

Widening the Scope of Software Product Lines – From Variation to Composition

Rob van Ommering¹ and Jan Bosch²

¹Philips Research Laboratories
Prof. Holstlaan 4, 5656 AA Eindhoven, The Netherlands
Rob.van.Ommering@philips.com

²University of Groningen, Department of Computing Science,
PO Box 800, 9700 AV Groningen, The Netherlands
Jan.Bosch@cs.rug.nl

Abstract. Architecture, components and reuse form the key elements to build a large variety of complex, high-quality products with a short lead-time. But the balance between an architecture-driven and a component-driven approach is influenced by the scope of the product line and the characteristics of the development organization. This paper discusses that balance and claims that a paradigm shift from variation to composition is necessary to cope with an increasing diversity of products created by an ever-larger part of an organization. We illustrate our claim with various examples.

Keywords: product population, software components, variation, composition.

1 Introduction

As a software engineering community, we have a long history of building software products: we have been doing this for half a century now. *Architecture* has always been an important success factor [7], though it has not received the attention that it deserves until the last decade. *Components* (or modules) serve as a means to cope with complexity (divide and conquer), and to streamline the evolution of systems [24]. We have been *reusing* software for a long time: operating systems, graphical and mathematical libraries, databases and compilers are rarely written from scratch for new products.

It is therefore not surprising that people developed the vision that the future of software engineering lies in the use of reusable components: new applications would be built by selecting from a set of existing components and then clicking or gluing them together. McIlroy was one of the first to advocate this idea [16], and many have followed since. We have indeed seen some successes in this area, for instance in the use of ActiveX controls (components) and Visual Basic (glue) to quickly build interactive applications for PCs. But we have also seen many failures, especially when reuse is applied *within* companies to create a diversity of products.

This is where software product lines enter. They provide proactive and systematic approaches for the development of software to create a variety of products. Again, software product lines have been around for decades, but have only recently begun to

receive the attention that they deserve. Note that product lines are not specific to software only: they are also commonplace in the car, the airplane and the food industry [8].

Most software product line research focuses on an early analysis of commonality and variation [1], followed by the creation of a variant-free architecture [25] with a sufficient number of variation points [13] to deal with diversity. While this works well for small product families in a small organization, we have experienced problems when applying this to larger ranges of products, especially when development crosses organizational boundaries.

This paper is about widening the scope of product lines to encompass a larger diversity of products (a *product population* [21]), where the development takes place in different organizational units but still within the same company. We believe that the successful building of product populations is the next hurdle to take in software engineering. Figure 1 illustrates this, showing on the one hand that the focus of our community shifts over time from products to components to families to populations, and on the other hand that we are seeking a balance between variation (as in product families) and composition (as in software components).

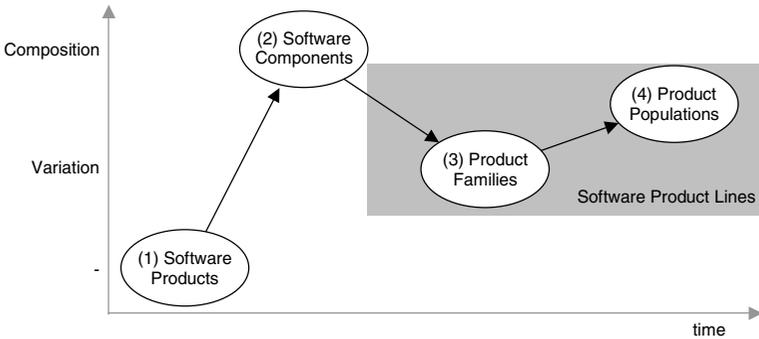


Fig. 1. Towards product populations

This paper is organized as follows. First we recapitulate the main drivers for software product lines. Then we discuss the influence of product scope and characteristics of the organization on software product lines. We proceed with introducing variation and composition as separate dimensions, and then explain variation and composition in more detail. We end with some concluding remarks. We take examples from the domain of one of the authors, the control software for consumer electronics products, but include examples from other domains as well.

2 Why Use Software Product Lines?

Let us first summarize *why* we want to create software product lines. In Figure 2, we show the main four drivers as we identified them within Philips at the left-hand side. We show three key solutions in the middle. But it's the combination of these three that forms the most promising solution to our problems: software product lines.

In the next subsections, we discuss each of the drivers and their relation to the solutions. Please note that arrows in Figure 2 show main dependencies only; arguments can be found for other dependencies as well, but adding them would obscure the main flow of the picture.

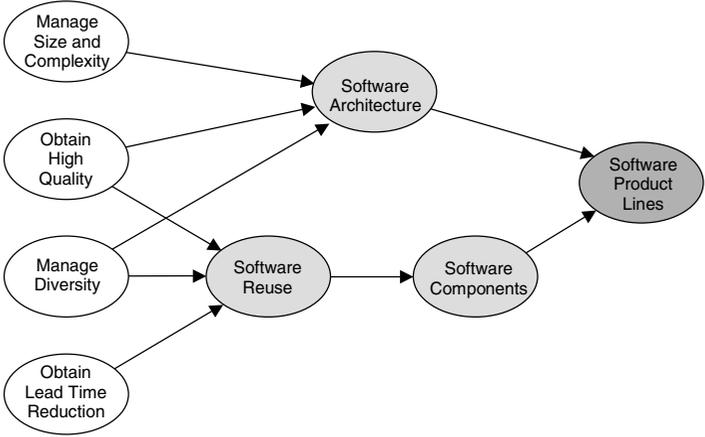


Fig. 2. Main drivers for software product lines

2.1. Manage Size and Complexity

Software plays a role of ever-increasing importance in many products of today. In Philips televisions, software was introduced in the late 70s when a 1-kilobyte 8-bit microprocessor was first used to control the hardware signal processing devices. The size of the software grew following Moore’s law: today, a TV typically has 2 megabytes of software. The same trend can be observed in other products, though the point in time when software starts to play a role differs. Typically, televisions run 10 years behind PCs, while shavers run 10 years behind televisions.

