

# Visualising Class Cohesion with Virtual Worlds

Neville Churcher

Warwick Irwin

Ron Kriz

Software Visualisation Group, Department of Computer Science,  
University of Canterbury, Private Bag 4800,  
Christchurch, New Zealand  
E-mail: {wal,neville}@cosc.canterbury.ac.nz

VT-CAVE,  
Virginia Tech  
Blacksburg VA24061  
E-mail: rkriz@vt.edu

## Abstract

High cohesion, or module strength, indicates that a system has been well partitioned into components which have strong internal relationships. An understanding of cohesion is an important factor in software design. However, cohesion is difficult to define precisely and to quantify, particularly for object oriented (OO) software, and attempts to develop metrics have had limited success. We advocate the use of visualisation techniques to provide a richer view of cohesion than is possible with simple numeric values. Our approach, based on visualising the raw data, side-steps many of the difficulties associated with the formulation of precise definitions of cohesion phenomena, avoids the loss of information associated with existing aggregate cohesion metrics and supports exploratory analysis. The pipeline-based visualisations of class cohesion described in this paper use the ANGLE 3D graph layout engine and XSLT transformations both to select the ingredients for visualisations and to determine their presentation details. We discuss our experiences with the use of virtual worlds as a presentation medium both on the desktop and in immersive environments and report early results from ongoing empirical work.

*Keywords:* software visualisation, cohesion, VRML, virtual reality, software metrics, software engineering

## 1 Introduction

Decomposition is one of the oldest and most fundamental techniques employed in software engineering. It results in a “good” system design which consists of relatively independent components which interact only in order to achieve system goals.

During analysis we model a real-world problem by breaking it up into more manageable sub-problems until these are small enough to be represented accurately in our preferred modelling notation. The problem is then defined by the representations of the sub-problems and the relationships between them. Similarly, our design and subsequent implementation will also consist of a number of related parts.

Although the terms used to denote the parts (sub-systems, modules, classes, components, ...) and relationships have changed over the years the benefits of a modular system have endured. The basic idea is simple: put things that belong together in one place. This leads to independent ‘pluggable’ modules which

are easy to develop, maintain and replace and which present simple interfaces to clients.

As always, there is a catch. There is no single correct way to perform decomposition and different choices can lead to different designs with very different properties. For example, techniques based on decomposition with respect to data flow tend to produce different results to those of techniques based on decomposition with respect to control.

Concepts such as information hiding (Parnas 1972) help address the problem of deciding what should be in a module. Good choices lead not only to ‘better’ modules but also to simpler relationships between them.

A related problem is assessing the degree to which each module’s contents ‘belong together’ and the relative independence of modules. Cohesion is a measure of internal module strength while coupling describes the strength of inter-module relationships. They are not independent—coupling is lowered (and cohesion raised) when communicating activities are located in the same module. Cohesion and coupling are an ingredient in many design techniques (Yourdon & Constantine 1979, Page-Jones 1988).

The advent of object-oriented software development has brought many benefits but has added further complexity to the concepts of cohesion and coupling by increasing the number of available component and relationship types. For example, in conventional procedural languages such as Pascal or C the dominant relationship is *procedure invokes procedure*. In OO languages components include methods, classes and packages while relationships include association, inheritance and containment in addition to invocation.

Coupling and cohesion for OO systems have been discussed extensively (Berard 1993, for example) and underlie design heuristics such as the Law of Demeter (Lieberherr & Holland 1989).

Only for very small systems is a subjective assessment of the levels of cohesion possible and sufficient. Measures which are more objective and amenable to automated collection are required for systems of realistic size. Cohesion is multi-faceted, involving complex interactions between many components, and it is not possible to represent it adequately with only scalar quantitative metrics without significant loss of information.

In our experience, software visualisation excels in problem domains having many variables, complex interactions and large data volumes. Software visualisation involves applying information visualisation principles to the software engineering domain. In previous work (Churcher, Keown & Irwin 1999, Hartley, Churcher & Albertson 2000, Churcher & Creek 2001, Irwin & Churcher 2002) we have explored the use of 3D virtual worlds to deliver visualisations to the desktop via VRML (Carey & Bell 1997). We are primarily

interested in “just in time” visualisations, requested and configured under the control of the user, rather than polished visualisations requiring time and expertise not normally available.

In this paper we develop 3D graph-based visualisations which expose cohesion. The bindings which give rise to cohesion are represented by edges and aspects of cohesion may be inferred directly from the graph layout. The ANGLE layout engine is used and XSLT transformations (Kay 2000) describe the mappings from the software system to the graph model and from the graph model to the final visualisation.

The remainder of this paper is structured as follows. In the next section we discuss previous cohesion metrics. Our visualisation strategy is introduced in section 3 and its application to the visualisation of class cohesion is described in section 4. Some results are presented in section 5 and in section 6 we discuss the use of immersive environments to go beyond desktop presentation. Finally, our conclusions and future directions are presented in section 7.

## 2 Measuring coupling and cohesion

Software metrics have been developed for measuring many aspects of software artifacts and processes (Henderson-Sellers 1996, Fenton & Pfleeger 1997) including coupling and cohesion.

Measures of coupling typically involve inter-module quantities such as fan-in and fan-out while measures of cohesion are generally derived from intra-module quantities. Ideally, metrics should be developed in the context of a framework which embodies the conceptual model of the paradigm (Churcher & Shepperd 1995*b*, Briand, Daly & Wüst 1998, for example) but this has proved to be an elusive goal. A comprehensive theory of software is still lacking.

Yourdon and Constantine identified several kinds of cohesion and coupling and used numeric values to indicate the desirability of each (Yourdon & Constantine 1979). Strictly speaking, this scale is ordinal even though the values are numeric. Cohesion measures based on program slices have been developed (Ott & Bieman 1998, for example).

The most commonly used suite of OO metrics, proposed by Chidamber and Kemerer, includes three relevant metrics: RFC (method invocation), CBO (data and object use) and LCOM (cohesion). Precise definition and interpretation of these metrics has proved not to be straightforward (Churcher & Shepperd 1995*a*). LCOM has been particularly problematic (Briand et al. 1998, Henderson-Sellers, Constantine & Graham 1996, Henderson-Sellers 1996) and a number of alternatives have been proposed.

