

Visualization and Graphical Layout in Object-Oriented Systems*

Volker Haarslev

Xerox Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA 94304, USA
haarslev@parc.xerox.com

and

University of Hamburg, Computer Science Department
Bodenstedtstr. 16, D-2000 Hamburg 50, F.R. Germany
haarslev@rz.informatik.uni-hamburg.dbp.de

Ralf Möller

University of Hamburg, AI Laboratory
Bodenstedtstr. 16, 2000 Hamburg 50, F.R. Germany
moeller@rz.informatik.uni-hamburg.dbp.de

Abstract

This report describes a new approach to visualizing program systems within the object-oriented paradigm. This approach is based on a \TeX -like notation which has been extended and generalized for specifying graphical layout of arbitrary objects. Our simplest scheme offers specifications similar to \TeX 's box-and-glue metaphor. Size and position of boxes and glue can be specified by constraints. The CLOS meta-level architecture is used to associate visualization and application objects. We propose several useful techniques such as indirect values, slot and method demons, and instance-specific meta-objects. Our techniques require no modifications to the systems which are selected for visualization. We demonstrate the feasibility of our approach using application domains such as CLOS debugging and constraint systems.

* This report combines and extends the following two papers:

A Declarative Formalism for Specifying Graphical Layout, published in: Proceedings, *1990 IEEE Workshop on Visual Languages*, Skokie/IL, Oct. 4-6, 1990, IEEE Computer Society Press, 1990.

A Framework for Visualizing Object-Oriented Systems, published in: Proceedings, ECOOP/OOPSLA'90, *European Conference on Object-Oriented Programming and Object Oriented Programming: Systems, Languages and Applications*, Oct. 21-25, 1990, Ottawa/Canada, ACM Sigplan Notices, 1990.

Table of Contents

1	Introduction	1
2	Layout Specifications	2
2.1	Vertical and Horizontal Boxes	2
2.2	Filler Specification	3
2.3	Local Variables in Layout Forms	4
3	Layout Protocols	4
4	Interaction Objects and Views	6
4.1	Interaction Objects	6
4.2	Referencing View Items	7
4.3	Named References of View Items	9
5	Extended Class Browser and Inspector	11
6	The Box Layout Algorithm Revisited	13
7	Optimizing Graphical Output	15
8	Meta-Level Techniques for Separating Application and Visualization	16
8.1	Indirect Values for Visualization Objects	17
8.2	Slot Demons for Application Objects	18
8.3	Method Demons	19
8.4	Instance-Specific Meta-Objects	20
9	Related Work	22
10	Summary and Future Work	23
	References	24

List of Figures

1	A DAG graph of a standard class hierarchy.	3
2	A general layout specification in combination with a box-style layout. ¹	5
3	Referenced Location of Item B.	7
4	Location of edges defined by node references.	8
5	EBNF Syntax of Layout Specifications.	10
6	A DAG graph of a class hierarchy with local class slots.	11
7	CLOS class inspector window (default size).	12
8	CLOS class inspector window (vertically and horizontally enlarged).	12
9	Schematic representation of nested boxes. Arrows represent fillers.	13
10	Water jugs modeling fillers.	14
11	The right window shows the state after a black circle has been moved.	15
12	The upper and lower bounds are indicated by shaded rectangles [Winston & Horn 89]. The gauges show the estimation interval from 0 (bottom) to 1 (top).	16
13	Outline of a wrapper method	19
14	Meta-objects for instances with metaclass <code>extensible-standard-class</code>	20

1 Introduction

Although programming has mostly been done in textual terms users have always had a notion of visualizing their programs. Programs have been entered as lines of text but soon users started to indent their programs and also used comments for separating or emphasizing particular program parts. Tools were developed which pretty-print or format source code. Modern programming environments offer debugging tools such as browsers and inspectors providing users with views of program structure and execution states. But these views display their information more textually than visually (pictorially). A further disadvantage of these environments is their lack of offering program designers adequate tools for visualizing and animating programs, which support both structural and conceptual visualization.

This report discusses within the paradigm of object-oriented programming the use of structural and conceptual visualization techniques. We describe a new approach combining both techniques, which is based upon \TeX -like layout specifications. Furthermore, we discuss the usefulness of meta-level architectures for implementing visualization techniques. We implemented a prototyping environment consisting of a set of extensible fundamental components which offer varying degrees of support. It is implemented in Macintosh Allegro Common Lisp and based upon the PCL implementation of the Common Lisp Object System [Bobrow et al. 88, Keene 89] (CLOS). We applied our visualization techniques to several application domains such as CLOS debugging, graphical editing, constraint systems, line routing, and concurrent logic programming (see also *Möller 90*).

Conceptual visualizations of programs are mostly created by hand. This hand-design is basically caused by the fundamental problem that geometrical and graphical information necessary to create suitable visualizations cannot automatically be derived from corresponding data. The major problem is to define interesting events which should be visualized. But very often interesting events are only indirectly reflected by algorithms. We refer to *Brown 88* for a detailed discussion of these problems.

Structural visualization uses program and data structures to generate relevant geometrical information. An important problem related to structural interpretation is that conceptual information about data can only indirectly be derived (e.g. from naming of identifiers). A very common approach to structural visualization is to guide the visualization process by underlying programming styles or computational models. Many approaches to visualizing imperative systems use flow charts or diagrams. The Transparent Prolog Machine [Eisenstadt & Brayshaw 88] is an example for relational or logic systems. A more radical approach is presented by Pictorial Janus [Kahn & Saraswat 90]. It defines complete visualizations of concurrent logic programs and captures static as well as dynamic information about these programs.

There exist many approaches to visualizing data flow of functional systems, e.g. VIPEX [Haarslev & Möller 90], Pluribus [Wright et al. 85], and Prograph [Matwin & Pietrzykowski 85, Cox et al. 89]. A diagramming approach to tracing object-oriented systems as an extension to a Smalltalk-80 debugger is described in *Cunningham & Beck 86*. GraphTrace [Kleyn & Gingrich 88] is also intended for understanding behavior of objects. It provides graphical traces of program executions. Both approaches are primarily focused on structural visualizations. In contrast to our approach they offer no support for conceptual visualizations.

