

A Browser for Incremental Programming

Nathanael Schärli and Andrew P. Black

Software Composition Group, University of Bern, Switzerland

OGI School of Science & Engineering, Oregon Health and Science University, USA

Abstract

Much of the elegance and power of Smalltalk comes from its programming environment and tools. First introduced more than 20 years ago, the Smalltalk browser enables programmers to “home in” on particular methods using a hierarchy of manually-defined classifications. By its nature, this classification scheme says a lot about the *desired* state of the code, but nothing at all about the *actual* state of the code as it is being developed. We have extended the Smalltalk browser with dynamically computed *virtual categories* that dramatically improve the browser’s support for incremental programming. We illustrate these improvements by example, and summarize the algorithms used to compute the virtual categories efficiently.

Key words: Smalltalk browser, incremental programming, intentional programming, method reachability, requires set.

1 Introduction

The most important of the Smalltalk programming tools is the *Browser*, which allows the programmer to examine, modify and extend the code of applications and of the system itself. The Browser was revolutionary when it was first introduced, and over the intervening years it has been improved in several ways. For example, semi-automated refactoring have been added, leading to a tool known as the Refactoring Browser [10], and in many Smalltalk dialects some form of package construct has subsumed the original primitive categorization of classes. However, the way in which today’s browser organizes the methods of a class is essentially the same as it was in 1980: a hierarchy of manually assigned “protocols”.

In the meantime, the concept of the “Integrated Development Environment”—for that is what the Smalltalk toolset would now be called—has proved to be so successful that similar environments have been created for other programming languages. For example, IBM’s VisualAge for Java [7] was essentially a re-targeting of

Email address: schaerli@iam.unibe.ch, black@cse.ogi.edu (Nathanael Schärli and Andrew P. Black).

the Smalltalk IDE to Java; more recently the cross-language environment Eclipse [3] has been making similar tools available for many popular languages. As these IDEs have evolved, they have started to perform real-time analysis of the code, and now typically provide the programmer with dynamic feedback such as syntax coloring and message prompting.

The raw material for real-time feedback is adequate computational power. Today's Smalltalks run on a commodity laptop about a thousand times faster than they ran on commodity machines of the early 1980s, and about fifty times faster than on the custom hardware of the Dorado. This power has been harnessed to write Smalltalk applications for real-time music and motion picture manipulation that would have been unthinkable twenty years ago. But the power has *not* been used to build tools that dynamically re-compute automatically defined views of an evolving program.

In this paper, we present one such tool. It derives its power from a real-time analysis of the program being modified. This enables us to infer and display system-wide information about the structure of the program, actively supporting incremental programming and program understanding.

The basic idea is to display a classification of the methods of a class C into “virtual categories”, *i.e.*, categories that are computed automatically by the browser and always kept up to date. One of the most useful is the *requires* category, which lists all of the messages that are sent to **self** by the methods of C , but for which methods are neither defined explicitly nor inherited. We also display the *supplies* category, which contains the concrete methods of C that implement inherited abstract methods, and the *overrides* category, which contains the concrete methods of C that override inherited concrete methods. The last virtual category, *sending super*, contains the methods that perform super-sends.

Of these virtual categories, the first—the set of required methods—is the one that we have found to be most useful. The display of this category supports “programming by intention” [6], a top-down style that encourages the programmer to think about *what* has to be done rather than about *how* to do it. The idea behind programming by intention is to imagine methods that do “the hard part” of one’s task, so that all that one has to do to complete the task is send the corresponding messages. Of course, later one applies the same idea to defining the “hard” methods. In this paradigm, the *requires* category provides a constant reminder of what is left for the programmer to do.

We have implemented all of these virtual categories in a browser for Squeak, an open-source dialect of Smalltalk [4,14]. In addition to being the most useful, the *requires* category also turns out to be the trickiest to define and implement efficiently; to the best of our knowledge there is no other browser that displays it. In contrast, the definition and computation of the other virtual categories is mostly straightforward, and recent versions of other Smalltalks (*e.g.*, VisualWorks [16]

and Squeak [14]) also compute some similar properties and use the results to decorate method names in the browser. However, no other browsers use this information to group methods into virtual categories that supplement the programmer-defined categories and provide the programmer with multiple views on the code. A Squeak image that demonstrates our browser, and contains its source code, is available [11].

This paper makes three contributions. First, we analyze the information that a programmer needs in order to be effective at incremental programming (section 2). Second, we describe the way that our browser supplies this information, and report on our experience using the extended browser and the way in which it changes the programming process (section 3). Third, we define the *required* set and describe an algorithm to compute it that runs in real-time (section 4).

2 Incremental Programming and its Demands

Even though incremental programming has been part of the Smalltalk ethic from the earliest days, it seems not to have been singled-out as an important feature. Ingalls' seminal paper on the design principles behind Smalltalk [5], for example, does not mention support for incremental programming. So it is necessary for us to ask ourselves what incremental programming really is, and how a development environment should support it.