These include LCOM\* which ranges from 0 for perfect cohesion (each method uses all properties) to 1 for the case where each method accesses only one property (Henderson-Sellers et al. 1996).

An alternative approach was suggested by Hutchens and Basili (Hutchens & Basili 1985). They were interested in the large-scale morphology of systems written in conventional languages. By considering data bindings (triples  $(f_i, d_k, f_j)$  where data item  $d_k$  is accessible to both functions  $f_i$  and  $f_j$ ) they were able to define a dissimilarity measure between functions.

This allowed the application of statistical clustering techniques (Everitt 1993, for example). The resulting dendrograms are trees showing how the most closely related (i.e. least dissimilar) functions form clusters which in turn interact with other clusters to form larger clusters until finally the entire system becomes a single cluster. Measures of cluster strength and distinctness are available.

The initial clusters, near the leaves of the dendrogram, suggest features such as ADTs where a group of functions interacts strongly with core data structures. Towards the middle of the dendrogram the coarser system structure is revealed while at the root all systems appear as a single blob.

Hutchens and Basili identified three characteristic ‘fingerprints’ from their experimental work.

**planetary** several subsystems ‘orbiting’ a central core.

**black hole** dominant cluster gradually accretes elements.

**gas cloud** shows no tendency to cluster.

Similar ideas have been explored for object oriented systems. Durnota and Mingins proposed a coherence metric from which dendrograms could be constructed (Durnota & Mingins 1993). Another application is the analysis of legacy code in order to identify structures which might become classes when migration to an object oriented version occurs (van Deursen & Kuipers 1999). Although these ideas are intuitively appealing they appear not to have been widely adopted.

Defining and measuring class cohesion involves multiple aspects. The internal cohesion of individual methods is analogous to the cohesion of functions in conventional languages and affects the coupling via invocation of other methods. In many ways, a class is a conventional program in microcosm. The use of class properties by methods is analogous to the access of global data by functions in conventional languages. Although encapsulation and inheritance ensure that there is no truly global data, the basic concept is the same. Indirect coupling arises between methods which access the same properties. Additionally, the invocation of methods within the same class lowers the coupling to other classes.

Our motivation stems from Hutchens and Basili’s approach of looking at a system as a whole and interpreting the complex picture which results. The alternative is to define specific metrics, with corresponding information loss, and attempt to find empirical evidence for their relevance. Rather than using statistical clustering, we apply visualisation techniques directly to the raw data in order to highlight the cohesion of classes. If we can develop OO counterparts of fingerprints then these will be particularly useful to software engineers.

## 3 Visualisation with Angle

The conventional visualisation pipeline involves several stages (Schroeder, Martin & Lorenzen 1998). Data capture and transformation are followed by the computation of geometry and finally the rendering of the images. In our previous work we have shown how this pipeline approach may be extended in the software visualisation domain (Churcher et al. 1999, Hartley et al. 2000, Churcher & Creek 2001, Irwin & Churcher 2002). In particular, we have explored the potential uses of XML (Martin, Birkbeck, Kay, Loesgen, Pinnock, Livingstone, P.Stark, Williams, Anderson, Mohr, Baliles, Peat & Ozu 2000) both for representing the products at each stage of the pipeline and for specifying the sequence of transformations which occur.

The pipeline used in our current work is shown in Figure 1. A central component is the ANGLE graph layout engine (Churcher & Creek 2001) which provides reliably “nice” layout of 3D graphs using a force-based approach. Nodes repel each other as if they

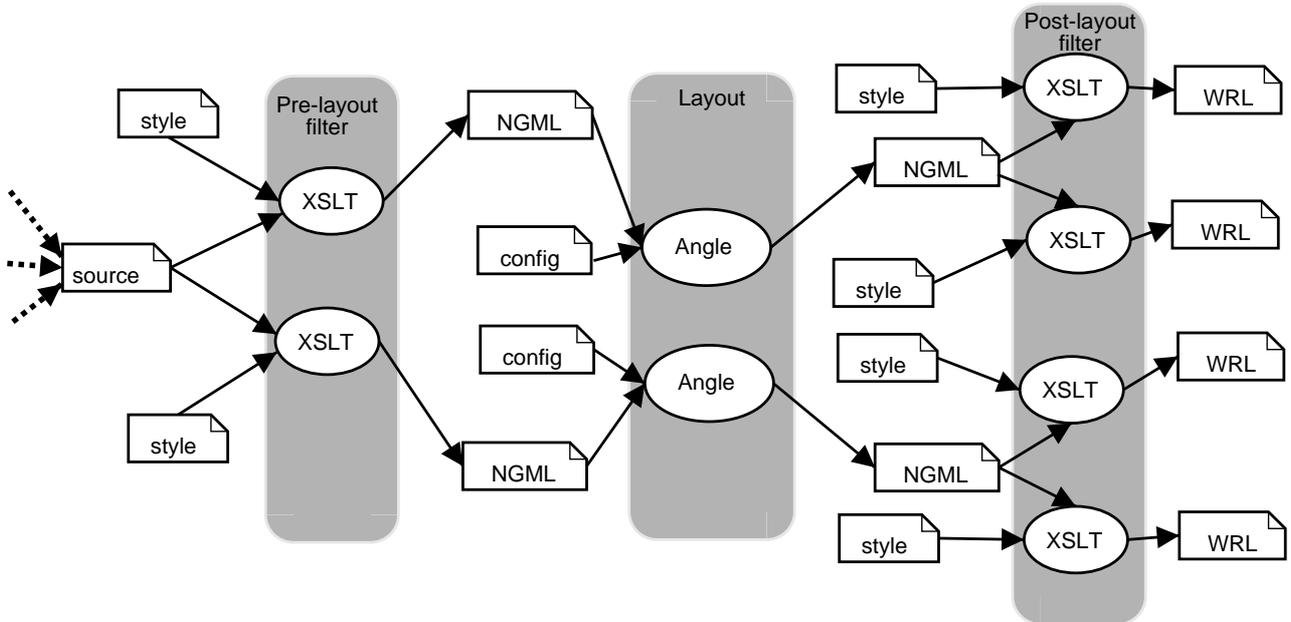


Figure 1: ANGLE in the visualisation pipeline

were charged particles while edges act like springs which pull connected nodes towards each other. Parameters selected by the user include the functional form and strength of the forces which, in turn, influence the equilibrium state which constitutes a “good” layout (Di Battista, Eades, Tamassia & Tollis 1999).