Besides structural and conceptual visualization techniques it is also important to support flexible schemes for aesthetically laying out combinations of units. With respect to forms-oriented user interfaces allocation of space and position is mostly constrained by the space globally available. Graphical interfaces usually also add local constraints. A typical application is a browser generating net-like representations of rule sets, classes, or objects. The spatial allocation of nodes may depend on adjoining nodes or the topology of edges (e.g. in order to avoid line crossing or long winding paths). This problem is addressed by many constraint-oriented systems. ThingLab I [Borning 81] and II [Maloney et al. 89] are examples for describing layout of graphical objects with constraints. Szekely & Myers 89 also presented a toolkit using constraints and active values. In contrast to constraint-oriented approaches we decided to provide a simpler but more compact and predictable notation for specifying layout. Furthermore, our approach has the advantage that it requires only $2n + \log(n)$ steps, be n the number of boxes. Thus, our algorithm has a computational complexity of $O(n)$ (see Section 6 for details).

The remainder of this report is structured as follows. The next two sections introduce our T_EX-like specifications. Section 4 introduces box items which are suited for more general applications and a mechanism for specifying references between box items. The next section demonstrates an extended CLOS class browser and inspector which serves as an example for the flexibility of our approach. Section 6 explains our box layout algorithm in more detail. Section 7 discusses some related optimization issues. Afterwards we discuss the use of the CLOS meta-object protocol for program visualization and demonstrate some of these considerations using a simple constraint system as example. Section 9 compares our approach with related work. This report concludes with a summary and a discussion of future work.

2 Layout Specifications

We adopted the “box-and-glue” metaphor of T_EX [Knuth 79] for specifying layout of objects. Layouts are composed of a set of rectangular regions, so-called *boxes*. Laying out boxes and positioning objects are associated with corresponding box types. Our system offers a set of predefined layout algorithms and box types. More general box types are discussed in Section 3.

2.1 Vertical and Horizontal Boxes

The fundamental scheme aligns boxes as a list of *horizontal* and *vertical* boxes. This layout technique has been found very useful for standard (forms-oriented) dialog windows (see Section 9 for a discussion of related work). A layout of a dialog window is specified as a combination of boxes with optional size specifications. Boxes may be arbitrarily nested. The size of boxes and the spatial relationship between them is expressed by an amount of *glue* or *filler* describing either a horizontal or vertical distance. Fillers can be specified as *fixed* (e.g. in pixel) or *variable*. Variable fillers depend on the space available to their enclosing box. We distinguish *relative* and *constrained* fillers. A relative filler is expressed as a fixed ratio to the size of its superior box. Constrained fillers can shrink (stretch) to a given lower (upper) limit. Default constraints are zero as lower limit and box size as upper limit. Several fillers as elements of the same box work together like springs. They share the available space and in general every filler claims the same amount of space which is only constrained by its lower and upper limit.

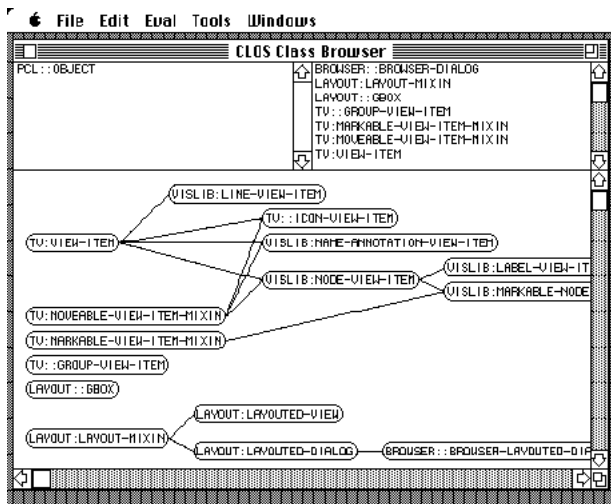


Figure 1: A DAG graph of a standard class hierarchy.

A vertical or horizontal box (`<box-type>` either `:vbox` or `:hbox`) is specified by the lisp form `<box-type> (:width h :height v) box-item-1 ...`. Its box items are laid out vertically resp. horizontally. If the size specification is omitted the width and height of a vertical or horizontal box are set to a filler with default constraints. In general this layout algorithm keeps the size of box elements unchanged. If elements require more space than available to their surrounding box they are allowed to extend beyond their box's boundaries. Boxes are also allowed to overlap one another. But this behavior is not always desired. Therefore, we introduced a *frame box* (`:fbox`) which constrains the size of its element in order to match exactly the frame box size. A frame box contains only one box item: `(:fbox (:width h :height v) box-item)`.

2.2 Filler Specification

The complete form specifying a filler is `(:filler :min m :max n)`, `:min` and `:max` are optional. We defined `:filler` as short-form of `(:filler :min 0 :max <box-size>)`. It is also possible to define the (minimal/maximal) size of a box with respect to its elements. Then, the size of this box is set to the result achieved by laying out its elements and shrinking fillers to their lower limit (see `:as-needed` in Figure 5). Thus, the box has a minimal size satisfying all lower bound constraints.