The best – perhaps only – way to manage size and complexity is to have a proper architecture. A software architecture defines, among other things, a decomposition of a system in terms of components or modules [12]. One common misconception (at least in our surroundings) is the assumption that components defined to manage size and complexity automatically help to satisfy the other drivers listed in Figure 2, most notably diversity. Experience has shown is not the case: defining components to manage diversity is an art per se.

2.2. Obtain High Quality

The second driver in Figure 2 is quality: the more complex systems become, the more difficult it is to maintain a high quality. Yet for systems such as televisions, high

quality of the software is imperative¹. Television users do not think that they are communicating with a computer, so they will not accept crashes. Also, although lives may not seem to be involved, a TV catching fire can have disastrous effects, and it is partly the software that prevents this from happening.

Quality can be achieved by having a proper architecture that guarantees (or at least helps to satisfy) certain quality attributes. Quality can also be achieved by reusing software that has already been tested in the field.

2.3. Manage Diversity

The third driver in Figure 2 is the first one specific to product lines: managing the efficient creation of a variety of products. For consumer products, we have ‘small-scale’ diversity within a family of televisions in areas such as the output device, the price setting, the region, the user interface, the supported standards, and interconnectivity. We have ‘large-scale’ diversity if we want televisions, VCRs, DVD/CD players/recorders, and so forth to be part of one product line approach. There is indeed sufficient commonality to justify the sharing of software between those products. We use the term *product population* to denote the large-scale diversity. Other examples of product populations can be found in the medical domain [14] and for printers and scanners [3].

There are actually two schools of thought for managing diversity, which we believe are two extremes in a spectrum of solutions. One is to build an architecture that expresses the commonality between products, with variation points to express the diversity. The other is to rely on reuse of software components for the commonality, but to use different compositions of these components to implement diversity. This paper is about balancing between these two solutions (variation and composition) depending on the scope of the products and the characteristics of the organization.

2.4. Lead-Time Reduction

The fourth driver in Figure 2 is lead-time reduction. Although in principle this is also relevant for single product development, we think that it is especially relevant in product line contexts, since the goal of product lines is to produce a range of products quickly. The main way to obtain lead-time reduction is to reuse software. Experience has shown that higher level languages, modern development methods, better architectures, and advanced software development environments have indeed improved software productivity, but not with a factor that is sufficient to obtain the speed required in product lines.

2.5. Does Software Reuse Imply Software Components?

Figure 2 may suggest that software reuse implies the deployment of reusable software components. Indeed the reuse of software components is one (compositional)

¹ Repairing software shipped in televisions sold to a million customers is as difficult as in a rocket sent a million miles into space.

approach to building product populations, and one that we will examine in this paper. But we should note that we can also have reuse *without* software components, for instance when using inheritance in an object-oriented (OO) framework. And, perhaps surprisingly, we can also build a product line that utilizes software component technology without actually reusing the components! An example of the latter is a framework with component plug-ins, where the plug-in components implement *diversity* rather than *commonality*. Figure 3 illustrates this: the left-hand side shows a reusable framework with product-specific plug-ins, while the right-hand side shows reusable components, composed in a product-specific way.

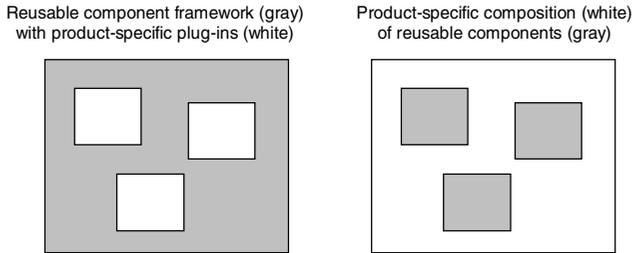


Fig. 3. Variation (left) and composition (right) are complementary approaches

2.6. Finding the Balance

In the rest of this paper, we concentrate on the balance between the use of variation and composition in software product lines, depending on the scope of the product population and the characteristics of the organization building it.

3 The Influence of Scope on Software Product Lines

As we have said before, many software product line approaches advocate an early analysis of the commonality and variation of the products to be built, resulting in a variant-free architecture with variation points. This is indeed a valid approach, but there may be circumstances in which the preconditions for such an approach are not satisfied. We shall discuss two categories of these in the following sections: (1) product related and (2) organization related.

3.1. From Product Family to Product Population

Traditional product lines cover families in which products have many commonalities and few differences. This enables the creation of a variant-free architecture with variation points to implement the differences. The variation points typically control the internals of a component, but they may also control the presence or absence of certain components, the selection of a concrete component to implement an abstract component (a place holder), or even the binding between components.

A typical product family in the consumer domain is the set of televisions (direct view, flat screen, projection) produced by a single company². Let us now extend the product range to include video recorders, CD and DVD players, CD, DVD and hard disk recorders, audio amplifiers, MP3 players, tuners, cassette players, and so forth. These products still have enough in common to warrant reuse of software. But they are also sufficiently different to require different architectures.

