than the product level.

Evolution of the parameters may seriously affect individual product architectures. If the evolution is limited to broading the types of the parameters, or perhaps upward-compatibly extending the parameters, then the individual product architectures should remain valid.

There are two limiting factors. First the kinds of the parameters allowed may seriously affect how well the generic architecture serves to cover the necessary products. If the kinds of first class objects are too limited, then one may not have sufficient descriptive power to satisfactorily describe the product line. Second is the question of whether parameterization covers all the kinds of variation that one might need to have among the products in a product line. We have seen examples above that suggest that parametric approaches are not sufficient in and of themselves.

# 6 A Service Oriented Architecture as a Generic Architecture

One of the typical kinds of problems found in developing such large and complex systems such as telephone switches is the need to provision the various products with different features. Provisioning these systems is not the kind of thing that can be done with parametric or variation independent approaches. One can always of course do it with either styles or under–constrained descriptions, but that does not help much if one wants these provisioned features to be architectural features.

Thus an approach to describing a product line architecture is one in which the various architectural services that may be provisioned are defined as part of the architecture and are then selected in an instantiation process to define a particular product. One advantage of this approach is that the possibilities are explicit in a more tangible way than in a parametric approach. Moreover, if done properly, the architectural dependencies of these services are also made explicit and the implications of choices are thus more explicit.

As with the parametric approach, instantiation is accomplished with well-understood technology. Analysis and planning both can be done relative to the product line description with the added advantage that the planning of a specific product can be derived from the product line planning itself via the selection mechanism of provisioning.

As long as the evolution of the product line architecture is done via the addition of new services, existing product architectures will remain valid instances of the new product line architecture.

While this goes a long way towards a useful approach for provisioned products, it is likely to be insufficient in itself for a complete product line specification.

# 7 Putting The Pieces Together

I think it is clear at this point that a comprehensive approach to defining a generic architecture for a product line requires all of these different ways of addressing various product line issues.

Styles are certainly needed for aspects of the product line that are orthogonal to the specific component structure. For example, one may want to define a style for initialization or fault handling that must be satisfied by all the components in a product line to ensure appropriate cross–product use.

Under–constrained descriptions always provide a wider degree of flexibility than over–constrained ones. Clearly some aspects of a product line will be best served by this approach where large degrees of design and implementation freedom are useful to respond to such things as changes in technology.

The variation–independent architecture is certainly needed where you want to delay such considerations as platform or distribution until build time or even execution time.

Parametric and provisioning approaches are again obviously useful for various kinds of generic descriptions and provide the most direct means of deriving their product architecture from the product line architectures.

## 8   Summary

I have considered a variety of useful ways of 'genericizing' architectural descriptions (or prescriptions). I claim that a generic architecture is a fundamental requirement for a product line and that each of these approaches is needed as a means of defining some important elements in such a generic architecture.

# Generic Software Architecture Descriptions

Dewayne E. Perry

*Bell Laboratories*
*Room 2A-429*
*600 Mountain Ave*
*Murray Hill NJ 07974*

*dep@research.bell-labs.com*
*www.bell-labs.com/~dep/*

September 1998     **Newcastle '98**     Lucent Technologies

---

# Models of SW Architecture

- Perry & Wolf 89/92 model of SWA
- SWA = ( Elements, Form, Rationale )
- Elements : process, data and connecting
- Form is the set of properties of, and relationships among, the elements
- Rationale is the justification for the elements and form

September 1998     **Newcastle '98**     Lucent Technologies

---

# Outline

- *Software Architecture - Overview*
- Product Line Architecture - Overview
- PLA Description Issues and Generic Approaches

September 1998     **Newcastle '98**     Lucent Technologies

---

# Styles

- An incomplete architectural prescription
- Focuses on certain aspects of the architecture
  - architectural elements
  - formal characteristics
  - constraints on architectural elements
  - constraints on formal characteristics

September 1998     **Newcastle '98**     Lucent Technologies

## Styles

- Problem: Restrict the architectural structure
  - for example, strict layering of the architecture
- Solution: layered architecture style
  - constrain the interactions
    - any interaction at elements on the same level
    - no interactions at more than one level away
    - level below: initiate interactions only
    - level above: react interactions only

## State of Current Work

- Pretty much agree about process, data and connecting elements as first class entities
- Models differ primarily with respect to Form
- Few models pay attention to rationale
- Styles tend to focus on element and form restrictions

## Styles

- Useful rule of thumb: a style for a domain
- Problem: multiple domains in any significant architecture
- Challenge: integrating the styles consistently