Figure 1 shows a dialog window for a simple CLOS browser which displays a class hierarchy. The right table contains all direct subclasses of the class listed in the left table. The tables can be replaced by direct super resp. subclasses, scrolled, and shifted to focus on "interesting" classes. The lower part of the window displays a graph of the selected class hierarchy. The dialog window results from the following (schematic) layout specification.

```
(let ((left-table (make-dialog-item ...))
      (right-table (make-dialog-item ...))
      (graph-view (make-layout-view ...)))
```

```
(make-layout-dialog :layout
  (:vbox (:width :filler :height :filler)
    (:hbox (:height 1/4 :width :filler)
      (:fbox () left-table) (:fbox () right-table))
      (:fbox () graph-view)))
  (setf (layout graph-view) ...)) ←
```

The browser dialog is specified as a vertical box with `:filler` as width and height. Its first item is a horizontal box whose height is set to 1/4 of that of the vertical box. The horizontal box contains two scrollable tables which are enclosed by frame boxes. The height of these frame boxes is constrained by their surrounding horizontal box. Their width is not explicitly specified, therefore the default value `:filler` is chosen and half of the width of the horizontal box is assigned to each frame box (and its inferior table). The second item of the vertical box is a frame box surrounding the box element `graph-view` which generates the class graph. A layout form which is similar to that defining `graph-view` is shown in Figure 2.

2.3 Local Variables in Layout Forms

We propose *local variables* in layout forms as a useful extension. These variables could represent box attributes such as the actual width and height of a box. These variables could serve as constants within the scope of a form. The following form exemplifies this feature.

```
(let ((item-1 (make-dialog-item ...))
      (item-2 (make-dialog-item ...)))
  (:vbox (:width (:filler :bind ?total-width)) ; ?total-width is local variable
    :filler
    (:hbox ()
      20
      (:fbox (:height (truncate (/ ?total-width 2))) item-1) ; use local variable
        :filler
        (:fbox (:height (truncate (/ ?total-width 2))) item-2) ; use local variable
          20)
      :filler))
```

One interpretation of this specification might be as follows. The evaluation of the form `:bind ?total-width` depends on its lexical context. In this case the value of `?total-width` is set to the actual width of the vertical box. Then, this value is used to determine the height of `item-1` and `item-2`. Therefore, the height of `item-1` and `item-2` depends on the total width of their outermost vertical box. This dependency cannot be expressed in our current box layout scheme.

3 Layout Protocols

Our basic layout algorithms are based on an abstract protocol for manipulating boxes and box items. Therefore, every object conforming to this protocol can be laid out and every (rectangular) region can be interpreted as a box. The protocol is implemented as a set of generic functions. Multi-object methods can be supplied for different kinds of boxes and items,

```

(let ((upper-offset 10)
      (left-offset 10))
  (:vbox ()
    upper-offset
    (:hbox () left-offset
      (:gbox (:dag *roots*
                 #'successors
                 *max-depth*
                 #'appearance
                 ...))))))

```




Figure 2: A general layout specification in combination with a box-style layout.²

which may represent position and size in different ways. The use of multi-object methods also has the advantage that layout of objects may depend on their context. Our layout protocol also allows to specify whether the layout algorithm has to be reapplied if global constraints (e.g. by resizing the window) have been changed.

Apparently it is not reasonable to describe every layout with the box-and-glue metaphor. An obvious example is a set of nodes arranged as a graph. Therefore, our layout language has the notion of a *general box*: (`:gbox (<layout-name> <arg-1>...<arg-n>)`). For instance, this box is used to specify the layout of directed acyclic graphs (DAGs) (see Figure 2).

In this example the items to be arranged are defined inductively by a set of roots, a successor function and a maximal depth. The appearance function is used to compute the graphical representation of nodes (e.g. class objects for an inheritance graph). Position and size of the general box (`:gbox`) are defined implicitly by a closure rectangle around all graph items (see Figure 2). This rectangle defines a box which may be arranged using the box layout specifications already known.

One may also think of other arrangements of box items in a general box. We use the DAG example mentioned above in order to explain our protocol for supplying a new layout specification interpreter.³

```

(defmethod layout-spec-p-using-key ((key (eql ':dag)))
  t)

(defmethod parse-layout-spec-using-key ((key (eql ':dag)) layout-specs)
  "Returns (generated and) laid out :gbox items."
  (interpret-dag-layout layout-specs))

```

The layout name (e.g. `:dag`) of a general box form is used as a key to discriminate the corresponding layout interpreter method, the rest of the form is bound to the parameter `layout-specs`.

² The indentation of the boxes is used for demonstration purposes only.

³ Layout forms are usually defined as macros evaluating the right expressions (in the right scope).

This extension scheme exploits that CLOS methods are not only attached to objects (or their classes) but can be also discriminated on every Lisp object. Layout forms are represented as lists, i.e. layout descriptions can easily be manipulated (e.g. by pattern matching algorithms).

4 Interaction Objects and Views

We introduce box items which are suited for more general applications and describe their application to our class browser which serves as a simple example for describing the basic features of our layout language. We refer to *Möller 90* for a discussion of more sophisticated user interfaces.

4.1 Interaction Objects

Interaction objects represent box items of layout specifications. Interaction objects (e.g. all standard elements of the Macintosh Toolbox, graph nodes, graph edges) are instances of CLOS classes. We defined an additional interaction object, a so-called *view*. Views provide a framework for handling non-standard interaction components. These components are called *view items*. View items can be freely added to and removed from views. The interactive behavior of view items can easily be modified by adding certain predefined superclasses (mixins) to their class definitions. Typical desired behaviors are to move, select, or mark items. The algorithms ensuring a consistent image on the screen are provided by views. Views can also be declared as scrollable.

The system evaluates generic functions to draw and delete visible items if required. For the nodes in the class browser example (Figure 1) we might use the following definitions.

```
(defclass label-view-item (movable-view-item-mixin view-item)
  ((label :initarg :label :accessor label)))

(defun make-label (...) ...)
```

Class `label-view-item` is defined by inheriting the base class for view items and a ‘mixin’ which allows to move items interactively. The generic function `view-item-draw` is evaluated to draw view items:

```
(defmethod view-item-draw :after ((item label-view-item) view dialog)
  (let ((position (view-item-position item))
        (size (view-item-size item)))
    (frame-round-rect dialog std-gcontext ...)
    (move-to dialog ...)
    (draw-string dialog std-gcontext (label item))))
```

The drawing functions use *graphical contexts*⁴ defining coordinate transformations and drawing attributes. The position and the drawing vector of a view item define a *drawing rectangle* (see Figures 3 and 4) which is used as clipping rectangle. Specialized methods depending on particular drawing contexts (e.g. views or dialogs) can easily be defined.

