

Package Surface Blueprints: Visually Supporting the Understanding of Package Relationships

Accepted at ICSM'2007: International Conference on Software Maintenance

Stéphane Ducasse* Damien Pollet Mathieu Suen Hani Abdeen Ilham Alloui
Language and Software Evolution Group — Université de Savoie, France

Abstract

Large object-oriented applications are structured over large number of packages. Packages are important but complex structural entities that may be difficult to understand since they play different development roles (i.e., class containers, code ownership basic structure, architectural elements...). Maintainers of large applications face the problem of understanding how packages are structured in general and how they relate to each others. In this paper, we present a compact visualization, named Package Surface Blueprint, that qualifies the relationships that a package has with its neighbours. A Package Surface Blueprint represents packages around the notion of package surfaces: groups of relationships according to the packages they refer to. We present two specific views one stressing the references made by a package and another showing the inheritance structure of a package. We applied the visualization on two large case studies: ArgoUML and Squeak.

This paper makes heavy use of colors in the figures. Please obtain and read an online (colored) version of this paper to better understand the ideas presented in this paper.

1 Introduction

To cope with the complexity of large software systems, applications are structured in subsystems or packages. It is now frequent to have large object-oriented applications structured over large number of packages. Ideally, packages should keep as less coupling and as much cohesion as possible [25, 5], but as systems inevitably become more complex, their modular structure must be maintained. It is thus useful to understand the concrete organization of packages and their relationships. Packages are important but complex structural entities that can be difficult to understand since they play

different development roles (i.e., class containers, code ownership basic structure, architectural elements...). Packages provide or require services. They can play core roles or contain accessory code features. Maintainers of large applications face the problem of understanding how packages are structured in general and how packages are in relation with each others in their provider/consumer roles. This problem was experienced first-hand by the first author while preparing the 3.9 release of Squeak, a large open-source Smalltalk [8]. In addition, approaches that support application remodularization [1, 20, 22] succeed in producing alternative views for system refactorings, but proposed changes remain difficult to understand and assess. There is a good support for the algorithmic parts but little support to understand their results. Hence it is difficult to assess the multiple solutions.

Several previous works provide information on packages and their relationships, by visualizing software artifacts, metrics, their structure or their evolution [6, 7, 10, 19, 23, 28]. However, while these approaches are valuable, they fall short of providing a fine-grained view of packages that would help understanding the package shapes (the number of classes it defines, the inheritance relationships of the internal classes, how the internal class inherit from external ones...) and help identifying their roles within an application.

In this paper, we propose Package Surface Blueprint, a compact visualization revealing package structure and relationships. A package blueprint is structured around the concept of surface, which represents the relationships between the observed package and its provider packages. The Package Surface Blueprint reveals the overall size and internal complexity of a package, as well as its relation with other packages, by showing the distribution of references to classes within and outside the observed package. We applied the Package Surface Blueprint to several large case studies namely Squeak the open-source Smalltalk comprising more than 2000 classes, ArgoUML and Azureus.

Sections 2 & 3 present the challenges in supporting package understanding, and summarize the properties wanted for effective visualizations. Section 4 presents the structuring principles of a package blueprint, which are then declined

*We gratefully acknowledge the financial support of the french ANR (National Research Agency) for the project "COOK: Réarchitecturisation des applications industrielles objets" (JC05 42872).

to support a reference view and an inheritance view in Section 5. Section 6 then describes some recurring patterns. In sections 7 & 8, we discuss our visualization and position it w.r.t. related work before concluding.

2 Challenges in Understanding Packages

Although languages such as Java offer a language mechanism for modelling the dependencies between packages (*i.e.*, via the import statement), this mechanism does not really support all the information that is important to understand a package. We present a coarse list of useful information to understand packages. Our goal here is to identify the challenges that maintainers are facing and not to define a list of all the problems that a particular solution should tackle.

Size. What is the general size of a package in terms of classes, inheritance definition, internal and external class references, imports, exports to other packages? For example, do we have only a few classes communicating with the rest of the system?

Cohesion and coupling. Transforming an application will follow natural boundaries defined by coupling and cohesion [5, 2]. Assessing these properties is then important.

Central vs. Peripheral. Two correlated pieces of information are important: (1) whether a package belongs to the core of an application or if it is more peripheral, and (2) whether a package provides or uses functionality.

Developers vs. Team. Knowing who are the developers and maintainers of the application and packages helps in understanding the architecture of the application and in qualifying package roles [13, 24]. Approaches such as the distribution map may help in this task [9].

In addition, packages reflect several organizations: they are units of code deployment, units of code ownership, can encode team structure, architecture and stratification. Good packages should be self-contained, or only have a few clear dependencies to other packages [5, 2, 18]. A package can interact with other ones in several ways: either as a provider, or as a consumer or both. In addition a package may have either a lot of references to other packages or only a couple of them. If it defines subclasses, those can form either a flat or deep subclass hierarchy. It can contain subpackages.