Fortunately, we do not have to look far for our answers. Over the last few years, Kent Beck and others have codified a set of practices called Extreme Programming (XP) [1], which capture very clearly the essence of incremental work: not just incremental programming, but also incremental requirements gathering, incremental design, incremental planning, and incremental deployment. This means that instead of bringing the work on each of these activities to an end before proceeding to the next, programmers iterate through all of the activities multiple times, each time doing only as much work as is necessary to achieve the current (incremental) goal. We will take the "extreme" style of programming as a paradigmatic example of what a good incremental programming environment should support.

2.1 *Programming in an Extreme Environment*

In an "XP shop", the incremental and iterative style of work is applied not only to the major development activities, but also to the work within each activity. During implementation, this means that classes are not written sequentially, one after the other, and finally executed only when they are all complete. Instead, a programmer works on several classes at a time, and combines and tests them as soon as a fraction of the functionality is implemented.

The following patterns of work seem to be central to incremental programming.

- (1) *Programming with limited knowledge.* A programmer starts implementing a part before he has full knowledge of the whole system.
- (2) *Working in multiple contexts.* A programmer may be working on several classes in parallel. For example, he may start implementing a new method before finishing the implementation of another. As a consequence, he often switches working context.
- (3) *Understanding how classes collaborate together.* When working with multiple incomplete classes at once, the programmer needs to understand how these classes work together and how changes in one class affect related classes.
- (4) *Understanding what is still missing.* With incremental programming, it is important to know what parts are still missing in order to make a program, or at least a part of it, complete. The missing parts are a good indication of what the programmer should work on next.
- (5) *Refactoring.* A consequence of starting to program with limited knowledge is that it is likely that the implementation will have to change as the programmer becomes more knowledgeable. XP says that we should embrace change, not be fearful of it. Refactoring is the secret weapon that enables us to combat the increasing entropy and code rot that would otherwise be the result of continual redesign and re-implementation.
- (6) *Testing.* Writing and maintaining tests is a corner-stone of incremental and iterative programming. Tests are used to capture the intended behavior of a feature as soon as that feature is implemented—or even before. Tests make it possible to refactor quickly and safely, and help us understand what is still missing.

2.2 *Supporting Incremental Programming in Smalltalk*

Smalltalk was designed to support experimental programming and its integrated environment encourages the programmer to alternate between design, coding, testing and debugging [15]. Therefore, it is not surprising that the Smalltalk language and its development environment provide rich support for four of the six patterns of work identified above.

Since the Smalltalk language is not explicitly typed, the programmer does not have to specify the type of method arguments or instance variables when they are introduced. As a consequence, programming in Smalltalk requires less knowledge about the design of the whole system in order to start implementing (1).

The earliest implementations of Smalltalk-80 came with a windowing system that allowed one to open multiple browser windows and work concurrently in multiple contexts (*e.g.*, with multiple classes and methods) (2). More recently, this capability has been drastically improved with new browsers such as Whisker and the Star Browser, both of which allow one to work in multiple contexts without having to open multiple windows. See Section 5 for more information about Whisker, the

Star Browser, and all the other Smalltalk extensions mentioned in the section.

Since the introduction of the Refactoring Browser Smalltalk has been able to command a rich set of tools for semi-automated refactoring (5). Incremental testing was originally supported by the practice of writing executable comments and example methods. Because of incremental compilation [15], and because no checks for completeness are performed at compile time, it is possible to test a method immediately after it is written. Since the introduction of the SUnit testing environment, it is now also possible to effectively organize these tests (6).

Unfortunately, Smalltalk provides only limited support for the two remaining patterns of work, which means that it is relatively hard to understand how different classes relate to each other (3) and what is still missing to make them complete (4).

We believe that these problems can be solved by providing the programmer with statically accessible and always up-to-date information about ways that the different classes of a system collaborate. The task of providing this information can be broken into several subtasks. First, we need to identify what kind of information is needed to understand how the classes collaborate. Second, we need to be able to compute this information in real-time. Finally, we need to find a way to make this information readily accessible to the programmer.

In the rest of this paper, we explain how we completed these tasks for the most common form of collaboration between classes: inheritance. In Section 5, we briefly point out how the same techniques can be applied to help the programmer understand aggregation and delegation, the other major forms of collaboration in object-oriented languages.

3 The Browser

The previous section identified some limitations in the support offered for incremental programming by existing Smalltalk browsers. Now we describe our new browser and how it overcomes these limitations. First, we discuss the categories of information presented by our browser, and explain why this information makes it easy to understand inheritance collaborations. Then we give an “illustrated walk-through” of a programming session, showing how the browser facilitates incremental programming and helps programmers to understand existing class hierarchies.