⁴ The notion of a graphical context is defined by views.




Figure 3: Referenced Location of Item B.

4.2 Referencing View Items

View items may reference other view items. References to other view items are also declaratively described. The corresponding coordinates are automatically computed. This scheme defines edges as view items which reference their corresponding nodes. For instance, the standard drawing function of an edge retrieves the references to its nodes and computes coordinates used to draw the connecting line. References are described by a so-called *reference box* (`:rbox`). A reference box consists of two view item, the referencing item (e.g. A) and the referenced item (e.g. B), and a pair of horizontal and vertical coordinates specifying the location which is referenced. Figure 5 shows the complete syntax.

```
(:rbox A B
  (:horizontal :filler :reference :filler)
  (:vertical :filler :reference :filler))
```

This box defines a reference from item A to item B. The keyword `:reference` specifies the referenced point. Figure 3 shows an example for this type of reference. The start point of item A (gray vector) is defined as reference to the center of item B (circle). The reference box of item B is shown as gray rectangle. It is defined as the drawing rectangle of item B. The usage of filler specifications in the horizontal and vertical coordinates ensure the centering of A's reference point to B. The next example defines a reference point which is 3 pixel left and above to the lower right corner of item B.

```
(:rbox A B
  (:horizontal :filler :reference 3)
  (:vertical :filler :reference 3))
```

Edges of graphs are also represented as view items. We apply the concepts introduced insofar to define a class `line-view-item` representing edges. The end points of an edge are specified as references to its nodes. The drawing function of `line-view-item` uses the predefined functions `references-of-this-item` and `reference-position` in order to compute the coordinates of the two points defining the edge (see Figure 4).



Figure 4: Location of edges defined by node references.

```
(defclass line-view-item (view-item)
  ())

(defmethod view-item-draw :after ((item line-view-item) view dialog)
  (let* ((references (references-of-this-item item))
         (p1 (reference-position (first references)))
         (p2 (reference-position (second references))))
    (move-to dialog p1)
    (line-to dialog std-gcontext p2)))
```

References may also be used to define an erase function specialized for edges. Only the line representing an edge is erased.

```
(defmethod view-item-undraw ((item line-view-item)
                             view dialog
                             position size references)
  (using-gcontext ((eraser-gcontext :pen-pattern *white-pattern*))
    (let ((p1 (reference-position (first references)))
          (p2 (reference-position (second references))))
      (move-to dialog p1)
      (line-to dialog eraser-gcontext p2)))))
```

If the location of the rounded rectangles (Class-1, Class-2) is changed the reference points P1 and P2 are recomputed and the line is redrawn.

The layout specification of our class browser uses reference boxes to define node-connecting edges. The following specifications replace (at the position marked by ←) and supplement the specification given in Section 2.2.

```
(setf (layout graph-view)
  (:vbox ()
    10 ; 10 Pixel upper border
    (:hbox ()
      10 ; 10 Pixel left border
      (:gbox (:dag (list (find-class class-name)) ; list of DAG roots
                #'class-direct-subclasses ; successor function
                *hierarchy-depth* ; max. expansion depth
                #'(lambda (class) t) ; expansion predicate
                #'(lambda (class) ; node-creating function
                    (make-label (class-name-as-string class)))
                #'make-line-view-item ; edge-creating function
                #'western-reference ; start point of edge
                #'eastern-reference)))))) ; end point of edge
```

```

(defun western-reference (referencing-object referenced-object)
  "Definition of reference points with :rbox form"
  (:rbox referencing-object
         referenced-object
         (:vertical :filler :reference :filler)
         (:horizontal :reference :filler)))

(defun eastern-reference (referencing-object referenced-object)
  "Definition of reference points with :rbox form"
  (:rbox referencing-object
         referenced-object
         (:vertical :filler :reference :filler)
         (:horizontal :filler :reference)))

```

4.3 Named References of View Items

A useful extension might be to assign names to descriptions of reference boxes. They would enable the user to define higher-level abstractions such as directed references. For instance, the following specifications define the start and end point of an object `arrow` which points from the lower right corner of object `node-1` to the upper left corner of object `node-2`.

```

(:rbox :arrow-start arrow node-1
       (:horizontal :filler :reference)
       (:vertical :filler :reference))

(:rbox :arrow-end arrow node-2
       (:horizontal :reference :filler)
       (:vertical :reference :filler))

```

The predefined function `reference-description` can be applied to a reference and may return the corresponding reference form which defined this reference and/or the reference name (e.g. `:arrow-start`). For instance, this feature could be used by an arrow-drawing method for determining the arrow direction.

```

(defmethod view-item-draw :after ((item arrow-view-item) view dialog)
  (let* ((references (references-of-this-item item))
         (rd1 (reference-description (first references)))
         (rd2 (reference-description (second references)))
         (p1 (reference-position (first references)))
         (p2 (reference-position (second references))))
    (cond
      ((and (eq rd1 ':arrow-start) (eq rd2 ':arrow-end)) ; draw arrow
       (draw-arrow dialog std-gcontext p1 p2)) ; from p1 to p2
      ((and (eq rd2 ':arrow-start) (eq rd1 ':arrow-end)) ; draw arrow
       (draw-arrow dialog std-gcontext p2 p1)) ; from p2 to p1
      (t (error ...))))