Figure 1 shows situations where the same group of classes can be dispatched. Note that for the purpose of illustration, Figure 1 only shows references but the same idea holds for inheritance between classes distributed in several packages. In both cases (a) and (b), there are only two packages but in case (a) most of the classes of P4 inherit directly from a class in P1 while in case (b) all the classes of P4 inherit

internally from B2 which is a root of an inheritance hierarchy. Revealing this difference is important since we want to understand if we can change the relationships between P1 and P4 during a refactoring process. In cases (a) and (c), we have exactly the same relationships between classes but the package structure is different. As mentioned by R. Martin importing a class equals importing the complete package [21], therefore importing two classes from the same package is quite different from importing them from two different packages since in the latter case we import all the classes of the two packages.

Note that understanding packages is also important in the context of remodularization approaches [1, 20, 22]. There it is important to understand how the proposed remodularisation compares with the existing code. This problem is particularly stressed in presence of legacy applications that consist of thousands of classes and hundreds of packages.

3 Visualization Challenges

We researched the characteristics that an efficient visualization should hold [3, 30, 32]. As our focus is on providing a first impression of a package and its context, we want to exploit the gestalt principles of visualization and preattentive processing¹ as much as possible to help spotting important information [29, 14, 15, 32].

To support the understanding of packages, we want the visualization to highlight the characteristics of a package in terms of its internal size, internal and external references. In particular we want to spot classes or dependencies that stand out in a given package. We stress that our visualization should take into account the following properties:

Good mapping to reality. The visualization should offer a good representation of the situation that the maintainer can trust and from which it can draw and validate hypothesis.

We want the visualization to highlight the general tendency of a package in terms of its internal size, internal and external references. In particular we want to spot classes or dependencies that stand out in a given package.

Scalability and simple navigation. The maintainer should easily access the information. The visualization should

¹ Researchers in psychology and vision have discovered a number of visual properties that are preattentively processed. They are detected immediately by the visual system: viewers do not have to focus their attention on a specific region in an image to determine whether elements with the given property are present or absent. An example of a preattentive task is detecting a filled circle in a group of empty circles. Commonly used preattentive features include hue, curvature, size, intensity, orientation, length, motion, and depth of field. However, combining them can destroy their preattentive power (in a context of filled squares and empty circles, a filled circle is usually not detected preattentively). Some of the features are not adapted to our needs. For example, we do not consider motion as applicable.

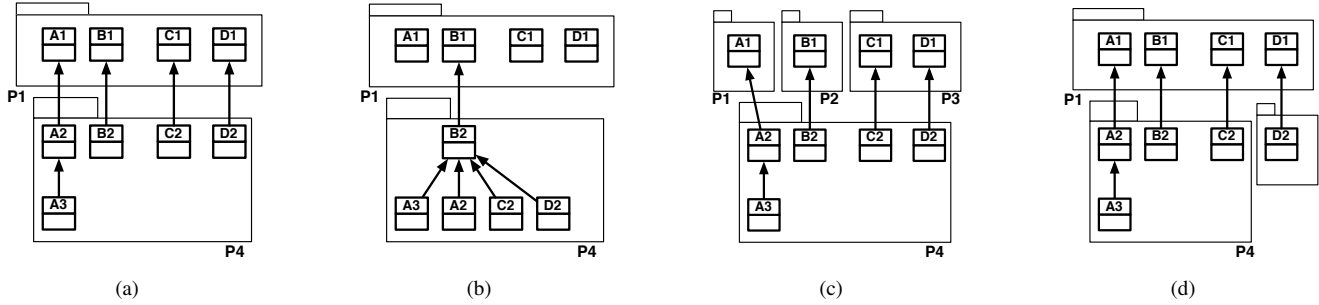


Figure 1. Different package configurations over the same number of classes.

scale *i.e.*, we should be able to have system overview as well as focusing on a particular package. We want a visualization that scales well with the number of packages and of dependencies, so we prefer to avoid depicting dependencies with graphs. Given that the graph will contain more than thousands of nodes and much more edges, this will result a unusable view [16].

Low visual complexity. By being regular and well structured, *i.e.*, reusing the same conventions of color or position, the visualization should help the maintainer to learn it and understand it. In addition, while the visualization should offer a lot of information, it should not be complex to analyze.

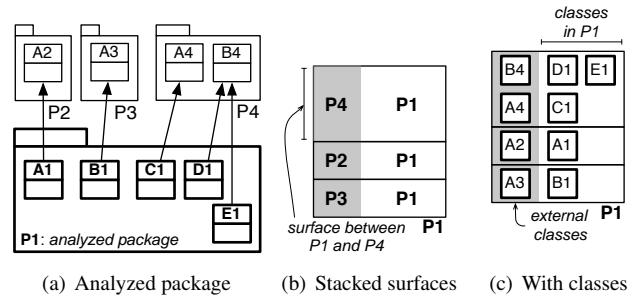
4 Package Surface Blueprints

A package blueprint represents how the package under analysis references other packages. Figure 2 presents the key principles of a Package Blueprint. These principles will be realized slightly differently when showing direct class references or inheritance relationships.