The current version of ANGLE is strongly based on XML technology. Its input and output files are in attributed NGML format—an XML vocabulary in which the graph structure is represented by node and edge elements—with the output files augmented with layout information. More expressive XML-based vocabularies for graph representation and exchange are being developed (Winter 2001, Brandes, Eiglsperger, Herman, Himsolt & Marshall 2001). However, NGML is currently sufficient for our needs in this domain and we can make use of XSLT transformations in order to translate to and from such alternative representations. We believe that the separation of a model from its many possible views is at least as important in visualisation as it is in other areas of software engineering.

Designing a visualisation is not an exact science and there are generally competing solutions. The key activity is the establishment of mappings between features in the system and elements of the display presented to the user.

Software visualisation requires the construction of artificial geometry (nodes and edges in our case) based on semantics extracted from the source data by appropriate mappings. ANGLE provides a flexible and extensible approach to specifying and applying such mappings. A given data source may be passed through any number of pre-filters, each of which is an XSLT transformation determining which elements and relationships will be included in the structure and semantics of the graph which underlies the visualisation. Similarly, post-filters transform the graph layout into the desired form for presentation. Figure 1 shows a single data source with 2 pre-filter and 4 post-filter transformations (i.e. two different graphs, each presented in 2 different ways).

#### 4 Visualising cohesion

It is perhaps not surprising that such a multi-faceted concept as cohesion has resisted representation by a single numeric metric. Loss of information is inevitable when multi-faceted data is aggregated into a single value.

We contend that it is preferable to deliver a holistic “big picture” representation of the class structure which can be readily interpreted by software engineers. Since cohesion arises from the bindings among the class components we should try to visualise these explicitly and use them to assess cohesion levels.

Of course, quantitative metrics may be included in such views. The use of 3D virtual worlds enables more variables to be included in our visualisations. Thus, the size, shape or colour of symbols can be used to represent the values of numeric metrics.

The basic steps in visualisation construction are:

1. Form graph model by extracting node and edge data from source data
2. Compute 3D graph layout
3. Use semantic information to add presentation detail.

These steps correspond respectively to the pre-layout filter, layout and post-layout filter stages in Figure 1.

The source data for the pipeline is represented by the symbol at the left of the figure. In our work, this typically arises from the use of yacc-based static analysis tools (Irwin & Churcher 2002). The details of the data format are defined by an appropriate Document Type Definition (DTD). For example, the data used to produce the worlds shown in this paper is expressed in terms of classes, which have properties (attributes, data members, fields) and methods (function members, services), which in turn use properties. Figure 2 shows an example of this format.

XSLT stylesheets are used in the transformation from the specific source data format to the more general attributed NGML input for ANGLE. One considerable advantage of this approach is that a single source data set may be used unchanged in multiple visualisation pipelines, each of which will emphasise different aspects of the data.

```

...
<class>
  <name>Customer</name>
  <property name="address"/>
  <method signature="getAddress" accessMode="public">
    <name>getAddress</name>
    <use property="address"/>
  </method>
..
</class>

```

Figure 2: Class description format

	source	ngml
$\gamma_c$		class $\mapsto$ node
$\gamma_p$		property $\mapsto$ node
$\gamma_m$		method $\mapsto$ node
$\gamma_{cp}$		class-has-property $\mapsto$ edge
$\gamma_{cm}$		class-has-method $\mapsto$ edge
$\gamma_{mp}$		method-uses-property $\mapsto$ edge

Table 1: Pre-filter mappings

Table 1 shows some typical mappings between the elements of the source data and those of the corresponding attributed NGML file. The nodes and edges resulting from the selected mappings will participate in the layout computation and be available to subsequent pipeline stages. The worlds shown in this paper represent uses as edges ( $\gamma_{mp}$ ) between method ( $\gamma_m$ ) and property ( $\gamma_p$ ) nodes. Some include nodes representing the class itself ( $\gamma_c$ ) and edges connecting it to properties ( $\gamma_{cp}$ ) and methods ( $\gamma_{cm}$ ).

Applying the XSLT file corresponding to some pre-filter mapping set  $\Gamma$  yields an NGML description of the graph, together with various other data which have been attached as attributes of node and edge elements. Such attributes are ignored by ANGLE and simply pass out the other end where they are available to transformations further down the pipeline. In this example, three separate kinds of node and three kinds of edge are present.

ANGLE then generates a 3D layout consistent with the configuration data (algorithm, terminator and iteration parameters) provided. The format of the output is essentially the same as that of Figure 3 with the addition of coordinate elements for each node.

A given output NGML file may then be further transformed in various ways, as specified by XSLT files corresponding to post-filter mappings,  $\Delta$ , to produce VRML worlds or other formats. For example, Table 2 contains post-filter mappings used to produce worlds shown in this paper. Note that nodes and edges which participate in the layout computation will not necessarily appear in the presentation unless corresponding mappings are used. Although not shown here, it is straightforward to add further mappings (such as *accessmode(public)*  $\mapsto$  *transparency 0.5*). In this way, changes to presentation details, such as colour and shape of symbols, are neatly separated from the structural data about the underlying graph.

	ngml	VRML
$\delta_c$	node (classNode)	$\mapsto$ purple cylinder
$\delta_p$	node (propertyNode)	$\mapsto$ green cube
$\delta_m$	node (methodNode)	$\mapsto$ red sphere
$\delta_{cp}$	propertyEdge	$\mapsto$ yellow cylinder
$\delta_{cm}$	methodEdge	$\mapsto$ purple cylinder
$\delta_{mp}$	useEdge	$\mapsto$ cyan cylinder