The latter is actually not a fundamental problem. In principle, it seems perfectly possible to create a single architecture for a ‘super consumer product’, that would be the combination of all products mentioned above, and to strip this architecture for any subset to be sold. While this is certainly possible in hindsight, it is very difficult to predict which combinations of audio and video elements will actually be interesting next year, and to include that in the architecture in advance.

Instead, the market itself thrives on novel combinations of existing functions (for instance a mobile telephone with built-in digital photo camera). We believe that in such an evolving ‘compositional’ market, the software architectures must also be compositional. Put differently, variation can be used to implement diversity known ahead of time, while composition can be used for open-ended diversity.

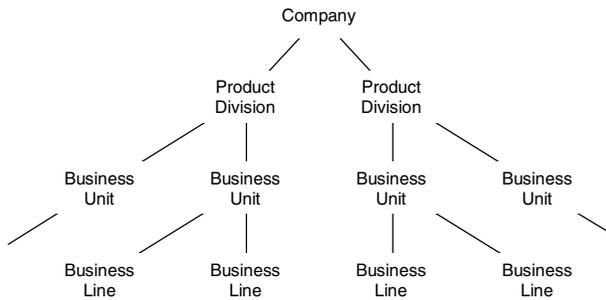


Fig. 4. Schematic Organization of a Large Company

3.2. From Business Line to Product Division and Beyond

The second argument in favor of composition has its roots in organization. A large company like Philips is typically organized in product divisions (Semiconductors, Consumer Electronics), which are split into business units (video, audio), which again are split into business lines (flat TV, projection TV). See Figure 4 for an illustration.

Within a business line, setting up a software product line is generally relatively easy. One person is in charge of the business, which has its own profit and loss calculation, and a product line is demonstrably an efficient way to create a variety of products. The products are reasonably alike (e.g., all high-end Flat TVs). Moreover, the product roadmaps in a business line are usually carefully aligned, thus enabling the creation and use of shared software. A single architecture with variation points is very feasible.

² For technical reasons, low-end televisions usually belong to a different family than high-end televisions.

Compare this with a business unit. Although – of course – one person still is in charge, profit and loss responsibility is delegated to the business lines, making sharing software between business lines at least cumbersome. Each business line manager has his own agenda for planning software development. Protection of intellectual property (IP) is sometimes an issue, but more frequently the misalignment of roadmaps causes problems: one business line is changing software with its product release still far away, while another product line needs a mature version of the software for its immediate next release. A careful and explicit roadmapping is required here [22]. Products show more differences, but in general a single architecture with variation points is still be feasible.

Problems get worse at the level of a product division. Not only do the products become more different (e.g., TV versus set-top box), but the managers become more apt to look after themselves instead of after their common interest, the roadmaps become harder to align, and the cultures between organizations become different. In Philips, the TV business is a traditional and stable consumer business (low profit per product but high number of products, targeted directly at the person on the street), while the set-top box business is new and evolving (currently targeted at large providers who deliver the boxes free with their services). TVs are typically programmed in C in a few megabytes, while set-top boxes are programmed in C++ in many megabytes.

Finally, product divisions act as independent companies (at least at Philips), and while sharing of software between them is in many cases not feasible (as between medical systems and consumer electronics), it may in certain situations still be desirable. Philips Semiconductors produces for instance many of the chips for Philips Consumer Electronics, and it also sells the same chips to third parties, so certain software can be better leveraged if developed by the semiconductors division. This puts extra demands on protection of IP. Cultures between product divisions are even more different, and the levels of software engineering maturity may differ.

3.3. How to Share Software in Large Organizations?

It should be obvious by now that the top-down creation of a single architecture with variation points to cover all consumer products within Philips is not feasible. Instead, we should find a way to let different business lines create the software components in which they are specialized (e.g., the business line DVD creates a DVD unit), and find ways to ‘arbitrarily’ compose them.

Should we then strive for a bottom-up approach as advocated by many believers in software component markets (i.e., (2) instead of (4) in Figure 1)? The answer is no: we still need a careful balance between top-down and bottom-up, because we need to make sure that components fit easily, as explained in the next section.

3.4. Balancing between Architecture- and Component-Driven Development

Classical product and product line approaches ((1) and (3) in Figure 1) are architecture-driven, top-down, planned, intra-organization, with all software to be developed in the project. Component market evangelists ((2) in Figure 1) on the other hand, claim that software development should be component-driven, bottom-up,

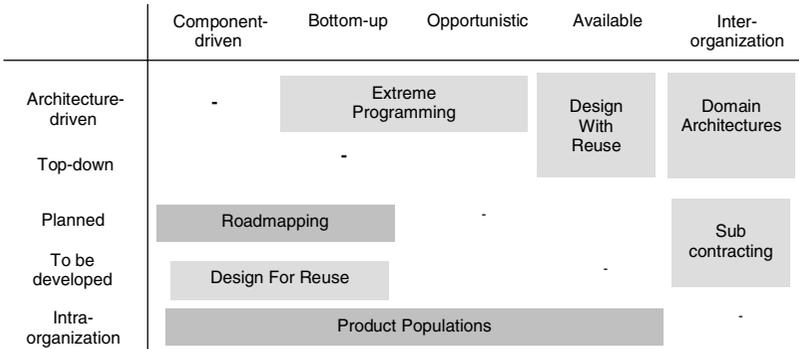


Fig. 5. Balancing between Architecture- and Component-Driven Development

opportunistic, inter-organization and using components that are already available. Figure 5 shows both sets of these opposite characteristics along different axes. Note that they are not all dependent – we can make cross combinations:

Extreme programming could be seen as an opportunistic bottom-up yet architecture-driven method.