4.1 Basic Principles

The package blueprint visualization is structured around the “contact areas” between packages, that we name *surfaces*. A *surface* represents the conceptual interaction between the observed package and another package. In Figure 2(a) the package P1 is in relation with three packages P1, P2, and P4, via different relationships between its own classes and the classes present in the other packages, so it has three surfaces.

A package blueprint shows the observed package as a rectangle which is vertically subdivided by each of the package’s surfaces. Each subdivision represents a surface between the observed package and a referenced package, and will be more or less tall, depending on the strength of the relation between the two packages. In Figure 2(b), the package blueprint of P1 is made from three stacked boxes because



(a) Analyzed package (b) Stacked surfaces (c) With classes

Figure 2. Consider P1 that references four classes in three other packages (a). A blueprint shows the surfaces of the observed package as stacked subdivisions (b). Small boxes represent classes, either in the observed package (right white part) or in referenced packages (left gray part) (c).

P1 references package three other packages. The box of the surface between P1 and P4 is taller because P1 references more classes in P4 than in P2 or P3.

In each subdivision, we show the classes involved in the corresponding surface. By convention, we *always* show the classes in the referenced packages in the leftmost gray-colored column of each surface, and the classes of the observed package on the right. In Figure 2(c), the topmost surface shows that classes D1 and E1 reference class B4, and that C1 references A4. If many classes reference the same external class, we show them all in an horizontal row; we can thus assess the importance of an external class by looking at how many classes there is in the row: in Figure 2, the row of B4 stands out because the two referring classes D1 and E1 make it wider.




Figure 3. Surface package blueprint detailed view.

4.2 Detailed Explanation

To convey more information, we add variations to the basic layout described above, as illustrated in Figure 3.

Internal References. To support the understanding of references between classes inside the observed package, we add a particular surface with a thick border at the top of the blueprint. We name this surface the head of the blueprint, and the rest its body. In the head, the first column represents the internal classes that are referenced from within the package itself: here A1 and G1 are the classes referenced respectively by B1 and C1 and H1 and I1. The height of the head surface indicates the number of classes referenced within the package.

Position. Internal classes are arranged by columns: each column (after the leftmost one) refers to the same internal class for all the surfaces. The width of the surface indicates the number of referencing classes of the package. Figure 3 shows that class C1 internally references A1, and externally references A3 and B3.

We order classes in both horizontal and vertical direction to present important elements according to the (occidental) reading direction. Horizontally, we sort classes from left to

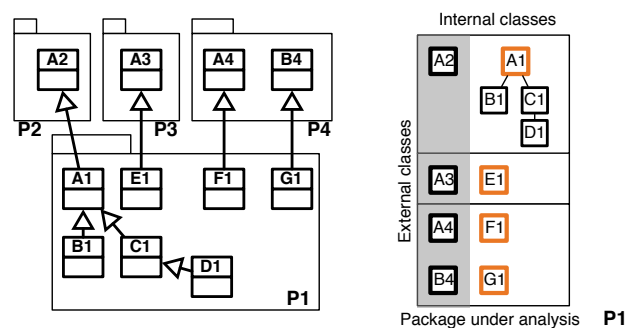


Figure 4. Inheritance package surface blueprint. Orange bordered classes inherit from external classes directly.

right according to the number of *external* classes they reference from the whole package. Hence classes referencing the most occupy the nearest columns from the gray area.

We apply the same principle for the vertical ordering, both of surfaces within a blueprint, and of rows (*i.e.*, external classes) within a surface. Within a package, we position surfaces that reference the most classes the highest. Within a surface, we order external classes from the most referenced at the top, to the least referenced are at the bottom of the surface. This is why in Figure 3 the surface with P3 is the highest and why the surface with P2 is above P4, since there are more classes references from P2 than from P4.

Color. We want to distinguish referenced classes depending on whether they belong to a framework or the base system, or are within the scope of the application under study. When a referenced class is not part of the application we are currently analyzing, we color its border in cyan. In addition the color intensity of a node conveys the number of references it is doing: the darker the more references. Both intensity and horizontal position represent the number of references, but position is computed relative to the whole package, while intensity is relative to each surface. Thus, while classes on the left of surfaces will generally tend to be dark, a class that makes many references in the whole package but few in a particular surface will stand up in this surface since it will be light grey.