```

```

<layout-form> ::= <hbox-form> | <vbox-form> | <fbox-form> |
                <gbox-form> | <rbox-form>

<hbox-form> ::= (:hbox <box-size-spec>
                {<layout-spec> |
                 <layout-form> | <item> | <splice>})

<vbox-form> ::= (:vbox <box-size-spec>
                {<layout-spec> |
                 <layout-form> | <item> | <splice>})

<fbox-form> ::= (:fbox <box-size-spec> <item>)

<layout-spec> ::= <size-spec> | <filler-spec>

<box-size-spec> ::= ([:width <layout-spec>] [:height <layout-spec>])

<size-spec> ::= <integer> | <rational> | <float>

<filler-spec> ::= :filler |
                (:filler [:min <filler-size-spec>]
                 [:max <filler-size-spec>])

<filler-size-spec> ::= <integer> | :as-needed5

<item> ::= <S-expression>

<splice> ::= (:splice <S-expression>)

<gbox-form> ::= (:gbox (<user-defined-layout-name>
                    <arg-1>...<arg-n>))

<rbox-form> ::= (:rbox <referencing-object> <referenced-object>
                <horizontal-ref> <vertical-ref>6)

<horizontal-ref> ::= (:horizontal <distance-spec> :reference <distance-spec>)

<vertical-ref> ::= (:vertical <distance-ref> :reference <distance-ref>)

<distance-ref> ::= {[:filler] {<abs-distance>} {<rel-distance>}}

<abs-distance> ::= <positive integer>

<rel-distance> ::= <rational> ∈ [0,1] | <float> ∈ [0,1]7

<referencing-object> ::= <view-item>

<referenced-object> ::= <view-item>

```

Figure 5: EBNF Syntax of Layout Specifications.

⁵ Only valid in <hbox-form> or <vbox-form>.

⁶ Both components may be exchanged.

⁷ A relative size of a reference box refers to the enclosing rectangle of its referenced view item.

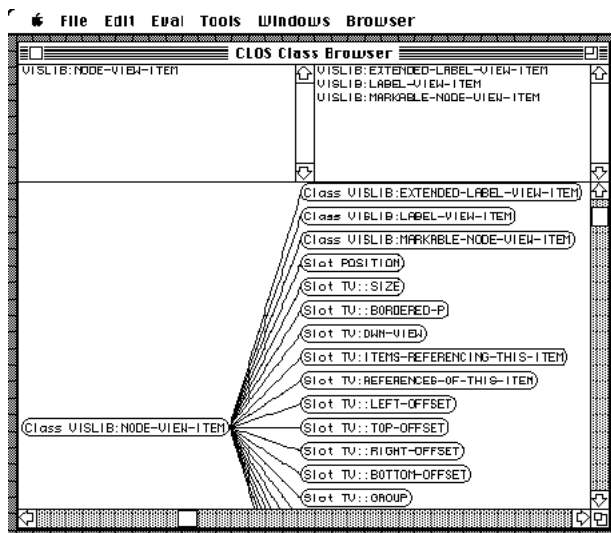


Figure 6: A DAG graph of a class hierarchy with local class slots.

5 Extended Class Browser and Inspector

This section shows an extended class browser and inspector. The modified class browser additionally displays direct slots of classes (see Figure 6). These extensions were easily achieved by defining appropriate methods for the generic functions `successors` and `appearance` (see layout form in Figure 2). These modifications serve as an example for the flexibility of our approach.

```
(defmethod successors ((any t))
  "Do not show arbitrary objects"
  nil)

(defmethod successors ((class standard-class))
  "Show direct subclasses and slots"
  (append (class-direct-subclasses class) (class-direct-slots class)))

(defmethod appearance ((class standard-class))
  "Create a graphical class representation"
  (make-label class ...))

(defmethod appearance ((slot standard-slot-description))
  "Create a graphical slot representation"
  (make-label slot ...))
```

The following example is part of a user interface of a CLOS inspector which has been implemented within our framework. The inspector displays a window consisting of tabular subwindows. Figure 7 shows information about a class `tv:window`. It displays the class precedence list, direct super and subclasses, slot information, and direct (locally defined) methods. If necessary, tables may be scrolled provided the space available to a table is not sufficient for displaying all table elements.

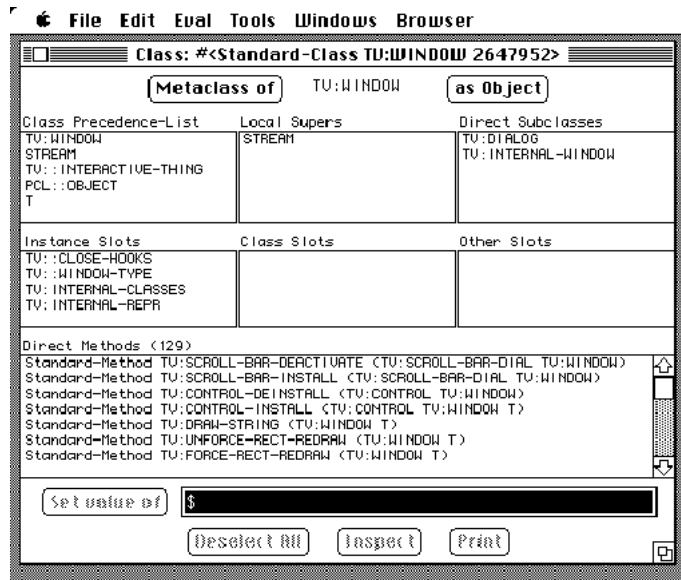


Figure 7: CLOS class inspector window (default size).