## Current Approaches to Form

- Configurations
- Types
- Patterns
- Properties

VII.10

## Configuration as Form

- Characterization
  - Basic box and lines approach
  - Components may be processes, subsystems, etc
  - Connections are defined by Provides/Requires clauses
- Approach to Style
  - Tend not to be interested in styles
  - Except in the context of dynamic arch's

## Types as Form

- Approach to Style
  - Emphasis on the basic classes or types of components and connectors
  - Perhaps, a slight more emphasis on connectors
  - Eg, pipes and filters; blackboard architecture

## Types as Form

- Characterization
  - Typically, an historical approach
  - Look for types and classes of architectural objects
  - Often organized hierarchically

## Patterns as Form

- Characterization
  - Emphasis on patterns of interactions
  - Tendency to focus on connections with components as endpoints
- Approach to Style
  - Architectural instances are specializations of styles

## Properties as Form

- Characterization
  - Properties of (or constraints on) data, process and connecting elements
  - Relationships among data, process and connecting elements
- Approach to Style
  - Selection of some critical elements
  - Selection of some properties and relationships
  - Constraints on properties and relationships

## Product Line - Overview

## Outline

- Software Architecture - Overview
- *Product Line Architecture - Overview*
- PLA Description Issues and Generic Approaches

## Product Line - Basic Aspects

- Begin with product instances
  - legacy based
  - use architecture recovery processes
- Focus on appropriate business domain
  - use domain specific architectural processes
  - map from recovered to domain architecture
- Abstract/Generalize to Product Line Architecture

## Product Line - Issues

- *Product Line Reference Architecture*
- Product Line Processes
- Asset Base
- Supporting Technology
- Organizational Issues

## PLA Description Issues

- What generic features do you need
- Relationships between PLA and PIA
  - Derivation
  - Conformity
  - Analysis
  - Planning
- How is evolution of PLA supported

## Outline

- Software Architecture - Overview
- Product Line Architecture - Overview
- *PLA Descriptions Issues and Generic Approaches*

## PLA and Generic Descriptions

Claim:

     *Generic descriptions are necessary for product line architectural descriptions*

## Generic PLD Approaches

- Style description
- Under-constrained description
- Variance-free description
- Parametric description
- Service/provision oriented description

## Style - Advantages

- Minimalist approach
- Add new products easily
- As long as they conform to style
- Some project planning for the PLA applies to the product instance architecture (PIA)

## Style

- Intuitive appeal
- Captures essential characteristics
  - basic components
  - minimum interactions
  - basic constraints
- Ignores variation

## Style - Disadvantages

- Not easy to refine PLA into PIA
  - by extension, addition
- PLA conformity analysis required
- When PLA evolves, must revalidate PIA conformance

## Style - Evaluation

- Possible, but not adequate
- better uses of styles than for PLA

## Under-Constrained - Advantages

- Easier to create PIA from PLA than Styles
- Analysis at PLA level applies to PIA level
- Planning at PLA level applies to PIA level
- Evolution via constraint relaxation easy

## Under-Constrained

- Difference in completeness
  - style focus: critical features, eliminate non-essential, non-stylistic
- Capture PL as completely as possible
- With variations not ruled out by overly constraining the architecture
- Variance within constraints, not within the aspects not defined

## Under-Constrained - Disadvantages

- Extending the PLA is a significantly more constraining task
- May not be possible to define all new products within constraints
- PLA evolution may cause conformity problems

## Under-Constrained - Evaluation

- Seems appropriate where primary difference is something like performance where the functionality remains the same
- Too confining for variance often needed for individual products

## Variance-Free - Advantages

- Analysis and planning at the PLA level
- Product variance depends on implementation and not on architecture
- PLA is the PIA
- Evolution of the PLA means evolution of the PIAs

## Variance-Free

- Architecture is not under-constrained
- Variance is not considered architecturally important
  - product difference a design or implementation issue not an architectural one
  - eg, platform or distribution independence
- There are implications for the PLA

## Variance-Free - Disadvantages

- Standard specification problem of talking about what is not there
- May not be able to isolate all variance this way

## Variance-Free - Evaluation

- Useful for range of options for a particular aspect (eg, fault tolerance, distribution ...)
- But may not be able to account for variance in functionality

## Parametric - Advantages

- Variations well-defined and well-known
- Instantiation of PIA from PLA is well-understood (possibly automatic)
- Analysis at PLA level
- Planning at PLA level
- Evolution by relaxing constraints or by upward compatible extensions OK

## Parametric

- Standard approach: parametric abstraction
- Limits depend on the constraints on the arguments
- Defines a family of possible instantiations

## Parametric - Disadvantages

- Kinds of parameters allowed may seriously affect how well the PLA covers the necessary variance
- Incompatible parameter evolution generates conformance problems