Table 2: Post-filter mappings

```

<graphset>
  <graph>
    <title>The Customer Class</title>
    <nodes>
      <node nodetype="classNode">
        <name>Customer</name>
      </node>
      <node nodetype="propertyNode"
        accessmode="private">
        <name>address</name>
      </node>
      <node nodetype="methodNode"
        nid="getAddress"
        accessmode="public">
        <name>getAddress</name>
      </node>
    ...
    </nodes>
    <edges>
      <edge edgetype="propertyEdge">
        <from>Customer</from>
        <to>address</to>
      </edge>
      <edge edgetype="methodEdge">
        <from>Customer</from>
        <to>getAddress</to>
      </edge>
      <edge edgetype="useEdge"
        usetype="unspecified">
        <from>getAddress</from>
        <to>address</to>
      </edge>
    ...
    </edges>
  </graph>
</graphset>

```

Figure 3: Attributed NGML

$\Gamma_1$	$\{\gamma_c, \gamma_p, \gamma_m, \gamma_{cp}, \gamma_{cm}, \gamma_{mp}\}$
$\Gamma_2$	$\{\gamma_c, \gamma_p, \gamma_m, \gamma_{cp}, \gamma_{mp}\}$
$\Gamma_3$	$\{\gamma_p, \gamma_m, \gamma_{mp}\}$
$\Delta_1$	$\{\delta_c, \delta_p, \delta_m, \delta_{cp}, \delta_{cm}, \delta_{mp}\}$
$\Delta_2$	$\{\delta_c, \delta_p, \delta_m, \delta_{cp}, \delta_{mp}\}$
$\Delta_3$	$\{\delta_p, \delta_m, \delta_{mp}\}$

Table 3: Sets of mappings used in this paper

For convenience, we define some frequently-used sets of mappings (Table 3) used in the pre- and post-filter transformations reported in this paper. Clearly, many other individual mappings and combinations are possible. These are straightforward to generate.

## 5 Results

In this section we begin by examining visualisations of some simple systems in order to identify characteristic features and illustrate the effects of different mappings. A more realistic example will then be examined.

Figure 4 shows four low cohesion situations for a class (represented by the purple cylinder) with three properties (green cubes) and three methods (red spheres). Edges represent the relationships *class-has-property* ( $\delta_{cp}$ ), *class-has-method* ( $\delta_{cm}$ ) and *method-uses-property* ( $\delta_{mp}$ ). Each figure is a snapshot of a VRML world and is labelled with the corresponding pre- and post-filter mappings.

Figure 4(a) shows the case where no method uses any property. The graph is symmetric, with the class

node in the centre, and is characterised by a radial fingerprint reminiscent of a *kina* (sea urchin). Properties which are never used are likely to be candidates for removal during design review. These will not be highlighted by metrics such as LCOM and it is valuable to be able to see these readily.

Now let us consider the effect of adding *method-uses-property* edges to our visualisations. Figures 4(b), 4(c) and 4(d) show three visualisations of the case where each method uses one property.

In Figure 4(b) all edges both contribute to the layout and are represented explicitly since  $\gamma_{cm} \in \Gamma_1$  while in Figure 4(c) all contribute to the layout but *class-has-method* edges are not present in the world since  $\delta_{cm} \notin \Delta_2$ .

In Figure 4(d) the edges representing *class-has-method* relationships neither contribute to the graph layout nor are they present in the world since  $\gamma_{cm} \notin \Gamma_2$ .

The effects of more than one method accessing a property and more than one property being accessed by a method are shown in Figure 5. In this example all methods access at least one property.

Figures 5(a) and 5(b) both include  $\gamma_{cm}$  but the corresponding edges are not rendered in Figure 5(b) since  $\delta_{cp} \notin \Delta_2$ . The 3-fold symmetry arising from the three methods each accessing the same property and the two-fold symmetry arising from the use of two properties by one method are clearly visible in each case. The presence of two clusters of accesses is highlighted in Figure 5(c) where the  $\Delta_3$  post-filter has suppressed the class node and all non-use edges.

Figure 5(d) results from transformations which exclude  $\gamma_{cm}$ —this generally leads to more “open” layouts.

At the opposite end of the class cohesion spectrum lies perfect cohesion, where each method uses each property. Figure 6 shows 4 presentations of this situation for a class with 6 properties and 6 methods. In Figure 6(a)  $\Gamma_1$  has included in the layout the class node (at the centre of the Figure) and all associated edges. The same pre-filter has been used in Figure 6(b), so the overall shape is the same, but  $\Delta_3$  has suppressed the class node and non-use edges. This gives rise to a more un-cluttered graph since the shape is the same but fewer edges need be displayed. When  $\gamma_{cm}$  is removed (Figure 6(c)) the class node (at left) moves from the centre but the overall symmetry is maintained. In the perfect cohesion case, symmetry is also maintained if the non-use edges are omitted from the layout (Figure 6(d)) since “everything is connected to everything else”.

Perfect cohesion results in highly symmetric polyhedral structures characterised by a “gas cloud” fingerprint.

One of the primary aims of cohesion and coupling measures is to guide refactoring activities. In OO this includes activities such as splitting a class into two or more classes, combining classes (the reverse of splitting) and moving properties or methods from one class to another. Figure 7 shows a potentially separable class which has 5 properties and 5 methods. One pair of properties is used by each of a pair of methods while the other three properties are accessed by each of the remaining three methods.

Two clusters are visible—each corresponding to a candidate class after refactoring. In general, separable classes are suggestive of the “planetary system” fingerprint. The class node connects the clusters in Figures 7(a), 7(b) and 7(c).

However, if the pre-filter mappings for a separable class do not include  $\gamma_c$ ,  $\gamma_{cp}$  and  $\gamma_{cm}$  then more than one connected component will be present. This situation is not generally welcome in force-based layout as the connected components will separate dramatically

under the effect of the repulsive force. It is possible to modify the functional form of the forces to keep separate connected components “nearby” but other aesthetic issues, such as distortion of the components, then arise. A reasonable compromise is illustrated in Figure 7(d) where the mappings used are such that the class node and its associated edges participate in the layout but are omitted in the visualisation.

The examples we have given thus far have been somewhat artificial. Classes encountered in the wild are likely to be larger and to exhibit more complex interactions between their properties and methods. Figure 8 shows several visualisations of a class selected from the results of our on-going study of Java software.