3.1 Information Presented by the Browser

Our browser is shown in figure 1. At first glance it looks like the standard Smalltalk browser, but a few extra features are visible. In the figure, the class Integer is selected in the class pane (the second from the left). The third pane, which in the

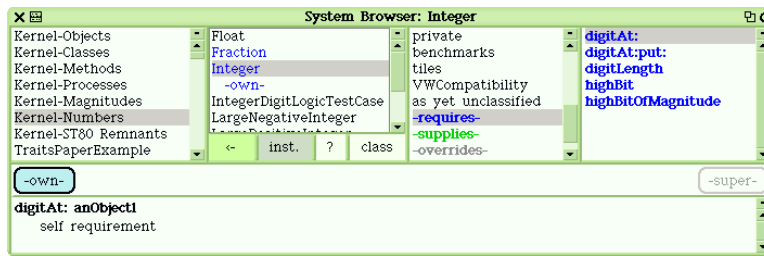


Fig. 1. The browser examining the class Integer.

standard browser contains a manual categorization of the methods of the selected class, now contains in addition some *virtual categories*. These categories present information that we identified as key to understanding how a class collaborates with its neighbors in the inheritance hierarchy.

The category *-requires-*, which is colored blue, includes all of the methods that the class Integer sends to itself but does not define or inherit. In our figure, the list of these required methods appears in the fourth pane, where we have selected `digitAt:`, which is consequently displayed in the code pane at the bottom of the browser. The implementation shown, `self requirement`, is a marker method generated by the browser to indicate that `digitAt:` is an unsatisfied requirement. The same holds for the requirements `digitAt:put:` and `digitLength`. The remaining two requirements, `highBit` and `highBitOfMagnitude`, do in fact have implementations: `self subclassResponsibility`. However, the browser recognizes this as a marker method and still categorizes them as required. The required methods tell the programmer how the class Integer is parameterized, and which methods are necessary in order to make it complete.¹

The next category, *-supplies-*, lists methods that are required by some other class and provided by the class that we are browsing. In the case of Integer, this virtual category contains 10 methods including `*`, `+`, `-`, `/`, `<`, `=`, and `hash`. This tells the programmer that Integer's superclass (Number) is parameterized by these 10 methods and shows the concrete implementations that Integer supplies for these parameters.

The third category, *-overrides-*, lists those methods provided by Integer that override methods inherited from its parents. In our example, there are 11 such methods, including `//`, `asInteger`, `even`, `floor`, and `isInteger`. The *-overrides-* category is important for two reasons. First, it gives the programmer a view on the class as a “delta”, that is, the methods in this category characterize the parts of the behavior of the parent class that are changed by this subclass. Second, the overriding methods, together with the supplied methods, are the most critical for the understanding of the class Integer and reasoning about its correctness. This is because inheritance breaks encapsulation [13]: subclasses collaborate with their superclasses through a much broader interface than the public interface of the superclass. Taken together, overridden and supplied methods represent the hooks through which the behavior of

¹ Note that in Smalltalk, Integer is the abstract superclass of the concrete classes SmallInteger, LargePositiveInteger, and LargeNegativeInteger.

Number collaborates with the behavior specified in Integer. In particular, this means that the behavior implemented in each of these methods of Integer needs to conform to the specification implied by Number, and that these are the methods in Integer that need to be adapted whenever the specification in Number changes.

The fourth virtual category is called *-sending super-*; it contains the methods that perform *super sends* (i.e., message sends that cause the message lookup to be started in the superclass). This category is important because it tells the programmer which of the methods collaborate with behavior from their parent class. Furthermore, all of these methods depend explicitly on their position in the inheritance hierarchy, so special care has to be taken if they are moved during refactoring. This category is not shown in the figure because it is empty. In fact, all the virtual categories are displayed in the browser only if they are not empty.

Although it is not clear from the grayscale figure, each of the generated virtual categories has a characteristic emphasis: blue for *requires*, green for *supplies*, grey for *overrides*, and underlined for *sending super*. Even when browsing methods using the ordinary, manually-defined method categories, the names keep their characteristic emphasis. So a supplied method that sends to super will always show up in green and underlined. The blue color-coding is also applied to the name of the class itself in the second pane whenever the set of required methods is not empty. This serves as a reminder that the class is incomplete, e.g., it may be an abstract class such as Integer, or the programmer may still be working on it.

3.2 Programming with the browser

Now that the reader understands the information presented by our browser, we will illustrate how this information significantly facilitates programming in general, and incremental programming in particular.

As a working example, we will assume that we are about to implement a hierarchy of classes for collections in Smalltalk. We start by implementing the class Collection, which serves as the abstract root class of our hierarchy. At this level, we will implement the most common interface methods such as *isEmpty*, *includes:*, *select:*, *add:*, *remove:*, etc.. While writing these implementations, we proceed in an “intentional” style: we need not worry about sending messages that are not yet implemented. This is because our browser updates all the virtual categories in real-time, and as a consequence every message that is sent to self but not yet implemented immediately appears in the *-requires-* category.