The Case of Inheritance. Up to now, we only discussed references, but inheritance is a really important structural relationship in object-oriented programming. We adapt the Package Surface Blueprint to offer a view specific to inheritance, as shown in Figure 4. In this variation, we consider only single inheritance so we don't need the head surface: we can display all classes and subclasses transitively inheriting from external classes on the same row. We distinguish the direct subclasses of external classes by showing them with

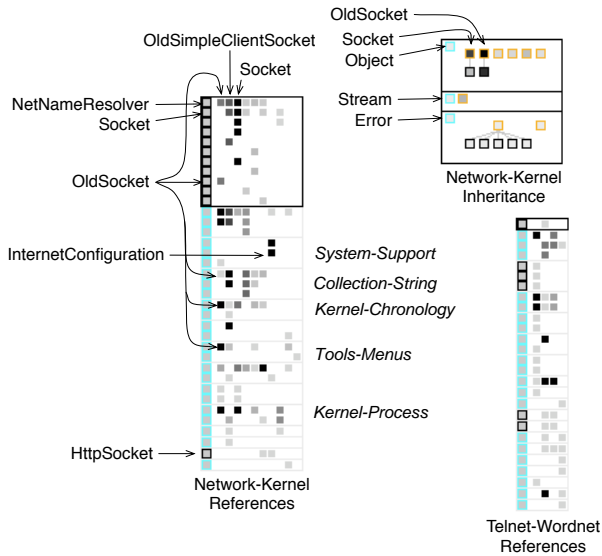


Figure 5. Analysing the Network-Kernel Package.

an orange border; indirect subclasses are black-bordered and arranged in trees under their superclass. In addition, root classes such as Object are filled in cyan and abstract classes in blue. In Figure 4, A1 inherits from A2 defined in package P2, while B1, C1, and D1 inherit from A1.

The fill color of classes in the inheritance view still represents the number of *references*, but relative to the *package* and not to the surface like in the references views. This makes it possible to correlate inheritance and references. For instance, the top-right view in Figure 5 shows that most references come from a subclass (Socket) of Object; in other cases, references might come from classes that are lower in the hierarchy as HTMLInput in Figure 6.

4.3 An Example: The Network Subsystem

We are now ready to have a deeper look at an example. The Squeak Network subsystem contains 178 classes and 26 packages — this package contains on the one hand a library and a set of applications such as a complete mail reader. The blueprint on the left in Figure 5 shows the references package blueprint of the Network-Kernel package in Squeak.

Glancing at it we see that the package blueprint of the Network-Kernel package has nearly a square top-red surface indicating that most internal classes are referenced internally. This conveys a first impression of the package’s cohesion even if not really precise [5]. Contrast it with the package blueprint of the Telnet-Wordnet package which clearly shows little internal references.

We see that Network-Kernel is in relation with thirteen other packages. Most of the referenced classes are cyan, which means that they are not part of the network subsystem. What is striking is that all except one of the referenced classes are classes outside the application (see (HTTPSocket) in Figure 5). However, since the package is named *kernel*, it is strange that it refers to other classes from the same application, and especially only one. We see that half of the referred packages have strong references (indicated by their dark color).

Using the mouse and pointing at the box shows using a fly-by-help the class and package names (indicated in italics in Figure 5). The Tools-Menus surface indicates some improper layering. Indeed it shows that Network-Kernel is referencing UI classes via the package Tools-Menus which seems inappropriate. We learn that the class making the most internal references is named OldSocket; this same class also makes the most external references, to three packages (Collection-String, Tools-Menus, and Kernel-Chronology). The second most referencing class is named OldSimpleClientSocket. It is worth to notice that OldSocket is only referencing itself and that even OldSimpleClientSocket does not refer to it, so it could be removed from this package without problems. The third most referencing class is Socket. Having two classes named Socket and OldSocket clearly indicates that the package is in a transition phase where a new implementation has been supplanting an old one. We learn that the most internally referenced class is NetNameResolver and the second most is Socket. So this is a sign of good design since important domain classes are well used within the package.

The inheritance package blueprint shows that the Network-Kernel package is bound to three external packages containing the three superclasses Object, Error, and Stream. In addition the package, while inheriting a lot from external packages, is inheriting from the same class, here Object. The difference between the two main surfaces is interesting to discuss: the topmost surface shows that most of the classes are directly inheriting from one external superclass (here Object), while the second one shows that errors are specialized internally to the package. All in all, this makes sense and provides a good characterization of the package.

5 Packages Within Their Application

Understanding a package in isolation (mainly as a consumer) is interesting but lacks information about the overall context *i.e.*, is a selected class used by other packages? which packages is a selected surface about? As shown in the following subsections, our approach also supports the understanding of the situation of a class/package within the context of a complete application.