- Top-down, architecture-driven development deploying available software is sometimes called *design with reuse*, while on the other hand component-driven bottom-up development is sometimes called *design for reuse*.
- A top-down, architecture-driven development spanning multiple organizations is possible, but only if there is consensus on a *domain architecture*.
- The planned development across organizations of new software is often called *subcontracting*, and should not be confused with the development of reusable software³.

Our vision is to achieve a component-driven, bottom-up, partly opportunistic software development using as much as possible available software to create products within our organization; we call this a *product population* approach. In reality, life is not so black and white. More specifically, we believe that such an approach should be:

- Component-driven (bottom-up) with a lightweight (top-down) architecture, since absence of the latter would result in architectural mismatches [10].
- Opportunistic in using available software but also planned when developing new software, with the latter carefully roadmapped [22].
- Intra-organizational but spanning a large part of the organization.

In a way, building product populations is more difficult than producing and utilizing software components on a component market. In the latter case, the laws of economics apply: component vendors can leverage development costs over many customers (as compared to a few in product populations), while component buyers can find second sources in case vendors do not live up to expectations. One would expect that the same economical model could be used *within* an organization, but in

³ Within Philips, we have seen a single team that had as a mission the development of reusable software act in practice as two teams subcontracted for two different product projects.

practice this does not work: contracts within a company are always less binding than those between companies.

3.5. Third-Party Component Markets

Does this mean that we do not believe in third-party component markets [31]? Yes, we do, but not as the only solution.

First of all, most software components on the market tend to implement generic rather than domain-specific functionality. While this ensures a large customer base, it does not help companies like Philips to build consumer products. Put differently: of course we reuse generic components, such as a real-time kernel from third parties, but there is still a lot of software to be created ourselves. This situation may change with the arrival and general acceptance of TV middleware stacks.

Secondly, many companies distinguish between *core*, *key*, and *base* software. Base software can and should be bought on the open market. Key software can be bought but is made by the company itself for strategic reasons. Core software cannot be bought because the company is the only company – or one of the very few – actually capable of creating such software. But there may still be a need to share this software within the company.

Thirdly, many components out on the open market are still relatively small (e.g., buttons, sliders), while the actual need for reuse is for large components. Of course there are exceptions: underlying Microsoft's Internet Explorer is a powerful reusable HTML displaying *and* editing engine that can be used in different contexts (such as Front Page and HTML Help). But such examples are still exceptions rather than the rule.

Finally, to be accepted in an open market, functionality must be mature, so that there is consensus in the world on how to implement and use such functionality. This process may take well over 20 years. Consider operating systems and windowing systems as examples.

4 The Dimensions of Variation and Composition

Figure 6 shows variation and composition as two independent dimensions. The variation axis is marked with 'As is' (no variation), parameterization, inheritance, and plug-ins as subsequently more advanced techniques of variation. The composition axis is marked with no composition (yet reuse), and composition. The table is filled with some characteristic examples, which we discuss briefly below. We also zoom into variation and composition with more elaborate examples.

4.1. Plain 'Old-Fashioned' Libraries

The classical example of software reuse is through libraries. Examples include mathematical libraries, graphical libraries, and string-manipulation libraries. The functions in the library are either used 'as is' or can be fine-tuned for their specific application using a (procedural) parameter mechanism.

| | | | |
|---|------------------|-----------------------------|--------------------------|
| | | COMPOSITION → | |
| | | Reusable but not composable | Composable |
| V A R I A T I O N ↓ | ‘As is’ | Libraries | LEGO |
| | Parameterization | | Visual Basic |
| | Inheritance | OO Frameworks | Frameworks as Components |
| | Plug-ins | Component Frameworks | |

Fig. 6. Variation and Composition as Two Independent Dimensions

Typically, the software deploying the library includes the header file(s) of the library, and thus becomes specifically dependent upon the library. It is possible to abstract from the library by defining a library-independent interface, for example, POSIX⁴. The build process may then select the actual library implementing this interface. However, in practice, most software built on top of libraries is still specifically dependent upon those libraries.

This is especially true for the library itself. Often, a library uses other libraries, but few of them allow the user to select which underlying libraries to use. Moreover, while some libraries are relatively safe to use (they define specific functions only, such as ‘sin’ and ‘cos’), other libraries may introduce clashing names (e.g., String), or have clashing underlying assumptions (e.g., they both define the function ‘main’).

The conclusion is that while libraries are a perfect low-level technical mechanism, more advanced techniques are necessary to achieve the variation and composition required for product lines.

4.2. Object-Oriented Frameworks

An OO framework is a coherent set of classes from which one can create applications by specializing the classes using implementation inheritance. An example is the Microsoft Foundation Classes framework [26]. The power lies in the use of inheritance as a variation mechanism, allowing the framework to abstract from specific behavior and to be implemented by the derived classes. A base class usually has a very intricate coupling with its derived classes, with the base class calling functions implemented in the derived class and vice versa.

This intricate coupling is also one of the weaknesses of this approach. A new version of the framework, combined with old versions of the derived classes may fail due to subtle incompatibilities; this is known as the *fragile base class problem* (see 29 for a treatment). Also, since many frameworks are applications themselves, combining two or more frameworks becomes virtually impossible.

⁴ POSIX stands for Portable Operating System Interface.

4.3. Component Frameworks

A component framework is part of an application that allows plug-in components to specialize behavior. Component frameworks resemble OO frameworks; the main difference is the more precise definition of the dependencies between the framework and the plug-in components. The plug-ins can vary from quite simple to arbitrarily complex; the latter case enabling the creation of a large variety of applications.