As a concrete example, consider the method *isEmpty*, implemented as

```
isEmpty  
↑ self size = 0
```

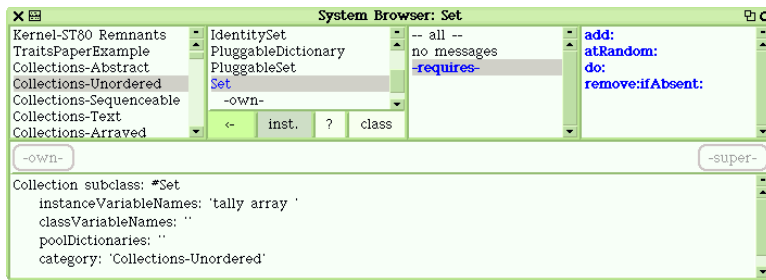


Fig. 2. The browser on the class `Set` just after it was created.

As soon as this method is accepted, a requirement size is created. It will not disappear until it is satisfied (*i.e.*, a size method is implemented) or all the methods that require size are removed. The same thing happens when we implement the method `select:`, which requires the methods `do:` and `emptyCopyOfSameSize`.

The constant availability of the *requires* category gives us the freedom to implement the methods in the order that they come to mind without having to worry about forgetting to implement the methods on which they depend. Furthermore, many errors, especially typographical errors, can be detected immediately because they cause weird requirements to suddenly appear. A new menu item, *local senders of ...* helps us discover why a method is required: it lists the methods in the current class hierarchy that send the selected message.

Once we have implemented the most essential `Collection` methods, we create the concrete subclass `Set` so that we can test our abstract implementation. Figure 2 shows the class `Set` just after it has been created as a subclass of `Collection`. At a glance, we see that this class is incomplete (the name `Set` is blue), that it does not yet implement any methods, and that it requires the methods `add:`, `atRandom:`, `do:`, and `remove:ifAbsent:` to make it complete. (Although `Set` does not yet have any of its own methods, it of course inherits methods from `Collection`; these requirements are inferred by analyzing the inherited methods.)

In order to properly satisfy these requirements, we will eventually need to implement some internal methods that map the contents of the set to the instance variables `array` and `tally`. Again, the automatically computed list of requirements allows us to implement these methods in whatever order we prefer, while keeping a high-level view of what is already implemented and what still needs to be done.

Figure 3 shows our browser on the class `Set` after it has been completed. Because there are no longer any unsatisfied requirements, the *-requires-* category has disappeared. Instead, the category *-supplies-* appears, showing how the parameters of `Collection` are satisfied by `Set`. We also see the category *-overrides-*, which contains methods such as `=`, `asSet`, `copy`, `includes:` and `size`. These are the methods that the class `Set` overrides in order to modify the default behavior or to provide a more efficient version.

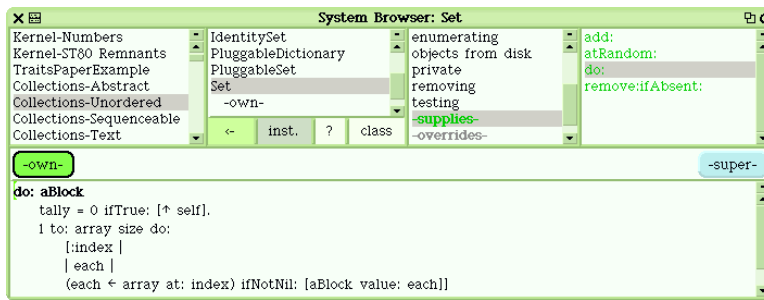


Fig. 3. The browser examining the class Set after it is completed.

Besides the fact that these two virtual categories are essential for understanding the relationship of Set to its parent, they also help the programmer to avoid “inter-level” errors between these classes. Because the methods in these categories capture the collaboration between Collection and Set, they are exactly the methods that must be examined if the specification in Collection is changed. For example, if we change the way that Collections are compared, we need to ensure that the overridden method = is updated accordingly.

After testing Set, let us suppose that we move on to create some sequenced collection classes such as OrderedCollection and Heap. Since these classes will have some code in common, we decide to make them both subclasses of a new abstract class SequenceableCollection, which will itself be a subclass of Collection. When looking at these three classes with the browser, we see that they all require the same methods, namely add:, atRandom:, do:, and remove:ifAbsent:.

Instead of implementing these three classes one after the other, we can now take advantage of the freedom granted by incremental programming and work on all three classes in incremental steps. This means that we go through these requirements and the necessary internal methods and decide for each method in which class it should be implemented. During this process, which may also include refactoring actions such as pushing methods up or pulling them down in the hierarchy, the features of our browser always give us up-to-date information about what methods are required for each of these classes and how they collaborate.

Concretely, we might start with the requirement do:, and decide that we should implement this in the class SequenceableCollection as follows.

```
do: aBlock
  1 to: self size do: [:index | aBlock value: (self at: index)]
```

As soon as we accept this code, the browser shows us that we now have a new requirement at: in SequenceableCollection and its two subclasses. Since at: obviously must be implemented differently in OrderedCollection and Heap, we then proceed to implement the two at: methods in the subclasses. This is reflected in the browser, which shows at a glance that at: is now a parameter of the abstract class SequenceableCollection and that it is *supplied* by different methods in both of the concrete subclasses.