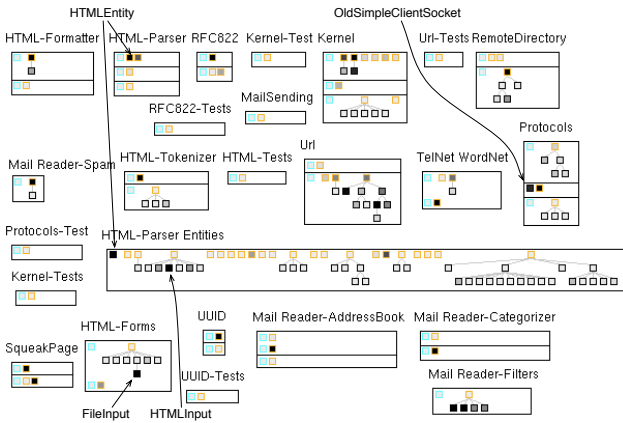


Figure 6. Inheritance global view in Network

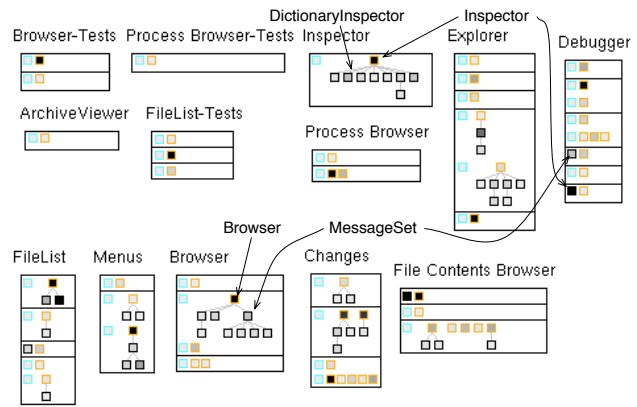


Figure 7. Inheritance global view in Tools

5.1 Inheritance package blueprint Overview

Overviewing all the package blueprints of an application gives a first impression of how the packages were built and structured. During our case studies, we identified a few remarkable usage patterns: a package can mainly contain big inheritance hierarchies (potentially a single one); classes in a package may inherit from superclasses within the application itself or from frameworks or the base system; or a package can specialize functionality and have few internal inheritance relationships.

First Case: Squeak’s Network. For example, Figure 6 shows all the package blueprints of the Network subsystem in Squeak, which groups library and application classes. It shows that there are only two places where classes inherit from classes within the Network subsystem scope: HTML-Entity and OldSimpleClientSocket. Note however that OldSimpleClientSocket has a lighter shade of gray than HTML-Entity; this indicates that the former is not referencing other classes as much as the latter.

Clicking on the HTML-Entity box, we can see that it is defined in the Network-HTML-Parser package, away of all its subclasses, and then directly consider that it is defined in the wrong package. We can immediately spot that some packages are heavily structured around inheritance, like the package Network-HTML-Parser Entities or Network-Mail Reader-Filters which define a single hierarchy.

The overview also shows classes doing a lot of references (indicated as black boxes) such as HTML-Entity, FileInput and HTMLInput. However, in the context of inheritance, we should pay attention to the fact that all the subclasses of a class inherit its behavior and references. While we can spot classes doing a lot of references, the view does not convey the tree ordering so it is difficult to evaluate the subclasses of a given class. The case of FileInput is interesting: while

it is a leaf in the inheritance tree, it makes a lot of direct references, indicating that the class is complex.

While the views are simple, they convey powerful information. If we analyze a bit, we can see that the percentage of black-bordered boxes reveals the amount of internal reuse. Orange-bordered classes that inherit from a cyan class indicate reuse of functionality from outside the application. Note that this is different from many orange-bordered classes inheriting from a black-bordered one (like with HTML-Entity in HTML-Parser Entities), since a lot of classes inherit from Object and indeed do not share the same domain. In contrast, inheriting from HTML-Entity clearly reuses its domain.

Second Case: Squeak’s Tools. Figure 7 shows the blueprints of the Tools packages which contain all the Squeak development tools: code browsers, debuggers... Without going into details, we immediately see different shapes. Here, the blueprints are thinner but often higher, showing that there is less internal reuse than in Network. Note that even if the Tools packages contain a large set of development tools, inheritance is actually used to reuse abstractions: The blueprint of Tools-Browser shows that the class Browser, even if it defines a tool, is inherited several times. Other tools reuse the abstraction of Browser: for instance, its subclass Message-Set allows one to browse a group of methods and is reused and extended in Tools-Debugger.

The blueprint of Tools-Debugger shows an interesting shape: it is narrow and has a nearly flat inheritance hierarchy. Moreover, all its classes are inheriting from classes outside the package. Note that this behavior makes sense because the package aggregates functionality defined elsewhere, and the view easily reveals it. The package Tools-FileList defines a tool to browse external files and shows a similar shape.

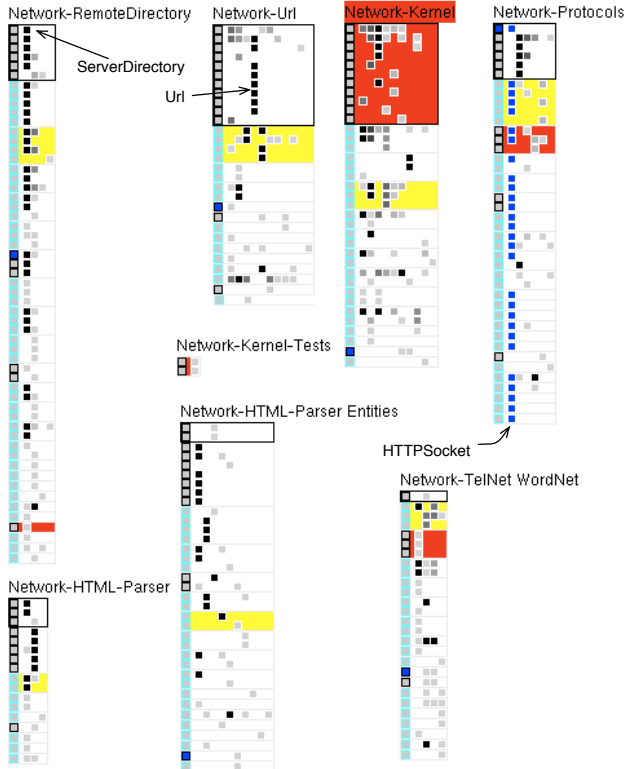


Figure 8. In this view, the `Network-Kernel` package was selected in red, surfaces with `Collections-Strings` annotated in yellow, and class `HTTPSocket` selected in blue.