Component frameworks share some of the disadvantages with OO frameworks. First of all, the plug-in components are usually framework dependent and cannot be used in isolation. Secondly, the framework often implements an entire application and therefore cannot be defined easily with other frameworks.

4.4. LEGO

The archetype non-software example of composition is LEGO, the toy bricks with the standard ‘click and play’ interface. Indeed, LEGO bricks are very composable, but they support no variation at all! This results in two phenomena.

First of all, it *is* possible to create arbitrary shapes with standard LEGO bricks, but one must use thousands of bricks to make the shapes smooth (see Figure 7 for a pirate’s head). As the components are then very small compared to the end product, there is effectively not much reuse (not more than reusing statements in a programming language).



Fig. 7. The pirate’s head with ordinary LEGO® bricks⁵

To solve this, LEGO manufactures specialized bricks that can only be used for one purpose. This makes it possible to create smooth figures with only a few bricks. The down side is that specialized bricks can only be used in specialized cases. A solution could be parameterized bricks, which indeed is a mechanical challenge! Maybe a LEGO printer, a device that can spray plastic, could be an answer here!⁶

[®] LEGO is a trademark of the LEGO Group.

⁵ From <http://www.lego.com>

⁶ Strictly speaking, a LEGO printer would be the equivalent of a code generator, which is a technique for building product lines that we do not tackle in this paper, although it can be very fruitful in certain domains (cf. compiler generators). See also [9].

4.5. Visual Basic

The archetype software example of composition is Visual Basic (VB). Actually, Basic is the non-reusable “glue” language, where ActiveX controls are the reusable components (later versions of VB also allow the creation of new ActiveX controls).

One of the success factors of Visual Basic is the parameterization mechanism. An ActiveX control can have dozens, even hundreds, of properties, with default values to alleviate the instantiation process. Another parameterization technique is the event mechanism, which allows the signaling of asynchronous conditions, where the application can take product-specific actions.

Still, ActiveX controls are usually quite small (the Internet Explorer excepted), and it is not easy to create application frameworks (e.g., supporting a document view control paradigm such as the Microsoft Foundation Classes do).

4.6. Frameworks as Components

The ultimate solution for variation *and* composition lies – we believe – in the use of OO or component frameworks as components. We have seen that frameworks provide powerful variation mechanisms. If a framework only covers part of an application domain, and if frameworks can be combined arbitrarily (making them *true* components), then we can create a multitude of products by selecting and combining them and then specializing them for a product.

There are some examples of OO frameworks used as components [4], but no examples of component frameworks used as components yet. The basic mechanism for achieving composition is to make every context dependency explicit and bindable by a third party, as we will see in a later section.

5 Variation Further Explained

Let us zoom into variation first, and then discuss composition in more detail.

5.1. Why Do We Need Variation?

The basic argument for variation is the one underlying all product line development: we need to implement diversity. We can achieve this with reusable components, but an interesting dilemma pops up. The more useful we make a component (read: larger), the less reusable it becomes. Conversely, the more reusable we make a component, the less useful it may become, the extreme case being the empty component: extremely reusable while not very useful! See also [29], Fig 1.1.

The solution for this dilemma is *variation* as a way to fine-tune components when instantiating them into a product. The archetype example is the Visual Basic ActiveX control, which usually comes with a long list of properties that can be modified at design-time and/or runtime. These properties also have default values to alleviate the effort of product creation; otherwise the burden on the designer may be too high.

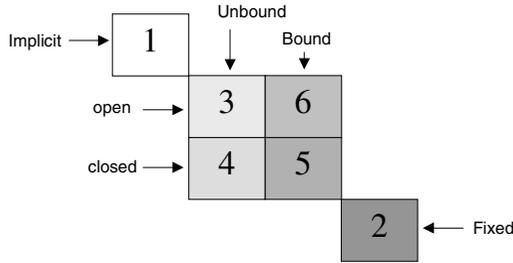


Fig. 8. A Life Cycle for Variation Points

5.2. Types of Variation

A variation point can be implicit or explicit, open or closed, unbound or bound, and flexible or fixed [4]. Not all sixteen combinations are meaningful. See Figure 8 for an overview. When development starts, a variation point is implicit; it must be made explicit during analysis and design. An explicit variation point can be open if the set of possible values can still be extended, or closed otherwise. A variation point is unbound if a value has not been selected, and bound otherwise. A variation point is fixed if it's bound, and the value cannot be changed anymore.

Traditional single-product development starts at (1) in Figure 8 and then immediately jumps to (2). Classical product-family development proceeds through (3), (4) and (5) (identify a variation point, identify its values, and then choose a value) before ending in (2). There is a tendency to make (5) the end point of development, allowing the variation point still to change at the customer site. There is even a tendency to make (6) the end point of development, allowing the customer to extend the set of values.

5.3. Example 1: Switches and Options

In our previous television software architecture, we distinguished between *switches* and *options*. A switch is a compile-time variation point (an `#ifdef`), closed, bound and fixed somewhere during development time. An option is a runtime variation point, a value stored in non-volatile memory and used in an if-statement, so closed yet unbound during development, and bound in the factory or at the user's home.

These variation points were used for adjustments of hardware, for the presence or absence of components, and in rare cases for selecting between components. The set of variation points was very specific for the product family and could not be used to achieve the large-scale structural variation that we needed to widen our family.

5.4. Example 2: The Windows Explorer

The Windows Explorer is an example of an application with many explicit, open, unbound and flexible variation points. It can be tuned into almost any browsing

application, though it cannot be used as a component in a larger application. In other words, it supports a lot of variation, but no composition.