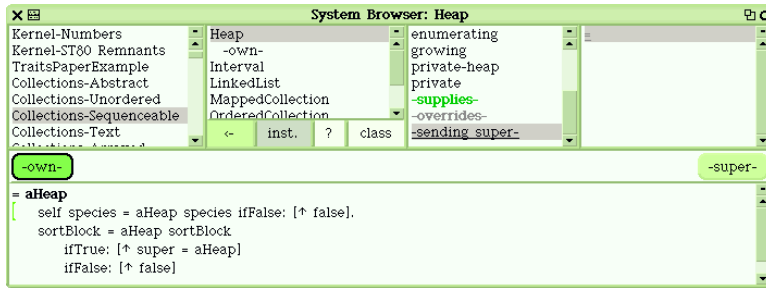


Fig. 4. The browser examining the class Heap.

Figure 4 shows the class Heap as it results from this process. In the figure, we have selected the virtual category *-sending super-*, which contains a single method =. This tells us that there is only one method that uses the keyword **super** to access an overridden implementation in the parent class and which is therefore vulnerable to refactorings. The two other virtual categories, *-supplies-* and *-overrides-*, tell us how Heap implements the parameters of its abstract parent class SequenceableCollection and which of the inherited methods it overrides.

3.3 Understanding and modifying existing classes

In the previous section, we have shown how our extended browser helps one to create a new class hierarchy. In addition, our experience has shown that the browser is also a great help when one must understand and modify existing classes and class hierarchies.

As an example, consider the collection hierarchy that we have just created. If a new programmer looks at this hierarchy, the virtual categories of the browser provide valuable guidance in understanding both the higher-level purpose and the lower-level implementation of each of the classes. This is because these categories separate information that is essential to understanding the collaboration from the rest of the code, and thus make the overall architecture more explicit. Once the programmer has a basic understanding of the parent class, she will need to look only at a small part of the code of a subclass in order to understand how that subclass relates to its superclass.

These features are even more important when it comes to extending a class in an existing hierarchy with some new features. Such an extension is quite common, particularly if the hierarchy was created incrementally, or following the XP doctrine of doing the simplest thing that could possibly work to satisfy the requirements that are currently available [1, Chapter 17]. Especially in complex hierarchies, it is quite easy to introduce “inter-level” errors, for example, by modifying or extending a class and forgetting to make all of its subclasses compatible with these changes. Not only does the browser help to avoid making such inter-level errors, it also shows us an overview of which subclasses accidentally became abstract, and are therefore

incomplete.

Although it might seem that these problems are easy to avoid without virtual categories, our experience has shown that this is not the case. When we introduced our new browser into Squeak, we immediately found dozens of abstract classes that should actually be concrete, as well as hierarchies that exhibited other inter-level errors such as sending unimplemented super messages. The most surprising thing was that we found these errors even in the core of the system, which has been used by thousands of users for many years. For example, classes like `Fraction`, `Bitmap`, `CharacterSet`, `Debugger`, and nearly all subclasses of the class `Morph`—which is the root class of Squeak’s user interface framework—are accidentally abstract, simply because programmers forgot to implement some methods.

We find this to be convincing evidence not only that our browser extensions are necessary to avoid such errors in the future, but also that they provide great help in finding and eliminating these errors in existing class hierarchies. Since abstract classes are “blue”, finding them is trivial. Having the requirements continuously available in an up-to-date virtual category also makes it straight-forward to identify the cause of the problem and eliminate it.

4 Implementation

Most of the virtual categories that are displayed in our browser can be computed straightforwardly. the exception is the *requires* set. In this section, we therefore focus on an algorithm for computing the *requires* set efficiently enough for it to be displayed in real-time.

To compute the required methods of a class we must consider all of the class’s *reachable* methods, that is,

- (1) all methods that are locally defined in the class,
- (2) all non-overridden methods defined in its superclasses, and
- (3) all methods that may be reached by super-sends from other reachable methods.

The *requires* set of a class contains all of the message selectors sent to `self` in one of the reachable methods, minus the selectors of the methods provided by the class (and by its superclasses). These definitions have been formalized in a technical report [12].

In order to compute this set, we first must find the self-sends and super-sends of a method. Whereas the super-sends can be immediately retrieved from the byte-code, computing the self-sends is more complicated, because they do not all emanate from a single syntactic construct. Consider, for example, the following method.

```

fasten
  | anObject |
  self hook.
  anObject := self.
  anObject button.
  self class new clip.

```

From the byte-code of this method, it is immediately clear that `hook` is sent to `self`. What about `button` and `clip`? These messages are also sent to an instance of the current class, and so they are really self-sent too. However, detecting this requires a deep analysis of this method, as well as of the method `new` on the class side.

Our current implementation does not carry out such an analysis. This means that in the above method, `hook` is the only self-send that we would detect. We compensate for this deficiency by allowing the programmer to declare *explicit requirements* by implementing a method with the body `self explicitRequirement`.