5.2 Interactively Querying the Blueprint

The maintainer can also query the system by clicking either on a class or on a surface. This highlights in red all occurrences of the class, or all surfaces referring to the same package. In addition, colors can be assigned to a surface to help the maintainer identify all the surfaces communicating with the same packages.

Figure 8 shows the blueprints of all the `Network` packages referencing and defining `HTTPSocket`. It is striking to see that `HTTPSocket` is a central class of the package `Network-Protocols` as it refers to most of the classes referred by that package. In addition, the surface referencing the package `Collections-Strings` is annotated in yellow and we see how all the packages refer to this package.

By clicking on the head surface, it gets colored in red and shows the package usage by coloring the surfaces referencing it in red. Figure 8 shows how the package `Network-Kernel` is used within the application.




Figure 9. A Sumo Blueprint: the `Critics` package in ArgoUML.

6 Striking Shapes

While applying blueprints to large applications we identified some striking shapes that the blueprint, a surface or a class within a blueprint would produce. We present here the most frequent ones.

6.1 Shapes of Packages and Surfaces

Sumo Package. A very large and tall reference blueprint denotes a package that makes a lot of references from many classes. Figure 9 shows an example: the package `Critics` of ArgoUML that defines all the rules for assessing the quality of models.

Small House Package. A small inheritance blueprint with only a couple of surfaces and few inheritance hierarchies often denotes a package that offers a well packaged functionality, like `Tools-Debugger` or `Tools-FileList` (Figure 7). These blueprints are usually taller than larger.

Flat Head Package. A reference blueprint with a wide but flat head indicates limited internal references. `Network-TelNet WordNet` and `Network-HTML-Parser Entities` in Figure 8 are flat head blueprints.

Exclusive External Referencer Package. When the first column in a blueprint is almost or completely cyan, the package makes most or all of its external references to classes outside the scope of the analyzed application. These pack-

ages typically extend a framework or a core library; Network-Kernel in Figure 8 is an example.

Loner Package. A loner is a package that contains only a couple of classes. It is often containing a single test case class. The blueprint Network-Kernel-Tests in Figure 8 or Network-Mail Reader-Categorizer, Network-UUID, Network-Mail Reader-Spam of Figure 6 are loners. Some of these packages are clearly good candidates for modularisation.

Large External Surface. When the topmost external surfaces are really large, like the four surfaces below the head in Figure 9, they identify packages that we must pay attention to, because changes in these external packages will very probably impact the package under analysis.

Square Head Package. A package that references all its own classes will have a blueprint with a square internal surface; this denotes a package that is quite cohesive. In Figure 8, Network-Kernel has a square head and appears to be relatively well packaged.

Tower Package. A reference blueprint with a small head and a thin body denotes a package with few internal references but that makes many external references. This package may not be cohesive but highly coupled with the external packages. The package peer in Azureus is an extreme of this shape, as shown in Figure 10. In Figure 8, Network-RemoteDirectory has a more cohesive head and three classes intensively referencing external packages.

6.2 Shapes of Classes

Main Referencer Class. A vertical alignment of dark squares in the body of a blueprint denotes a class that is responsible for many references to classes in other packages. The classes HTTPSocket and ServerDirectory are the main referencers in packages Network-Protocols and Network-RemoteDirectory; they are candidates to be central package classes (Figure 8).




Figure 10. Peer in Azureus: a Tower Blueprint

Main Internal Referencer Class. When vertical alignments are limited to the head, they reveal classes doing many internal and few external references. These classes often define the abstraction of the application. In Figure 8, the class Url only references classes within Network-Url.

Omnipresent Referenced Class. Classes of this kind are referenced by almost all the internal classes, and easily identifiable by filled rows in a surface. This makes sense for a facade class if it occurs a few times, but in ArgoUML we see this shape in most packages for Facade and Model (see Figure 9); we may thus assess that the Facade pattern is misused.

7 Evaluation and Discussion

7.1 Evaluation

The Package Surface Blueprint shows the internal number of classes as well as the number of classes externally referenced. Hence it conveys whether the package is using a lot of information or not.

Size. The Package Surface Blueprint shows the complexity of the observed package in several dimensions. The height of the body indicates the amount of external classes referenced, whereas the number of surfaces shows the number of referenced packages. The height of each individual surface shows how many classes are referenced in the corresponding package. This gives us an estimate of the coupling between the package and this surface; to further evaluate the coupling strength, we should also look at the intensity of referencing classes in the surface because it represents the number of references. In addition, the width of the surface indicates the number of referencing classes.