Several features are evident. The overall shape is somewhat planetary and open. A number of *kina* fingerprints indicate the presence of several methods which use only a single property. These are connected by a small number of common methods (such as those enclosed by the ellipse in Figure 8(c)) which use multiple properties.

Cohesion is far from “perfect”, since the structure is rather open, but it is not clear whether separation is warranted. Design review should concentrate on those *kina*, such as the pair at the upper- and lower-left of Figure 8(c), which are weakly connected (i.e. have few shared methods).

One method (rectangle in Figure 8(c)) uses only one property and is the only user of that property. The radial strand highlights this feature. This is also clearly visible in Figure 8(a) where (since  $\delta_{cm} \notin \Delta_3$  and  $\delta_{cm} \notin \Delta_3$ ) it appears as a separate cluster. A separate cluster also appears, for the same reason, in Figure 8(d), but is a long way away from the main cluster since non-use edges do not contribute to the layout. However, in Figure 8(b) it is much harder to identify (at the right hand edge) because of the larger number of edges present. Such a feature indicates a possible design anomaly or that an unseen ancestor class contains further uses.

Clearly, different mappings emphasize different features and we do not expect any particular set to be preferable in all situations. For example, the inclusion of  $\gamma_{cm}$  leads to more spherically symmetric layouts (Figures 8(a) and 8(b)) which can make it more difficult to spot weakly connected clusters. However, the flexibility of our approach allows users to develop rapidly custom visualisations and to compare the differences between alternative representations.

## 6 Beyond the desktop

In practice, classes can be very large and complex: one would expect frequently to encounter classes with tens of properties and hundreds of methods, particularly when inheritance is included. Another important aspect is the interaction of software engineers—refactoring is likely to be based on the outcome of collaborative discussions of artifacts such as our visualisations.

When worlds become very large they can become difficult to navigate and interpret with desktop facilities (typically a web browser with VRML plug-in). Similarly, as the number of collaborators increases it becomes more difficult for them to gather around a VDU. Immersive virtual reality environments such as CAVEs utilise stereo projections onto multiple surfaces to give users the impression of being inside the artifact being visualised.

Figure 9(a) shows the basic structure of the CAVE at Virginia Tech (<http://www.cave.vt.edu>). Figure 9(b) shows a snapshot from the CAVE console while the class of Figure 8 is being displayed. The

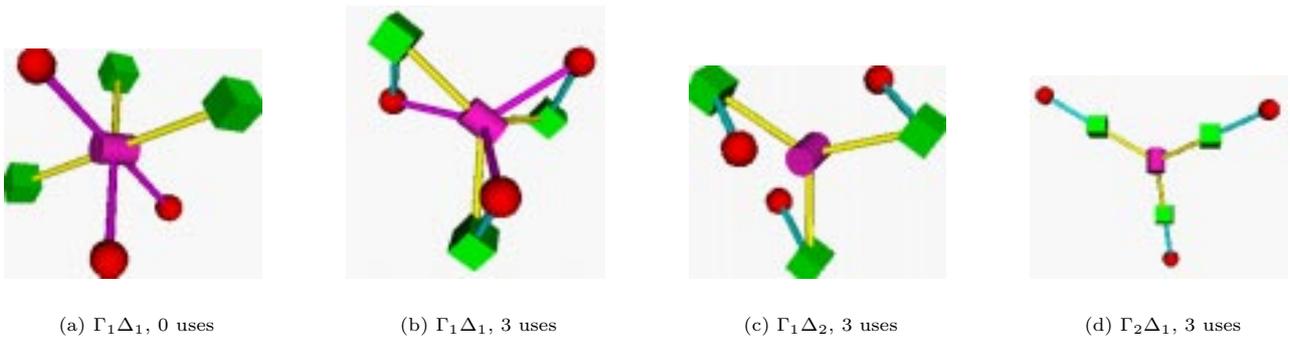


Figure 4: Low class cohesion

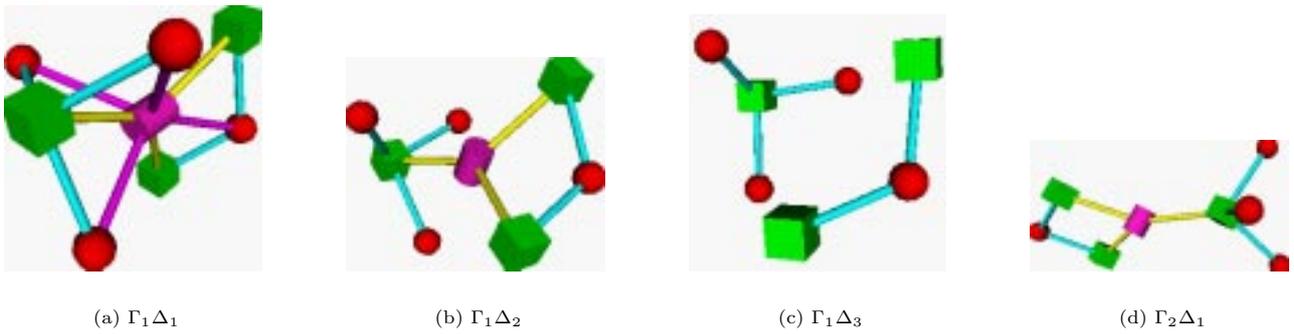


Figure 5: Multiple use participation

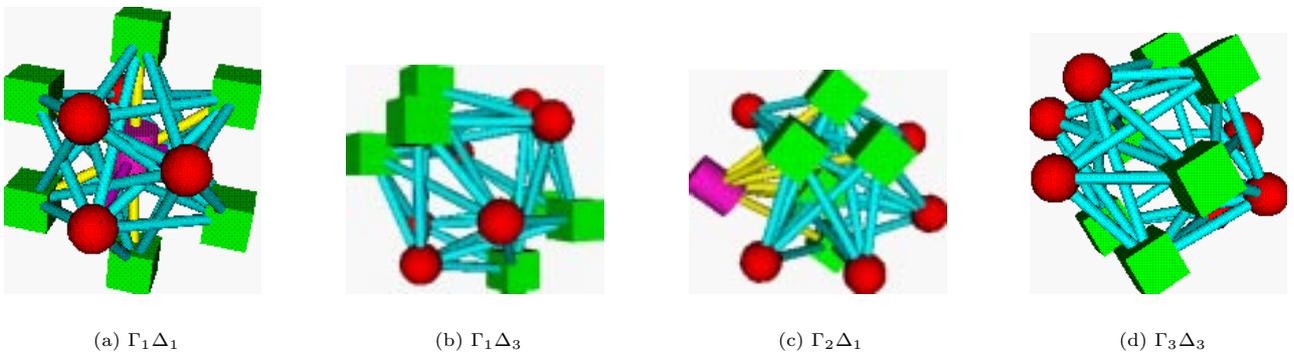


Figure 6: Perfect cohesion (6 properties, 6 methods)