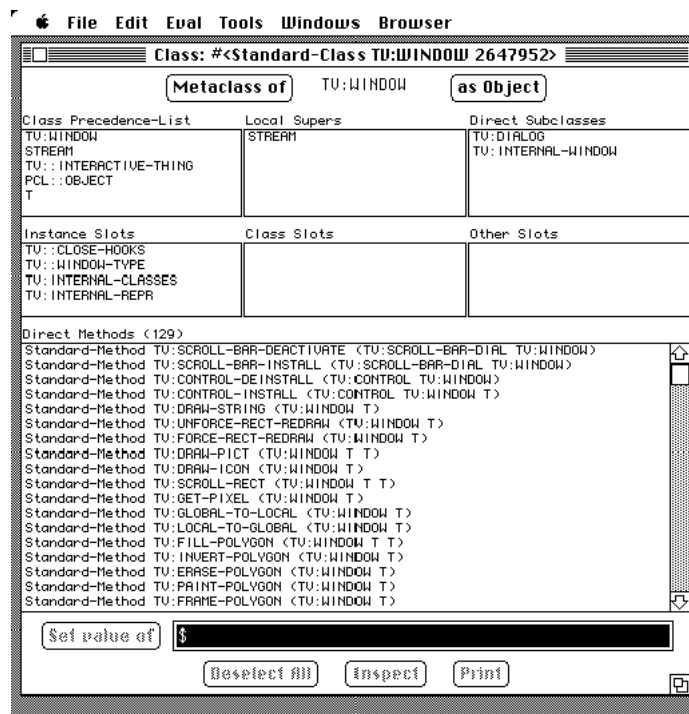


Figure 8: CLOS class inspector window (vertically and horizontally enlarged).




Figure 9: Schematic representation of nested boxes. Arrows represent fillers.

The corresponding layout specification is composed of several parts and the constraints are partially computed at runtime. Therefore, we only explain its general outline. The upper rows are specified as horizontal boxes containing three frame boxes which represent tables. The width and height of the frame boxes are defined by constrained fillers. Each filler acts like a spring. All fillers compete for the space available to them which, in fact, is constrained by their surrounding horizontal box. Therefore, after a relaxation process each filler acquires 1/3 of the available space. Roughly, the horizontal boxes compete for the vertical space in a similar manner. Minimal and maximal constraints guarantee that sparse tables have a visually appealing uniform shape.

Our notion of constrained fillers is important for describing a flexible window layout. For instance, Figure 8 shows the same window as in Figure 7 except the vertical and horizontal space available to the window has been enlarged. The table displaying the direct methods of class `tv:window` adapted to the new space constraint. The additional vertical space was completely consumed by this table since it is the only element whose upper filler constraint is still unsatisfied.⁸

6 The Box Layout Algorithm Revisited

This section discusses the basic box layout algorithm in more detail. Our formalism provides the user with a fast layout algorithm which is also easy to understand and anticipate. Important features of this algorithm are the allocation of arbitrary (interaction) objects conforming to the underlying protocol. The size of interaction objects may depend on predefined values or space constraints. These constraints can be expressed by minimal and maximal expansion values which are represented as fillers. Boxes may be arbitrarily nested but only fillers at the same level compete for the available space. Figure 9 exemplifies this feature.

The (vertical) fillers of level 1 are independent of the (vertical) fillers of level 3. This restriction

⁸ Actually, its upper limit for vertical space is set to a height which would allow to display the whole list of direct methods at once provided that sufficient screen space is available.




Figure 10: Water jugs modeling fillers.

reduces considerably the computational complexity of the layout algorithm. Another advantage is that the semantics of constrained fillers are easy to comprehend by the user. Even this scheme requires a simple relaxation algorithm for satisfying these constraints. Figure 10 illustrates our algorithm.

Each *water jug* models a filler with its minimal and maximal expansion value. The current extension of a filler corresponds to the water contents of a jug. The space available to fillers is modeled by an external *water reservoir* with capacity R_0 . The extension of a filler is computed by the following steps:

1. Initially each water jug j is filled to a level L_0^j . It is guaranteed that every water jug j is at least filled to its minimum \min_j . In case of an additional demand for water which cannot be supplied by an empty water reservoir⁹, this demand is satisfied by a “water pipe”. If there is any water R_1 in the reservoir left, the remainder of this water is distributed.
2. Each water jug gets a portion $P_i = R_i/N_i$ of the reservoir (be N_i the number of water jugs in the i th iteration). Step 1 may have caused for different water levels of the jugs. Therefore, every jug is filled to a common minimal level

$$L_i = M_i + P_i, \quad \text{with } M_i = \min_{j \in 1..N_i} L_{i-1}^j$$

3. Caused by the minimum satisfaction guarantee the level of a jug may already be higher than L_i or its maximum \max_j . The amount of water of a jug whose level¹⁰ exceeds either L_i or \max_j is returned to the reservoir. A jug which has reached its maximal level becomes inactive ($N_i = N_i - 1$).
4. Start again with step 2 ($N_{i+1} = N_i$; $i = i + 1$) provided the water reservoir is not empty and there is at least one jug j left whose water level is below \max_j . This algorithm terminates if all jugs have become inactive ($N_i = 0$) or the water reservoir is empty.

The algorithm’s termination is guaranteed since each cycle either makes one jug inactive or removes water from the reservoir. If no jug becomes inactive at least one jug is filled to the level M_i . Rounding errors are summed up. As final step this sum is rounded and added to the last filler.¹¹

⁹ This is an indication that the available space is not sufficient to fulfill the space requirements. Items could extend beyond the boundary of their surrounding box.

¹⁰ The capacity of a jug be temporarily unlimited.

¹¹ This is needed for frame boxes which may otherwise not exactly fit to their box’s lower border.

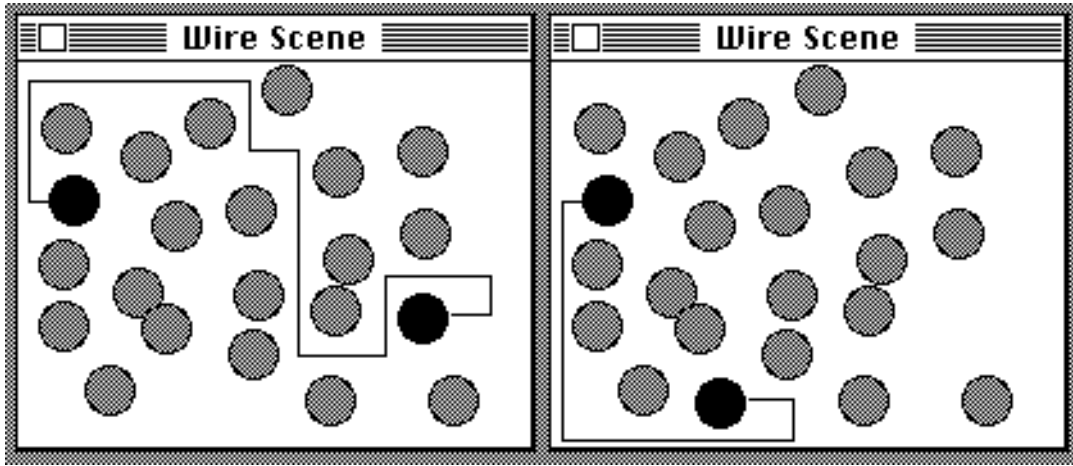


Figure 11: The right window shows the state after a black circle has been moved.