Those visual properties combine to give a quick impression not just about the visualized package, but also about its classes: a thin package with a long body depends on a lot of classes because of few internal classes. If moreover the blueprint is heavily lined, *i.e.*, it references a lot of packages, so some of its referencing classes may be complex and fragile.

Central or Peripheral. By looking at the border color of external classes (cyan or black), we can easily see if a package depends a lot on the framework or on the application. Also, by using the selection mechanism, we can interactively see if a package is imported by different subsystems (central) or just by specific ones (peripheral).

Cohesion and Coupling. The package blueprint also makes it possible to roughly compare how several packages are coupled with the observed one: larger surfaces indicate coupling to more classes and are positioned nearer to the head surface, while surfaces with more darker class squares

represent packages which are more coupled in term of sheer number of references. We can also estimate cohesion by comparing internal coupling (size and overall intensity of the head surface) and external coupling.

Co-changes and Impact Analysis. Because the package blueprint details how packages depend on each other, it hints at the fragility of the observed package to changes. Selecting a package or a class highlights surfaces or classes that reference the selected entity and are thus sensitive to its changes.

7.2 Discussion

Our approach has worked well on our case studies. It should be noted that we were *not* familiar with the case studies before applying our approach. We have been able to locate many conceptual bugs; for instance we found some clearly unwanted dependencies, like the package Network-Telnet WordNet referencing a class in the user interface framework. However one of our future works is to evaluate the view with users. The Package Surface Blueprint answers the main challenges proposed in Section 2 and in Section 3; we further intend to address some remaining challenges.

Position Choices. We grouped the internal references at the top of the package blueprint, then ordered the surfaces from the ones having the most external references at the top to the least at the bottom; inside a surface, we also ordered the rows from the most referencing ones to the least. This way, we do not force the reader to scroll through big visualizations, and use the fact that the reader pays more attention to the top elements than to the bottom ones. We also tried to layout surfaces compactly so that we can easily move them.

Seriation. Rows within a surface are sorted according to the number of references they contain. In an earlier version we applied the dendrogram seriation algorithm [17] to group lines having similar referencing classes. However the resulting views were not as meaningful as with a simple ordering. We plan to use seriation to group packages having similar surfaces *i.e.*, packages using similar packages.

Properties. Instead of the number of references, we could map different properties to the color of classes and surfaces. This can create new striking shapes, adapted to a specific maintenance problem.

Impact of Boundaries. We color classes that do not belong to the application in cyan; this is a bit limiting since we do not distinguish well the true root classes —*e.g.*, Object or Model in Squeak— from the classes of a domain library that the analyzed application would extend. We found it really effective to color surfaces so that the user can interactively

mark entities on which he wants to focus; this increases the usability of the tool and speeds up understanding packages.

Shapes. For the time being we represent the classes with squares only. We could convey more information by using several visually distinct shapes. But it is not clear which ones and how efficient the results will be.

Package Nesting. Currently we do not support the nesting of packages. A solution like the one proposed by Lungu *et al.* seems complementary to our approach and interesting to deal with package nesting [19]. We also consider two types of relationships between packages (direct reference and inheritance); therefore we can extend our approach to other types of relationships like method invocation.

Other Views. So far we only presented blueprints to understand how a package was referencing or inheriting from other packages and classes. However we developed the reverse view: blueprints that present incoming references made by external classes on the observed package. Due to space limitation we did not present it. This information is useful when supporting package splitting or merging.

8 Related Works

Several works provide or visualize information on packages. Many of these approaches treat software co-change, looking at coupling from a temporal perspective, whereas in this paper we focus on the static structure of relationships [4, 11, 12, 27, 31, 33].

Lungu *et al.* guide exploration of nested packages based on patterns in the package nesting and in the dependencies between packages [19]; their work is integrated in Software-naut and adapted to system discovery.

Sangal *et al.* adapt the dependency structure matrix from the domain of process management to analyze architectural dependencies in software [26]; while the dependency structure matrix looks like the package blueprint, it has no visual semantics. Storey *et al.* offer multiple top-down views of an application, but these views do not scale very well with the number of relationships [28].

Ducasse *et al.* present Butterfly, a radar-based visualization that summarizes incoming and outgoing relationships for a package [10], but only gives a high-level client/provider trend. In a similar approach, Pzinger *et al.* use Kiviat diagrams to present the evolution of package metrics [23]. Chuah and Eick use rich glyphs to characterize software artefacts and their evolution (number of bugs, number of deleted lines, kind of language...) [6]. In particular, the timewheel exploits preattentive processing, and the infobug presents many different data sources in a compact way. D’Ambros *et al.* propose an evolution radar to understand the package coupling based on their evolution [7]. The radar view is effective at identifying outliers but does not detail structure.

Those approaches, while valuable, fall short of providing a fine-grained view of packages that would help understanding the package shapes (the number of classes it defines, the inheritance relationships of the internal classes, how the internal classes inherit from external ones,...) and support the identification of their roles within an application.