## Parametric - Evaluation

- Means of abstraction well-understood
- Instantiation well-understood
- Good analysis and planning properties
- May not cover all forms of variance

## Service-Oriented - Advantages

- Instantiation is by selection
- Possibilities are explicit
- If done properly, architectural dependencies among services are explicit
- Analysis at PLA level
- Planning derived from PLA via selection
- Evolution via addition OK

## Service-Oriented

- In large complicated systems often need to provision individual products with different features
- Not doable with parameters or variation independence
- Architectural features selectable

## Service Oriented - Disadvantages

- Evolution via change/deletion causes conformity problems
- May not know all the services needed in advance

## Service-Oriented - Evaluation

- Simple/effective way of managing product line
- Likely to be insufficient for complete PLA

## Putting It Together

- Parameters where the ranges of solutions are well understood
- Provisioning where the possibilities are enumerable

## Putting It Together

- Comprehensive approach would require all these forms of generic description
- Styles useful for aspects distributed across sets of architectural components
- Under-constrain where flexibility is needed such as changes in technology
- Variations independence for delayed binding

# DISCUSSION

**Rapporteur**: Dr Paul Ezhilchelvan

## Lecture One

When Professor Perry was presenting the usefulness of Asset Base (AB) in developing Product Line Architecture (PLA), Professor Randell questioned about the problems and the difficulties involved in building an AB can pose, when a big organisation is developing a PLA. The speaker replied by pointing out that AB is meant to contain components that are re-used and deemed to be reusable. When building a multi-media application involving v-mail, e-mail etc, some operations remain identical throughout. Also, different interfaces have different implementations. Such components are accumulated in the AB during the PLA development. To Professor Randell's observation that PLA and AB are not directly connected in the Life-cycle figure presented, Professor Perry pointed out that the relation is indirect and mainly via the PL Processes Box.

Professor Brooks wondered about the effects of inevitable inconsistencies among the pieces in the AB. Professor Perry attributed these inconsistencies to different standards among products of same type which the PLA must take into account. Ideally, one would prefer a platform that is independent of communication architecture. But each communication medium has different standards which lead to different implementation.

Dr Kay wanted to know who the PLA is being developed for. It is primarily for the company who is responsible for products in product line; it also remains useful for project personnel managing the project. Basically, the aim of PLA is to disseminate much information about the product and its development so that the interval taken to produce and market the product is shortened. Professor Shaw observed that while one benefit of PLA to the company is faster production and marketing, it also simplifies incorporation of upgrade paths. Mr McVitie, observing the discussions on the usefulness of PLA, wondered whether PLA is to be developed with both AB and product line in mind. Dr Perry stopped short of endorsing the view, reiterating his view that PLA is to describe how best a product can be built.

Professor Randell compared the certain aspects of under-constrained PLA to specifying systems in terms of what they should not do (negatively) rather than what they should do (positively). The speaker opined that negatively specifying things work very well in certain cases.

Professor Johnson wanted to know how Product Instance Architecture (PIA) is described: formally or in terms of a programming language. The speaker replied that it can also be done in terms of constraints involving notations.

Dr Kay wondered how certain restrictions, such as a product which can only be run on some versions of Unix, can be expressed. It can be done through constraints, the speaker replied and elaborated on that point.

Professor Randell was of the opinion that it is the level of abstraction that distinguishes functional properties against non-functional ones. As an example, he cited, some of what can be non-functional aspects for a network user are functional aspects for a network manager. The speaker used these terms in the following sense: system activities/properties that can be clearly stated come under functional aspects, and others under non-functional aspects. Professor Johnson recalled that Fault Tolerance issues were once regarded as non-functional aspects some twenty years ago. Now, as we know how to specify them, they are a part of functional aspects of a system.

Professor Johnson opined that what is being proposed is the right thing and weighed the cost versus benefit. He said that many companies cannot and do not develop a PLA before production. They often have PLA in the form of power-point slides. He wanted the speaker

to shed light on the cost of developing PLA and the relative difficulties in working with it. The speaker replied that the cost incurred is not much in his own company and advantages are far more, and the architecture is thought to be an important ingredient in the development process. What is often missed is the architecture update as the product undergoes changes. Professor Shaw observed that the integration of new technologies into the development process may also be simplified.

Professor Malek observed that often there is a compromise between the rush to product marketing and maintaining product quality. He wondered whether having PLA enables or prevents such compromises. The speaker replied that PLA improves both and avoids the need to make such compromises. Professor Vogt wanted to know the influence of the new Telecom architecture TINA had in the speaker's thinking and considerations. The speaker was not aware of the TINA details and it was agreed to pursue further discussions about TINA off-line.