7 Optimizing Graphical Output

Our framework also supports the easy integration of optimization algorithms. We exemplify this with Figure 1. If the upper left graph node labeled with `TV:VIEW-ITEM` is moved, the edges referencing this node have to be informed about the location change. A simple solution might be that every edge concerned by this change reevaluates its reference to this node and redraws its line. The disadvantage of this solution is that each edge performs the same operation. Furthermore, the erasing and redrawing of all lines is performed in an uncoordinative way. But our combination of references and views supports a better strategy. Each edge receives only a translation vector which has to be added to the corresponding reference points.

After a graph node has been moved the edges referencing this node receive a corresponding translation vector. An uncoordinated redrawing usually causes “ping-pong” effects, e.g. redundant erasing and redrawing of lines which have again to be moved subsequently. Our strategy avoids this problem by using a *drawing cache*. This cache is associated with views and totally transparent to the user. Instead of immediately evaluating the generic drawing functions only drawing resp. erasing *orders* are buffered. Particular operations force to empty the cache and execute the buffered operations. For instance, when a graph node has been moved the cache will be emptied, i.e. after evaluating only *once* the coordinate transformations all buffered erasing functions are executed at first and then all drawing functions.

We conclude this section with another example which again emphasizes that our visualization algorithms can be applied to every CLOS object conforming to the predefined protocols. Figure 11 shows a visualization of a line routing algorithm. It is used to explore fast and visually appealing line routing algorithms which are essential for many visual programming applications. The windows contain two objects (black circles) which are connected by a line and an arbitrary number of other objects (gray circles). Each circle may be interactively moved. Each movement invokes the line routing algorithm in order to adapt the line path.

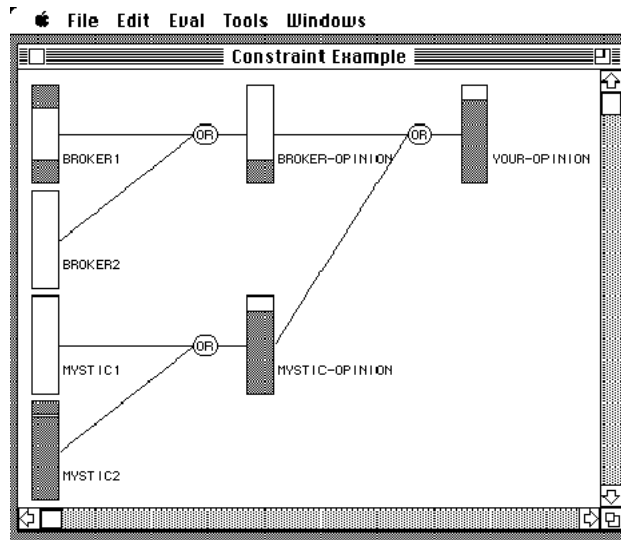


Figure 12: The upper and lower bounds are indicated by shaded rectangles [Winston & Horn 89]. The gauges show the estimation interval from 0 (bottom) to 1 (top).

8 Meta-Level Techniques for Separating Application and Visualization

Application and visualization objects have to be separated. In the following we discuss how to use the meta-object protocol of CLOS for separating application and visualization layers. We explain these considerations by using an animated visualization of a simple constraint system.

Our approach associates visualization objects with given application objects without requiring any modifications to the application. We support multiple views as well as controllers for manipulating the application's data structures. Several other mechanisms have been developed (Model-View-Controller-Scheme [Goldberg & Robson 83], CLUE [Kimbrough & LaMotte 89], Presentation-Types [Symbolics 88]). In this section we also discuss how basic features of these systems can be realized using our approach.

As example application we chose a simple constraint net. There is no need to present the application code since everything can be found in detail in *Winston & Horn 89*. Our visualization was generated without any modifications to the application code. The application provides a simple model of a stock exchange scenario. When are some stocks to split? The participants have uncertain knowledge and are influenced by one another. A constraint net models these influences by propagating certainty estimation intervals between 0 and 1. This interval of a 'broker' might be visualized by a gauge as found in *Winston & Horn 89*. The implementation distinguishes assertion objects (brokers, mystics, virtual intermediates, etc.) and constraint objects (or, and). Figure 12 shows an overview of an example configuration with gauges for assertions and simple nodes for constraints.

The visualization in Figure 12 can be described with the following layout descriptions.

```
(defun stock-exchange-connections-visualization (participants)
  "Opens a window and shows the connections of
  the given participants in a scrollable view."
```

```

(let ((stock-exchange-view
      (make-layout-view :scroll-bars ':both
                       :bordered-p nil
                       :auto-scrolling t)))
      (make-stock-exchange-dialog (:fbox () stock-exchange-view))
      (setf (layout stock-exchange-view)
            (:vbox ()
              10
              (:hbox () 10 (:gbox (:dag participants
                                  #'stock-exchange-wizard
                                  *max-connection-depth*
                                  #'application-visualization-coupler
                                  ...)))))))

```

The whole dialog consists of a view (laid out with an `:fbox`). The graph is defined by the set of participants and the successor function `stock-exchange-wizard`.