Even with this simplification, computing the requires set in real-time is quite challenging. The main problem is that a single change in a class may affect the requires set of all its subclasses. Thus, a change in `Object` may mean that we have to update the required methods of all of the classes in the system. A naive implementation based on the above definition would be far too slow to provide the programmer with useful feedback (see section 4.4 for a performance comparison). Updating the requires set in real-time required an optimized algorithm that caches critical data and takes advantage of the coherence of the inheritance hierarchy.

4.1 Caching Self-Sends and Super-Sends

Because computing the self-sends of a method requires some time-consuming analysis of the code, and because recent Squeak images contain more than 60 000 methods, we decided to cache the self-sends for every method. This means that the self-sends of a method are inferred only once: when the method is first created.

When computing the requires set, looking for methods that contain a certain self-send is far more common than looking for (or modifying) the self-sends of a particular method. Therefore we index the caches by the sent selectors: for each class C we maintain a dictionary whose keys are the selectors that are self-sent by the methods directly implemented in C , and whose values are the array of selectors that name the methods that perform those self-sends. For example, if the selector x is sent to `self` by the local methods on `a` and `b`, looking up x returns the array `#(a b)`.

Super-sends are critical for determining the set of reachable methods, so we also maintain a cache of the super-sends that are issued by the local methods of each class; this cache is similarly indexed by the super-sent selectors rather than by the selectors of the methods that perform the super-sends.

4.2 Using the coherence of the inheritance hierarchy

When a method is added, modified or removed, we need to check its class, and all its subclasses, to see whether there is any effect on the requirements. The heart of this computation is checking whether a given selector is self-sent in a given class. Especially in the case of large hierarchies, performing this check separately for each subclass proved to be far too slow. Therefore, we developed a recursive algorithm that takes advantage of the coherence that typically exists between neighbor classes in the inheritance hierarchy.

The main method of our algorithm is applied initially to the topmost class C ; it ascertains recursively, for C and each of C 's subclasses, whether that class self-sends a given selector x . As the method proceeds down the inheritance hierarchy, it keeps track of the set of methods that self-send x . Unless *all* of these methods are overridden in the next subclass C' , we immediately know that C' also self-sends x . Otherwise, we call a helper method that tries to find some methods that contain self-sends to x and are actually reachable from this subclass. The helper method first searches the local methods of C' , and then recurses through all the superclasses until it either finds such a method or reaches the top of the hierarchy. While going up the inheritance hierarchy, we maintain the set of all the unreachable methods in order to avoid false positives.

We now give a more detailed description of both the main method and the helper method.

Main method. The main method of the algorithm searches through the argument class C and all its subclasses and ascertains which of these classes self-sends the selector x . In addition to C and x , this method also takes an argument S , a set containing the superclass methods that are known to issue self-sends to x . This set is empty when the main method is called on the first class of the hierarchy. The method proceeds as follows.

- (1) We identify $S' \subseteq S$ that are reachable from the class C . This means checking which methods in S are not overridden in C , or, if they are overridden, are nevertheless reachable by a super-send from a method in C .
- (2) If S' is empty, we call the helper method discussed below to search for reachable methods that self-send the selector x . The result is stored as S' .
- (3) If S' is still empty, we indicate that C does not self-send x . Otherwise we indicate that the class C does self-send the selector x .
- (4) We perform a recursive call for each of the direct subclasses of C , passing S' as a parameter.

Helper method. The purpose of the helper method is to find a set of methods that contain self-sends to a selector x and are reachable from the class C . In addition to C and x , this method takes a third argument U , the set of all selectors that have

been found to be unreachable. This set is empty when the method is called from the main method. The computation proceeds as follows.

- (1) We check whether the class C contains local methods that issue a self-send to x and whose selectors are not in the set U of unreachable selectors. If we find some, we return them and exit.
- (2) Before we use recursion to search the superclass methods, we construct the set U' of all the superclass selectors that are unreachable from the class for which this helper method was initially called. Since all the unreachable selectors in C remain unreachable, the set U' includes all the selectors in U . In addition, U' contains all the selectors for which methods are defined in C (and therefore potentially override superclass methods) and which are not super-sent in C .
- (3) We use recursion to find superclass methods that contain a self-send to x . In order to avoid finding unreachable methods, we pass the set U' as a parameter. These methods are stored as R .
- (4) For each method in the set R , we check whether it is reachable *only* by a super-send. If so, we replace it by the local method that performs the super-send. Then we return the resulting set R .

4.3 Discussion

The above description conveys the basic idea of our algorithm, but omits some details. We now discuss some of these details and also make some general remarks.

Optimization of the helper method. In our implementation, we use an optimized version of the helper method. The main difference from the description above is that we employ temporary caches to avoid repeating the computation in step (3). Without this optimization, the helper method might, for example, compute the self-sends to a given selector in class `Object` over and over again, even though it may already be clear that no such self-sends exist.