The explorer can be tuned at a beginner's level and at an advanced level [18]. In the first case, simple additions to the windows registry suffice to change the file association (the application to be run for a certain type of file), the context menus or the icons. In the second case, handlers must be registered for specific shell extensions, and these handlers must be implemented as Component Object Model (COM) components. When such a handler crashes, Windows Explorer crashes, with all its windows including the desktop⁷.

The folder view of Windows Explorer can be modified by name space extensions; some FTP clients use this to add favorite sites to the name space. The folder view is part of the Explorer bar, which can also host other views (search, favorites, history). The right pane normally hosts a standard list view control, but it can also host the Internet Explorer to view (folders as) web pages, or any other user-defined window.

5.5. Example 3: A Framework Architecture

Wijnstra [32] describes a software architecture in the medical domain that deploys component frameworks to manage diversity. Figure 9 provides a schematic overview of this architecture. It consists of a fixed structure of *units*, COM components (mostly executables) running on a Windows NT computer. Each unit is a component framework in which COM components can be plugged to implement the diversity of the product family. Plug-in components may talk to each other, but the initial contact is made through the main units.

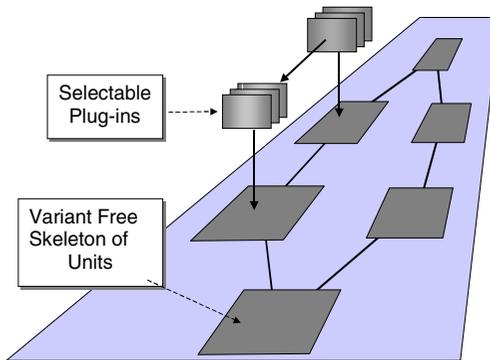


Fig. 9. Fusion

In this architecture, the fixed structure of units⁸ implements the common part of the family, while the plug-ins take care of the differences between the products. It is therefore an example of variation rather than composition (see the left-hand side of Figure 3). Of course, this is true if the plug-ins are relatively small and specific for one product only. In real life, some plug-ins are rather large and shared between

⁷ It is possible to run the desktop in a different process to prevent such crashes.

⁸ Some units are optional.

multiple members of the family. In fact, if one can create products by taking arbitrary combinations of plug-ins, then this system exhibits signs of composition rather than variation.

6 Composition Further Explained

We shall now discuss composition as a technique for managing diversity.

6.1. What Is Composition?

We speak of *composition* if two pieces of software that have been developed *without* direct knowledge of each other can be combined to create a working product. The components may have been developed in some common context, or they may share interfaces defined elsewhere. The essential element is that the first piece of software can be used without the second, and the second piece without the first.

The underlying argument for using composition to handle diversity is the large number of combinations that can be made if components are not directly dependent on one another. Note that composition excludes the notion of plug-ins, since these are directly dependent on the underlying framework, and cannot be used in isolation.

A nice analogy may be the problem of describing a language, where an infinite number of sentences must be produced with a finite (and hopefully small) set of ‘components’. Typically, such components are production rules in a grammar, where the production rules function both as a composition hierarchy as well as a set of constraints on the composition.

Garlan has illustrated how difficult composition can be in his well-known paper on architectural mismatches [10]. It appears that two components being combined must at least share some common principles to work together efficiently. We call this a lightweight architecture, but leave it to another paper to elaborate on how to create such an architecture.

Szyperski [29] defines components as ‘having explicit context dependencies only’, and being ‘subject to composition by third parties’. In Koala [20], the former are called (explicit) requires interfaces, while the latter is called third-party binding. Tony Williams recognized the same two elements in an early paper leading to open linking and embedding (OLE) [33], though an interface is called a ‘base class’ there, and the third party, a ‘creator’.

In the following subsections we provide examples of compositional approaches.

6.2. Example 1: The ‘Pipes and Filters’ Architectural Style

One of the oldest examples of the use of composition to create a diversity of applications with only few components is the pipes and filters architecture of Unix. With a minimum of common agreements (character based I/O, a standard input and output stream), and a new line as record separator, using a small set of tools such as *grep*, *sed*, *awk*, *find* and *sort* with essentially one simple composition operator, the

‘|’, one can program a large variety of tasks, including small databases, automatic mail handling and software build support.

Pipes and filters are considered an architectural style in [28], and indeed there are more examples with similar compositional value. Microsoft’s DirectShow [17] is a streaming architecture that allows the creation of ‘arbitrary’ graphs of nodes that process audio and video signals. The nodes, called *codecs*⁹, are implemented as COM servers and may be obtained from different vendors. The graph can be constructed manually with a program that acts as a COM client. A simple type mechanism allows the automatic creation of parts of such a graph, depending on the types of input and output signals. Other companies have similar architectures (e.g., Philips with the Trimedia Streaming Software Architecture (TSSA) [30]).

A third kind of pipes and filters architecture can be found in the domain of measurement and control. LabView [19] is an environment in which measurement, control and automation applications can be constructed as block diagrams and subsequently compiled into running systems. See also [5].

6.3. Example 2: Relation Partition Algebra

Software architectures are normally defined *a priori*, after which an implementation is created. But it is also possible to ‘reverse engineer’ an implementation to extract a software architecture from it, if the architecture has been lost. Even if the architecture has not been lost, it is still possible to extract the architecture from the implementation and verify it against the original architecture. Discrepancies mean that either the implementation is wrong, or the original architecture is ‘imperfect’ (to say the least).

Tools can be built to support this process, but a large variety of queries would result in a large range of tools. As it turns out, mathematical set and relation algebra can be used to model ‘use’ and ‘part-of’ relations in software effectively; queries can be expressed mostly in one-liners. A compositional tool set for this algebra is described in [23]; a similar technique is also deployed by [11].