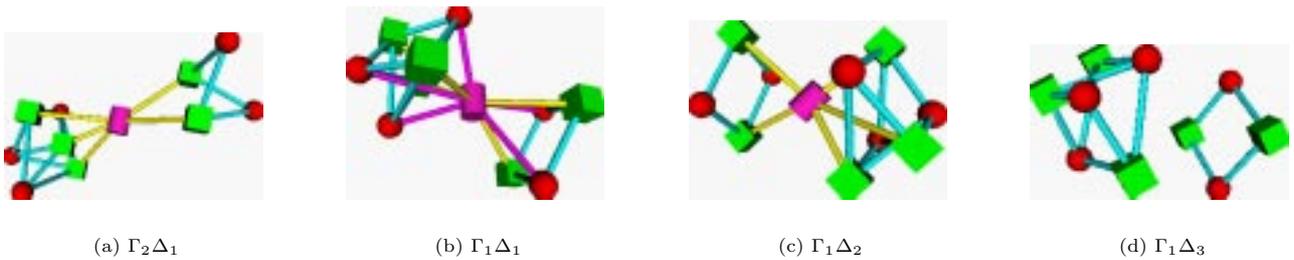
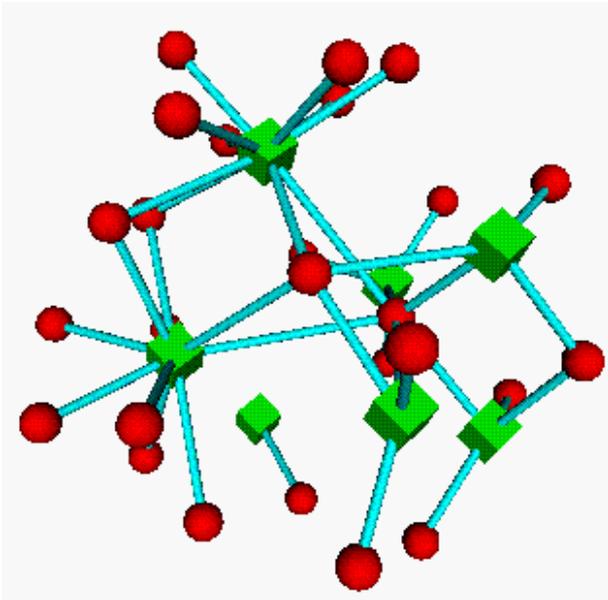
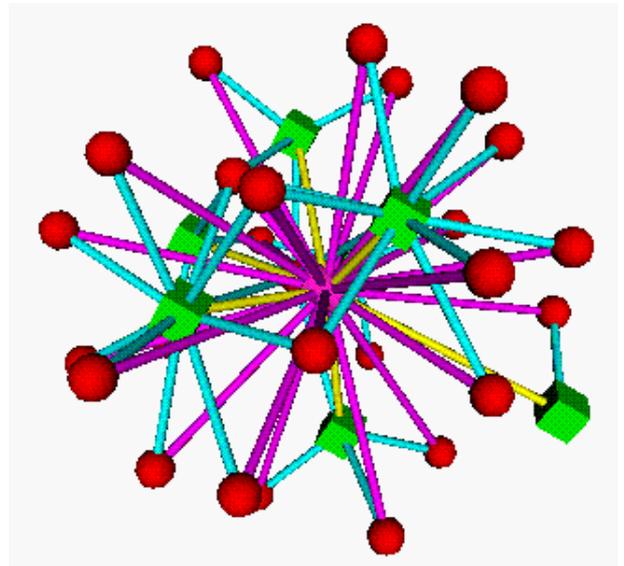