It is important to note that the helper method does not in general compute *all* the reachable methods that self-send the given selector. Our experience has shown that it is more efficient to stop at whatever level of the hierarchy we find the first self-sends. This is because overrides are relatively rare, and therefore the set S in the main method does typically not have to be very big in order to avoid multiple calls of the helper method.

Cache access. We see that the caches are well-suited to the algorithm. In step (1) of the main method and steps (2) and (4) of the helper method, we are able to check whether a selector is reachable via a super-send with a single lookup in the super-send dictionary. Similarly, in step (1) of the helper method, we can find all the local methods that issue a self-send to x by a single lookup in the self-send dictionary.

Cache consistency. The caches can be kept in a consistent state quite cheaply. This is mainly because the caches contain only local data, that is, the the cache for a particular class is independent of all the other classes in the hierarchy. Thus, modifying a class requires updating at most the local cache for that class. Furthermore, changing the place of a class in the hierarchy does not affect the caches at all.

4.4 Performance Comparison

Our first implementation of the browser deployed caches for the self-sends and super-sends of every method. However, unlike the approach presented above, these caches were indexed by the selectors of the methods performing the self-sends rather than by the selectors that were sent. Furthermore, our initial algorithm did not take advantage of the coherence in the class hierarchy. This meant that the method for finding out whether a class C self-sends a selector x was applied to each class separately. Using this implementation to find out which classes in the system required the selector + took several minutes (188 seconds)², which made it impossible to provide immediate feedback.

In a second version of the algorithm, we used the same caching strategy, but took advantage of the coherence in the class hierarchy. Performance improved significantly, but the same computation still took over 9 seconds. Finally, using the caching strategy and the algorithms presented above, the same test takes less than 100 milliseconds and thus meets our requirement for instantaneous feedback.

5 Related and Future Work

The idea of keeping an automatically updated list of things that remain to do dates back at least as far as the “grass catcher” of Trellis [9], and has been adopted in some form or other in many IDEs; the “Tasks” window of Eclipse is another example. However, such lists are typically a by-product of a global re-compilation, rather than being constructed modularly as the consequence of a change to a single method, as in Trellis and in our virtual categories.

Recently, there have been other extensions to Smalltalk browsers that provide the programmer with automatically updated information about the code. In the introduction, we mentioned that there are Smalltalk implementations such as Visual-Works that decorate method names in the browser to indicate things such as super-sends or overrides. However, these browsers neither use this information to group methods into virtual categories nor compute the required methods.

² All the performance data provided in this paper were measured in a Squeak 3.2 image consisting of 1860 classes, and were executed on a Mobile Pentium III 1.2GHz with 512MB RAM.

SOUL, the Smalltalk Open Unification Language [17], is an open, reflective logic programming language written in VisualWorks Smalltalk. Although it is not itself a browser, it can be used to extend a Smalltalk browser with many kinds of logical reasoning, triggered by the occurrence of an action such as accepting a method. This ability has been used to formulate queries that look for composite patterns, and to infer the types of instance variables. SOUL has also been used to create “virtual classifications” that map source-level artifacts to higher-level architectural components (and vice versa) [8].

The Refactoring Browser adds many useful features for semi-automated refactoring to the standard Smalltalk browser, and is available for several Smalltalk implementations including Dolphin Smalltalk [2], VisualWorks, and Squeak. Refactorings supported by this browser include moving methods up and down the inheritance hierarchy, to extracting a block of code into a method, and renaming instance variables and methods. The Refactoring Browser uses type inference to help the user to deciding which senders actually refer to the method being renamed

Whisker is a Smalltalk browser implemented in Squeak. Whisker’s main contribution is a screen layout that provides a simple and intuitive way to view the contents of multiple classes and methods simultaneously, while using the screen efficiently and avoiding the need to manually move and resize windows. Whisker does this by using subpane stacking, *i.e.*, dynamically stacking subpanes into a single column. Whisker also infers and displays information about the classes of objects that are bound to the instance variables of a class.

Another extended browser is the Star Browser [18], which is available for Squeak and VisualWorks. Like Whisker, the goal of the Star Browser is to allow the programmer to establish a working context without having to deal with too many windows. Unlike Whisker, the Star Browser is built on top of a lightweight classification model that allows one to categorize any sort of item such as classes, methods, method categories, *etc.*. Besides *extensional classifications* that are just bags of items, the model also allows one to express *intentional classifications*, which are defined by a description, and whose contents are updated automatically.

Despite these improvements, we believe that there remain many opportunities to improve the Smalltalk programming environment. One example is applying the techniques of this paper to provide the user with always up-to-date information about the collaborations between classes in aggregation and delegation relationships. We believe that identifying the key information necessary to understand these collaborations could lead to tool and process improvements similar to those that we have observed in the case of inheritance. For example, a browser might support a virtual category showing the methods that a class requires in order that its instances understand all the messages delegated to them.