6.4. Example 3: COM, OLE, and ActiveX

Microsoft’s COM [27] is a successful example of a component technology, but this does not automatically imply that it supports *composition* from the ground up. COM can be seen as an object-oriented programming technique, its three main functions, *CoCreateInstance*, *QueryInterface* and *AddRef/Release*, being, respectively, the new operator, the dynamic type cast, and the delete operator. COM’s main contribution is that it supports object orientation ‘in the large’: across language, process, and even processor boundaries.

If a COM client uses a COM service, the default design pattern is that it creates an instance of that service. It refers to the service by CLSID, a globally unique identifier that is nevertheless specific for that service, thus creating a fixed dependency between the client and the service. In other words, clients and services are (usually) not freely composable. To circumvent problems thus arising, some solutions have been found,

⁹ For *coding* and *decoding* streams of data.

most noteworthy the *CoTreatAs*, the *Category ID*, and the connectable interfaces. The latter provide a true solution, but in practice are only used for notifications.

Microsoft’s OLE [6] is an example of an architecture that allows the free composition of document elements (texts, pictures, tables) rather than software components. An OLE *server* provides such elements, while an OLE *container* embeds them. An application may be an OLE server and container at the same time, allowing for arbitrarily nested document elements. An impressive machinery of interfaces is needed to accomplish this, as a consequence, there are not many different OLE containers around (though there are quite a few servers).

Microsoft ActiveX controls are software components that are intended for free composition. In general, they do not have matching interfaces, making it necessary to glue them together. The first language for this was Visual Basic; later it became possible to use other languages as well. An impressive market of third-party ActiveX controls has emerged.

6.5. Example 4: GenVoca, Darwin, and Koala

GenVoca [2] is one of the few approaches that supports explicit *requires* interfaces to be bound by third parties. GenVoca uses a functional notation to combine components and is equipped with powerful code generation techniques. It has been applied to graph processing algorithms, where different combinations of the same set of components can be made to create a variety of applications.

Darwin [15] is a research language for the specification of distributed software architectures. It supports components with typed *provides* and *requires* interfaces, and it allows compound components to instantiate and bind other components. Darwin has been applied to the control and monitoring of coal mining, and - as Koala - to the control software for televisions (see below).

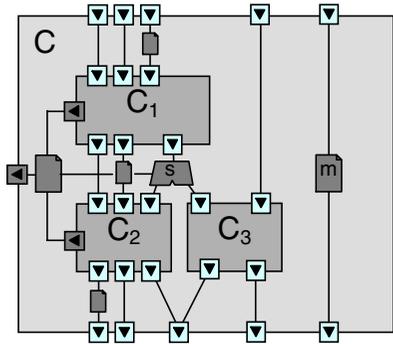


Fig. 10. The Koala component model

Koala [20] is an industrial architectural description language with a resource-friendly component model used for the creation of embedded control software for consumer products. Koala supports both composition and variation. Koala components can be instantiated and their (typed) interfaces can be bound by compound components as shown in Figure 10, thus forming a hierarchy of components. Variation is supported

in the form of diversity interfaces, a special form of requires interfaces that resemble properties in ActiveX controls; and switches, a kind of pseudo-dynamic binding. Various products are already on the market, created from different combinations of different subsets of a repository of reusable components. Koala supports adaptation in the form of glue modules as a fundamental technique for binding components.

6.6. Composability Is Actually a Form of Variation!

We have discussed variation and composition. We presented them as fundamentally different, but it is possible to unify the concepts. In all the examples that we have seen, components become ‘freely composable’ when context dependencies have been turned into explicit variation points. This may be in the form of an input stream or of connectable (requires) interfaces. Some approaches make every context dependency explicit (e.g., Koala), while others only make certain dependencies explicit (like in the pipes and filter architectures).

7 Concluding Remarks

Let us first recapitulate. Managing complexity and diversity while maintaining high quality and short lead times have become important software development issues. Reuse, architecture, and components are key solutions to achieve this. Originally, we envisaged reuse to happen ‘automatically’ once software components were available on the market. Later we discovered that sharing software actually requires a proactive, systematic approach: that of software product lines. The first product lines rely heavily on (limited) *variation* in an otherwise variant-free architecture. With product scope increasing, and software development starting to cross organizational borders, we believe that more powerful techniques are required to handle diversity and discuss *composition* as one such a technique.

This does not mean that we’re back to third-party component markets. Instead, a careful balance between architecture-driven top-down and component-driven bottom-up development is required, and we explain why and how. Also, composition does not replace variation, but rather they should be seen as two independent dimensions. We typify various points in the two-dimensional space with archetype examples, before explaining variation and composition individually in more detail. To make things as concrete as possible, we illustrate variation and composition with a rich set of well known examples.

Within Philips, we have moved from variation to composition in our architectures for certain sets of consumer products. We find that we can handle diversity better, and cover a larger range of products even though development is decoupled in time and space across organizations. We find that few other companies use composition, which originally surprised us. We then discovered that the few companies that do use composition also cover a wide product scope, cross organizational boundaries, or exhibit both phenomena.

Acknowledgements

We would like to thank attendants of the Software Architectures for Embedded Systems (ARES) workshops, the Software Product Line Conference, and the Philips ‘Composable Architectures’ project for many fruitful discussions on this topic.