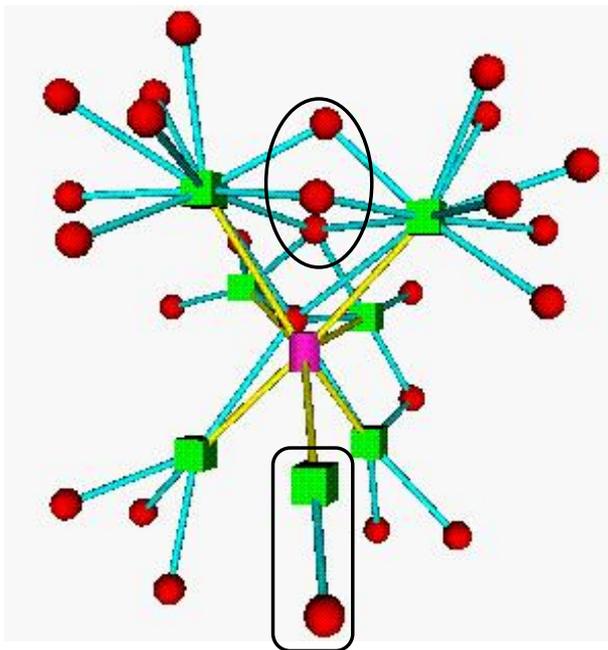
Figure 7: Separable class



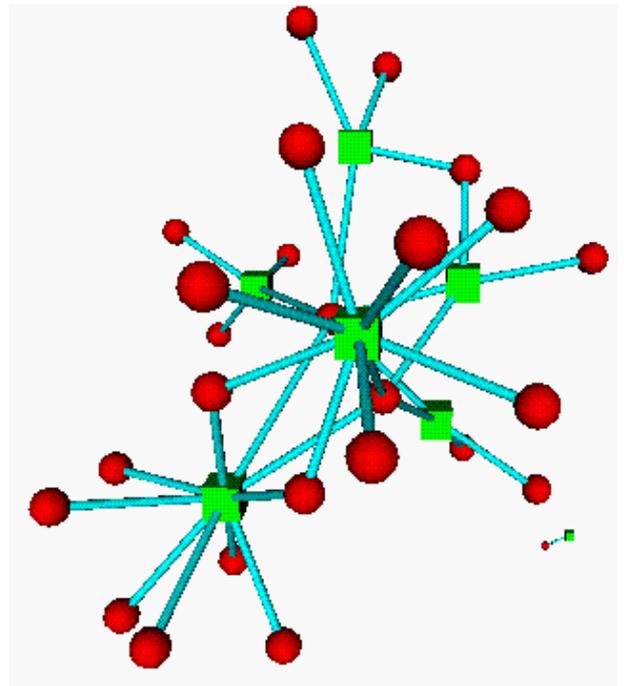
(a)  $\Gamma_1\Delta_3$



(b)  $\Gamma_1\Delta_1$

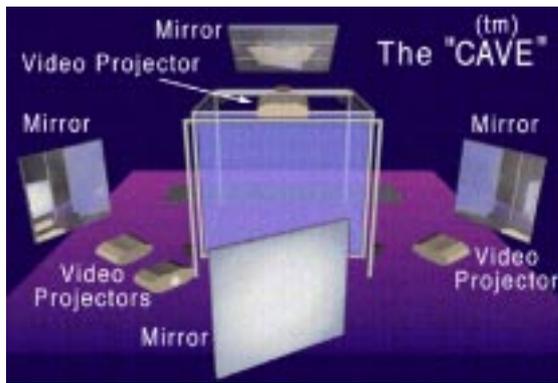


(c)  $\Gamma_2\Delta_1$

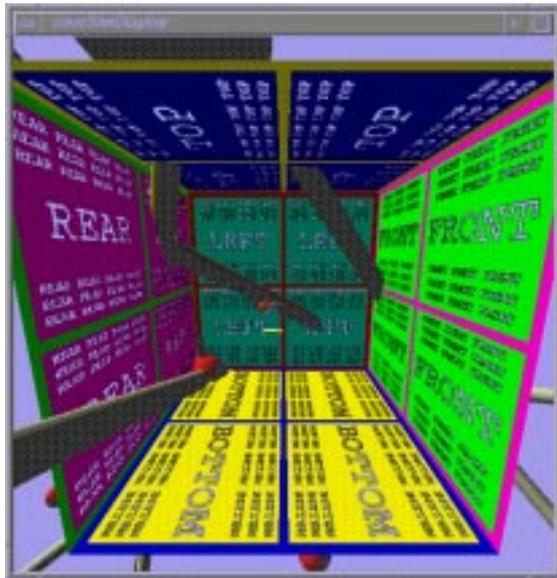


(d)  $\Gamma_3\Delta_1$

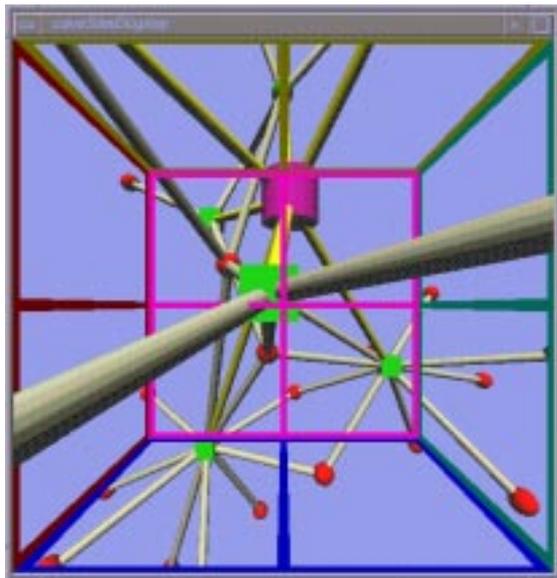
Figure 8: A “real” class with 7 properties and 26 methods



(a) Basic structure



(b) CAVE console snapshot showing walls



(c) CAVE console snapshot

Figure 9: The VT CAVE

labelled walls correspond to the projection surfaces indicate how much of the world is currently “inside” the CAVE. Figure 9(c) is another snapshot from the CAVE console and illustrates how the world extends beyond the CAVE walls, which are indicated by the grid lines. CAVE users do not see the walls and grid.

Experiencing these worlds in the CAVE is an impressive experience and leaves us in no doubt that comprehension of worlds such as these is greatly enhanced. In particular, the ability to interact with “life-size” artifacts is a significant advantage.

Although CAVE technology is currently rather expensive, lower-cost and portable alternatives and augmented reality systems are becoming more affordable and may be available routinely on the desktop in a few years. Our visualisations have also been used with the MagicBook (Billinghurst, Kato & Poupyrev 2001).

## 7 Conclusions and future work

Our initial results are encouraging. Our pipeline approach allows us to determine which components and relationships from a data source will be included in a graph structure, to generate a nice 3D layout of the corresponding graph and to control the presentation of the resulting visualisation using semantics derived from the original source data. XSLT transformations combined with ANGLE permit software engineers to explore not only the content of existing visualisations but also the creation of new kinds.

Our approach represents explicitly the bindings from which cohesion arises. A number of characteristic features have been identified and related to specific cohesion situations. We believe that our visualisations of class cohesion complement the conventional metrics-based approach. Data from our on-going empirical work allow us to make more confident predictions about which classes are most likely to exhibit cohesion anomalies.

Only a small set of mappings has been discussed in this paper. We will explore others (such as *method-invokes-method* edges) in order to learn more about what is necessary for a “realistic” visualisation and to build up our library of transformations.

Real systems exhibit many complicating factors. Constructors and methods such as Java’s `toString` tend to access many properties while setters and getters (mutators and accessors) tend to access only one. Inheritance and overloading are important. By providing a range of mappings we enable the user to explore the effects of such factors—we do not expect to find a single form which will be optimal in all situations.