However, the process of actually computing this information will be much harder

because, unlike inheritance, aggregation and delegation are dynamic collaborations. Especially because Smalltalk is not statically typed, it is not immediately clear which instance variables and method arguments may refer to instances of which classes. Nevertheless, this problem can be tackled by combining the results of type inference with information about the classes of instance variables and method arguments that is gathered dynamically. Although this will not always lead to completely accurate information, it seems likely that approximate information will be quite adequate in practice.

6 Conclusion

We have identified four virtual categories that, if updated in real-time, provide valuable insight on the program under development. In fact, this categorization gives the programmer a different view on the code: it structures the space of methods in a way that is quite different from the explicitly declared protocols. Whereas the protocols group the methods according to their role in the domain logic (*i.e.*, testing, printing, model access, *etc.*) the virtual categories group the methods according to their role in the composition, that is, the way that the class interacts with its neighbors in the inheritance hierarchy.

Our experience with the extended browser has shown that this viewpoint is very valuable both for writing and for understanding the code. While writing, it supports an incremental style of programming: the programmer can freely compose components and add methods, and rely on the browser to maintain an overview of what still remains to be done and where possible problems (*e.g.*, open requirements, conflicts, and overrides) might lie. Later, the same view helps a programmer to understand the code, because at a glance she can see all the critical methods that are essential for understanding the interaction between the various components. This stands in contrast to the conventional viewpoint, which leaves the programmer the task of finding these critical methods by looking through many methods (sometimes hundreds), spread over many protocols.

From an implementation point of view, most of these categories are quite easy to compute. By far the most challenging is the computation of the requires set. There are three reasons for this. First, the absence of explicit type information makes it hard to detect the requirements; second, finding whether a method is reachable turns out to be non-trivial; and third, heavy optimization is required to obtain real-time performance.

Nevertheless, our experience has shown that these obstacles can be overcome. Although we choose a very simple algorithm to infer self-sends, in practice most of the requirements are detected. Furthermore, caching of information about the requirements means that the requires set can be re-computed quickly, even after a change that affects many classes.

References

- [1] Kent Beck. *Extreme Programming Explained: Embrace Change*. Addison Wesley, 2000.
- [2] Dolphin Smalltalk. <http://www.object-arts.com/DolphinSmalltalk.htm/>.
- [3] Eclipse Platform: Technical Overview, 2003. <http://www.eclipse.org/whitepapers/eclipse-overview.pdf>.
- [4] Dan Ingalls, Ted Kaehler, John Maloney, Scott Wallace, and Alan Kay. Back to the future: The story of Squeak, A practical Smalltalk written in itself. In *Proceedings OOPSLA '97*, pages 318–326, November 1997.
- [5] Daniel H. Ingalls. Design principles behind Smalltalk. *Byte*, 6(8):286–298, August 1981.
- [6] Ron Jeffries, Ann Anderson, and Chet Hendrickson. *Extreme Programming Installed*. Addison Wesley, 2001.
- [7] Niraj Jetly. VisualAge for Java 2.0. *Java Developer's Journal*, 4(4):48–49, April 1999.
- [8] Kim Mens, Roel Wuyts, and Theo D'Hondt. Declaratively codifying software architectures using virtual software classifications. In *Proceedings of TOOLS-Europe 99*, pages 33–45, June 1999.
- [9] Patrick D. O'Brien, Daniel C. Halbert, and Michael F. Kilian. The Trellis programming environment. In *Proceedings Object-oriented programming systems, languages and applications, ACM SIGPLAN Notices*, volume 22, pages 91–102. ACM Press, October 1987.
- [10] Don Roberts, John Brant, and Ralph E. Johnson. A refactoring tool for Smalltalk. *Theory and Practice of Object Systems (TAPOS)*, 3(4):253–263, 1997.
- [11] Nathanael Schärli. Traits prototype in Squeak, February 2003. <http://www.iam.unibe.ch/~schaerli/smalltalk/traits/traitsPrototype.htm>.
- [12] Nathanael Schärli, Oscar Nierstrasz, Stéphane Ducasse, Roel Wuyts, and Andrew Black. Traits: The formal model. Technical Report IAM-02-006, Institut für Informatik, Universität Bern, Switzerland, November 2002. Also available as Technical Report CSE-02-013, OGI School of Science & Engineering, Beaverton, Oregon, USA.
- [13] Alan Snyder. Encapsulation and inheritance in object-oriented programming languages. In *Proceedings OOPSLA '86, ACM SIGPLAN Notices*, volume 21, pages 38–45, November 1986.
- [14] Squeak. <http://www.squeak.org/>.
- [15] L. Tesler. The Smalltalk environment. *Byte*, 6(8), August 1981.
- [16] Cincom VisualWorks Smalltalk. <http://www.cincom.com/scripts/smalltalk.dll/>.
- [17] Roel Wuyts. Declarative reasoning about the structure object-oriented systems. In *Proceedings of the TOOLS USA '98 Conference*, pages 112–124. IEEE Computer Society Press, 1998.
- [18] Roel Wuyts and Stéphane Ducasse. Unanticipated integration of development tools using the classification model. In *ESUG'2003 Academic Conference. Publication Pending*, 2003.