8 References

1. Felix Bachmann, Len Bass, *Introduction to the Architecture Based Design Method*, Tutorial at the First Software Product Line Conference (SPLC1), August 28-31, 2000, Denver, Colorado, USA.
2. Don Batory, Sean O'Malley, *The Design and Implementation of Hierarchical Software Systems with Reusable Components*, ACM Transactions on Software Engineering and Methodology, 1 no. 4, pp. 355-398 (October 1992).
3. Joe Bauman, *The Perfect Architecture is non-optimal, Winning with Chaos*, Proceedings of the 4th international workshop on Product Family Engineering, Bilbao, Spain, October 2001.
4. Jan Bosch, *Organizing for Software Product Lines*, Proceedings of the 3rd international workshop on the development and evolution of software architectures of product families, Las Palmas, March 2000.
5. Klaas Brink, *Interfacing Control and Software Engineering: a formal approach*, PhD thesis, Technical University, Delft, The Netherlands, June 24, 1997.
6. Kraig Brockschmidt, *Inside OLE Second Edition*, 1995, Microsoft Press, ISBN 1-55615-843-2.
7. Frederick P. Brooks Jr, *The Mythical Man-Month, Essays on Software Engineering*, Addison-Wesley Publishing Company, ISBN 0-201-00650-2.
8. Paul Clements, Linda Northrop, *Software Product Lines, Practices and Patterns*, Addison-Wesley, 2002, ISBN 0-201-70332-7.
9. Krzysztof Czarnecki, Ulrich Eisenecker, *Generative Programming: Methods, Tools, and Applications*, Addison-Wesley Pub Co; ISBN: 0-201-30977-7.
10. David Garlan, Robert Allen, John Ockerbloom, *Architectural Mismatch, or: Why it's hard to build systems out of existing parts*, ICSE 95, Seattle, Washington USA.
11. Richard C. Holt, *Structural Manipulations of Software Architecture using Tarski Relational Algebra*, Proceedings of fifth Working Conference of Reverse Engineering, WCRE'98, IEEE Computer Society, 1998.
12. *IEEE Recommended Practice for Architectural Description of Software Incentive Systems*, <http://standards.ieee.org/catalog/olis/se.html>
13. Ivar Jacobson, Martin Griss, Patrick Jonsson, *Software Reuse – Architecture, Process and Organization for Business Success*, Addison Wesley, New York, 1997, ISBN 0-201-92476-5.
14. Frank van der Linden, Jan Gerben Wijnstra, *Platform Engineering for the Medical Domain*, Proceedings of the 4th international workshop on Product Family Engineering, Bilbao, Spain, October 2001.
15. Jeff Magee, Naranker Dulay, Susan Eisenbach, Jeff Kramer, *Specifying Distributed Software Architectures*, Proc. ESEC'95, Wilhelm Schafer, Pere Botella (Eds.) Springer LNCS 989 pp. 137-153 (1995).
16. M. D. McIlroy, *Mass produced software components*, Proc. Nato Software Eng. Conf., Garmisch, Germany (1968) 138-155.
17. Microsoft *DirectShow*, <http://www.gdcl.co.uk/dshow.htm>, part of DirectX.
18. Microsoft, *Shell Programmers Guide*, <http://msdn.microsoft.com/library/>
19. National Instruments, *LabView*, <http://www.natinst.com/labview/>

20. Rob van Ommering, Frank van der Linden, Jeff Kramer, Jeff Magee, The Koala Component Model for Consumer Electronics Software, IEEE Computer, March 2000, p78-85.
21. Rob van Ommering, *Beyond Product Families: Building a Product Population?*, Proceedings of the 3rd international workshop on the development and evolution of software architectures of product families, Las Palmas, March 2000.
22. Rob van Ommering, *Roadmapping a Product Population Architecture*, 4th International Workshop on Product Family Engineering, Bilbao, Spain, October 3-5, 2001.
23. Rob van Ommering, René Krikhaar, Loe Feijs, *Language for Formalizing, Visualizing and Verifying Software Architectures*, Computer Languages 27 (2001) p3-18.
24. David L. Parnas, *On the Criteria to Be Used in Decomposing Systems into Modules*, Communications of the ACM, Vol. 15, No. 12, December 1972, p1053-1058.
25. Dewayne E. Perry, *Generic Architecture Descriptions for Product Lines*, Proceedings of the Second International ESPRIT ARES Workshop, LNCS 1429, Springer Verlag, Berlin Heidelberg, 1998, p51-56.
26. Jeff Prosise, *Programming Windows 95 with MFC*, Microsoft Press, 1996, ISBN 1-55615-902-1.
27. Dale Rogerson, *Inside COM, Microsoft's Component Object Model*, Microsoft Press, ISBN 1-57231-349-8, 1997.
28. Mary Shaw, David Garlan, *Software Architecture, Perspectives on an Emerging Discipline*, Prentice Hall, 1996, ISBN 0-13-182957-2.
29. Clemens A. Szyperski, *Component Software: Beyond Object-Oriented Programming*. Addison-Wesley, Harlow, UK, 1998, ISBN 0-201-17888-5.
30. The Trimedia Streaming Software Architecture, <http://www.trimedia.com/products/briefs/tssa.html>
31. Kurt Wallnau, Scott Hissam, Robert Seacord, *Building Systems from Commercial Components*, Addison-Wesley Pub Co; ISBN: 0-201-70064-6
32. Jan Gerben Wijnstra, *Supporting Diversity with Component Frameworks as Architectural Elements*, Proceedings of the 22nd International Conference on Software Engineering, Limerick, June 4-11, 2000, p. 51-60.
33. Tony Williams, *On Inheritance, What It Means and How To Use It*, Microsoft Internal Report, 1990, <http://research.microsoft.com/comapps/docs/Inherit.doc>