9 Conclusion

In this paper, we tackled the problem of understanding the details of package relationships. We described the Package Surface Blueprint, a visual approach for understanding package relationships. While designing Package Surface Blueprint, we tried to exploit gestalt visualization principles and preattentive processing.

We successfully applied the visualization to several large applications and we have been able to point out badly designed packages. To help users interpret views, we have identified a list of recurrent striking blueprint shapes. We also introduced interactivity to help the user focus and navigate within the system. We were however rather knowledgeable about both the visualization and the studied systems; in future work, we will validate the package blueprint usability by conducting tests with several independent software maintainers.

References

- [1] Anquetil and Lethbridge. Experiments with clustering as a software modularization method. In *WCRE*, 1999.
- [2] Arisholm, Briand, and Foyen. Dynamic coupling measurement for object-oriented software. *IEEE TSE*, 30(8), 2004.
- [3] Bertin. *Semiology of Graphics*. 1983.
- [4] Beyer. Co-change visualization. In *ICSM*, 2005.
- [5] Briand, Daly, and Wüst. A unified framework for coupling measurement in oo systems. *IEEE TSE*, 25(1), 1999.
- [6] Chuah and Eick. Information rich glyphs for software management data. *IEEE Computer Graphics and Applications*, 18(4), July 1998.
- [7] D'Ambros and Lanza. Reverse engineering with logical coupling. In *WCRE*, 2006.
- [8] Denker and Ducasse. Software evolution from the field: an experience report from the Squeak maintainers. In *ERCIM Working Group on Soft. Evolution*, vol. 166 of *Electronic Notes in Theoretical Computer Science*, Jan. 2007.
- [9] Ducasse, Gîrba, and Wuyts. Object-oriented legacy system trace-based logic testing. In *CSMR*, 2006.
- [10] Ducasse, Lanza, and Ponisio. Butterflies: A visual approach to characterize packages. In *Int'l Soft. Metrics Symposium (METRICS)*, 2005.
- [11] Eick, Graves, Karr, Mockus, and Schuster. Visualizing software changes. *IEEE TSE*, 28(4), 2002.
- [12] Froehlich and Dourish. Unifying artifacts and activities in a visual tool for distributed software development teams. In *ICSE*, 2004.
- [13] Gîrba, Kuhn, Seeberger, and Ducasse. How developers drive software evolution. In *Int'l Workshop on Principles of Soft. Evolution (IWPSE)*, 2005.
- [14] Healey. Visualization of multivariate data using preattentive processing. Master's thesis, Univ. British Columbia, 1992.
- [15] Healey, Booth, and T. Harnessing preattentive processes for multivariate data visualization. In *Graphics Interface*, 1993.
- [16] Herman, Melançon, and Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Trans. on Visualization and Comp. Graphics*, 6(1), 2000.
- [17] Jain, Murty, and Flynn. Data clustering: a review. *ACM Computing Surveys*, 31(3), 1999.
- [18] Lanza and Marinescu. *OO Metrics in Practice*. 2006.
- [19] Lungu, Lanza, and Gîrba. Package patterns for visual architecture recovery. In *CSMR*, 2006.
- [20] Mancoridis, Mitchell, Chen, and Gansner. Bunch: A clustering tool for the recovery and maintenance of software system structures. In *ICSM*, 1999.
- [21] Martin. Design principles and design patterns, 2000.
- [22] Mitchell and Mancoridis. On the automatic modularization of software systems using the bunch tool. *IEEE TSE*, 32(3), 2006.
- [23] Pinzger, Gall, Fischer, and Lanza. Visualizing multiple evolution metrics. In *SoftVis*, May 2005.
- [24] Pollet, Ducasse, Poyet, Alloui, Cîmpan, and Verjus. Towards a process-oriented software architecture reconstruction taxonomy. In *CSMR*, Mar. 2007.
- [25] Ponisio and Nierstrasz. Using context information to re-architect a system. In *Soft. Measurement Eur. Forum*, 2006.
- [26] Sangal, Jordan, Sinha, and Jackson. Using dependency models to manage complex software architecture. In *OOPSLA*, 2005.
- [27] Storey, Čubranić, and German. On the use of visualization to support awareness of human activities in software development: a survey and a framework. In *SoftVis*, 2005.
- [28] Storey, Wong, Fracchia, and Müller. On integrating visualization techniques for effective software exploration. In *IEEE Symposium on Information Visualization (InfoVis)*, 1997.
- [29] Treisman. Preattentive processing in vision. *Computer Vision, Graphics, and Image Processing*, 31(2), 1985.
- [30] Tufte. *The Visual Display of Quantitative Information*. 2001.
- [31] Voinea, Telea, and van Wijk. CVSScan: visualization of code evolution. In *SoftVis*, May 2005.
- [32] Ware. *Information Visualization*. 2000.
- [33] Xie, Poshyvanyk, and Marcus. Visualization of CVS repository information. In *WCRE*, 2006.