Formal user trials are required in order to provide quantitative support for our anecdotal evidence that comprehension and collaboration is enhanced by immersive environments.

User studies are needed in order to clarify which visualisations are most useful to software engineers in particular contexts. It may also be possible to identify potential cohesion problems from the presence of characteristic features and to highlight these for resolution by the software engineers. At this point, formal user trials may be conducted in order to provide quantitative support for our anecdotal evidence that comprehension and collaboration is enhanced by both immersive and non-immersive environments.

Virtual and augmented reality environments encourage and support the collaborative analysis of large complex systems and we anticipate their increased adoption as part of the software engineering tool set.

## References

- Berard, E. (1993), *Essays on object-oriented software engineering*, Vol. 1, Prentice Hall.
- Billinghurst, M., Kato, H. & Poupyrev, I. (2001), 'The MagicBook—moving seamlessly between reality and virtuality', *IEEE Computer Graphics and Applications* **21**(3), 6–8.
- Brandes, U., Eiglsperger, M., Herman, I., Himsolt, M. & Marshall, M. (2001), Graphml progress report, in Mutzel, Jünger & Leipert (2001), pp. 501–512.
- Briand, L., Daly, J. & Wüst, J. (1998), 'A unified framework for cohesion measurement in object-oriented systems', *Empirical Software Engineering* **3**(1), 65–117.
- Carey, R. & Bell, G. (1997), *The Annotated VRML 2.0 Reference manual*, Addison-Wesley.
- Churcher, N. & Creek, A. (2001), Building virtual worlds with the big-bang model, in P. Eades & T. Pattison, eds, 'Information Visualisation 2001', Vol. 9 of *Conferences in Research and Practice in Information Technology*, ACS, Sydney, Australia.
- Churcher, N., Keown, L. & Irwin, W. (1999), Virtual worlds for software visualisation, in A. Quigley, ed., 'SoftVis99 Software Visualisation Workshop', University of Technology, Sydney, Australia, pp. 9–16.
- Churcher, N. & Shepperd, M. (1995a), 'Comment on "a metrics suite for object oriented design"', *IEEE Transactions on Software Engineering* **21**(3), 263–265.
- Churcher, N. & Shepperd, M. (1995b), 'Towards a conceptual framework for OO software metrics', *ACM SIGSOFT Software Engineering Notes* **20**(2), 69–75.
- Di Battista, G., Eades, P., Tamassia, R. & Tollis, I. G. (1999), *Graph Drawing: Algorithms for the Visualization of Graphs*, Prentice Hall.
- Durnota, B. & Mingins, C. (1993), Tree-based coherence metrics in object-oriented design, in C. Mingins, B. Haebich, J. Potter & B. Meyer, eds, 'TOOLS12&9', Prentice Hall, Sydney, Australia, pp. 489–504.
- Everitt, B. (1993), *Cluster Analysis*, Edward Arnold.
- Fenton, N. & Pfleeger, S. L. (1997), *Software Metrics: A Rigorous & Practical Approach*, 2nd edn, International Thompson Computer Press.
- Hartley, D., Churcher, N. & Albertson, G. (2000), Virtual worlds for web site visualisation, in 'Proc APSEC2000, 7th Asia Pacific Software Engineering Conference', IEEE Press, Singapore, pp. 448–455.
- Henderson-Sellers, B. (1996), *Object-Oriented Metrics: Measures of Complexity*, Prentice Hall.
- Henderson-Sellers, B., Constantine, L. & Graham, I. (1996), 'Coupling and cohesion (towards a valid metrics suite for object-oriented analysis and design)', *Object Oriented Systems* **3**, 143–158.
- Hutchens, D. & Basili, V. (1985), 'System structure analysis: clusterings with data bindings', *IEEE Transactions on Software Engineering* **11**(8), 749–757.
- Irwin, W. & Churcher, N. (2002), XML in the visualisation pipeline, in D. D. Feng, J. Jin, P. Eades & H. Yan, eds, 'Visualisation 2001', Vol. 11 of *Conferences in Research and Practice in Information Technology*, ACS, Sydney, Australia, pp. 59–68. Selected papers from 2001 Pan-Sydney Workshop on Visual Information Processing.
- Kay, M. (2000), *XSLT Programmer's Reference*, Wrox.
- Lieberherr, K. & Holland, I. (1989), 'Assuring good style for object-oriented programs', *IEEE Software* pp. 38–48.
- Martin, D., Birkbeck, M., Kay, M., Loesgen, B., Pinnock, J., Livingstone, S., P. Stark, Williams, K., Anderson, R., Mohr, S., Baliles, D., Peat, B. & Ozu, N. (2000), *Professional XML*, Wrox Press.
- Mutzel, P., Jünger, M. & Leipert, S., eds (2001), *Graph Drawing: 9th International Symposium GD2001*, Vol. 2265 of *Lecture Notes in Computer Science*, Springer-Verlag, Vienna, Austria.
- Ott, L. & Bieman, J. (1998), 'Program slices as an abstraction for cohesion measurement', *Information & Software Technology* **40**(11–12), 691–699. Special Issue: Program Slicing.
- Page-Jones, M. (1988), *The practical guide to structured systems design*, Yourdon Press Computing Series, 2nd edn, Prentice Hall, Englewood Cliffs, N.J.
- Parnas, D. (1972), 'On criteria to be used in decomposing systems into modules', *Communications of the ACM* **14**(1), 221–227.
- Schroeder, W., Martin, K. & Lorenson, B. (1998), *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, 2nd edn, Prentice Hall.
- van Deursen, A. & Kuipers, T. (1999), Identifying objects using cluster and concept analysis, in 'Proceedings of the 21st International Conference on Software Engineering (ICSE'99)', ACM Press, Los Angeles, pp. 246–255.
- Winter, A. (2001), Exchanging graphs with GXL, in Mutzel et al. (2001), pp. 485–500.
- Yourdon, E. & Constantine, L. (1979), *Structured design : fundamentals of a discipline of computer program and systems design*, Prentice Hall, Englewood Cliffs, N.J.