The function `application-visualization-coupler` defines a mapping from application objects to visualization objects. Both functions are generic, i.e. different mappings may be specified for different classes of application objects.

```

(defmethod stock-exchange-wizard ((participant assertion))
  "Wizard's information about connection of assertion objects."
  (assertion-constraints participant))

(defmethod stock-exchange-wizard ((participant constraint))
  "Wizard's information about connection of constraint objects."
  (list (constraint-output participant)))

```

8.1 Indirect Values for Visualization Objects

Visualization objects have to refer to objects of the application side. There should exist a “dynamic” binding which could be easily maintained provided that classes of visualization objects offer support for some kind of active values [Bobrow & Stefik 83]. We present a simplified CLOS metaclass supporting non-nested active values which we call *indirect values*. Indirect values are defined by the form `#`(object reader writer)` where `writer` is optional. The following method sketches an implementation using a new metaclass and a corresponding meta-level method for the generic slot accessor function `slot-value-using-class`. Writing to slots with indirect values can be implemented analogously.

```

(defclass indirect-slots-class (standard-class)
  ())

(defmethod check-super-metaclass-compatibility ((x indirect-slots-class)
                                               (y standard-class))
  t) ; We do not care about that in this report.12

```

¹² We refer to *Graube 89*.

```
(defmethod slot-value-using-class ((class indirect-slots-class) object slot-name)
  (let ((direct-slot-value (call-next-method)))
    (if (indirectp direct-slot-value)
        (funcall (indirect-reader direct-slot-value)
                 (indirect-object direct-slot-value))
        direct-slot-value)))
```

Visualization objects have `indirect-slots-class` as metaclass. Using indirect slot values every slot access is delegated to the corresponding application object if required. Using the meta-object protocol it would be easy to determine all indirect objects or that indirect object referred to by a specific slot. The gauges for the stock exchange example use this metaclass to refer to the exchange participants. But what about the other direction: the gauges have to be “informed” if the participants’ estimations of stock splits change.

8.2 Slot Demons for Application Objects

An assertion object has one slot for the lower bound and one for the upper bound estimation. The corresponding visualization objects have to be informed when either of these slot values change. The most obvious way to achieve this is to define the assertion class with a metaclass that allows *demon functions* to be attached to slots. The “real” value of a slot is a structure that provides a value facet and an `if-modified` facet [Roberts & Goldstein 77]. All slot demon functions are evaluated when the slot value changes. The implementation of slot demons is similar to the one of indirect slot values. We introduce a metaclass `demon-slots-class` and define modified versions of `slot-value-using-class` and `(setf slot-value-using-class)` which access the value facet. The latter one evaluates the demons in the `if-modified` facet. Slot demons should be made removable.

Thus, the function `application-visualization-coupler` mentioned above can be defined as follows.

```
(defmethod application-visualization-coupler ((participant assertion))
  (let ((assertion-gauge (make-two-level-gauge ; indirect values
                          #`(participant assertion-lower-bound)
                          #`(participant assertion-upper-bound))))
    (add-slot-if-modified-demon
     participant ; object
     'lower-bound ; slot name
     #'(lambda ; demon function
          (assertion-obj name-of-modified-slot old-value new-value)
          (gauge-update assertion-gauge)))
    (add-slot-if-modified-demon
     participant ; object
     'upper-bound ; slot name
     #'(lambda ; demon function
          (assertion-obj name-of-modified-slot old-value new-value)
          (gauge-update assertion-gauge)))
    assertion-gauge))
```

Generic function methods

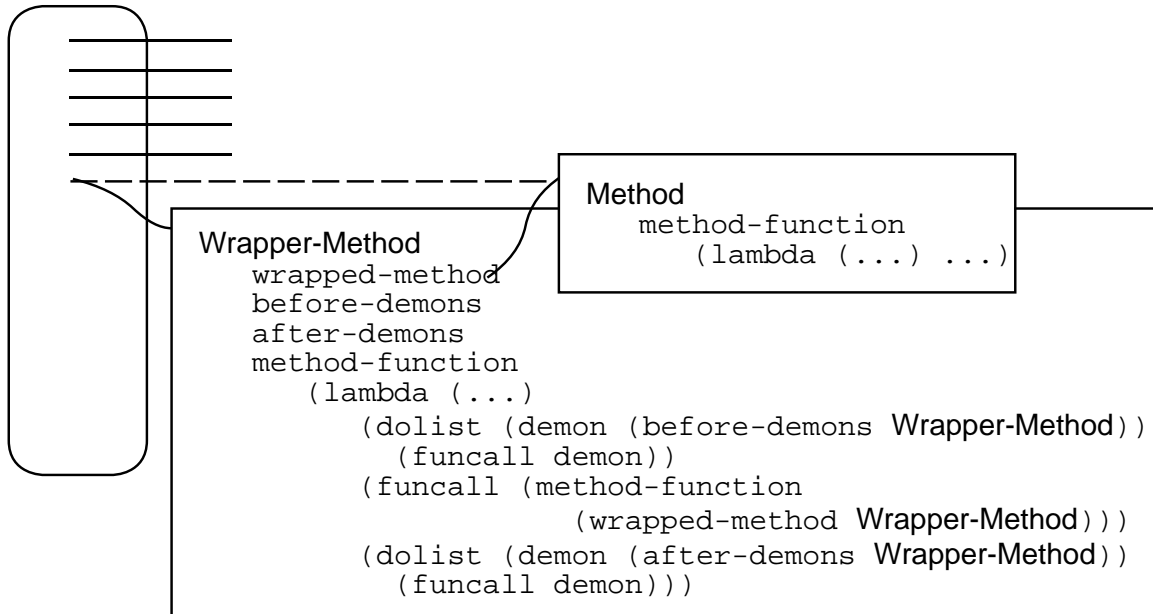


Figure 13: Outline of a wrapper method

```
(defmethod application-visualization-coupler ((participant or-box)
  (make-label "OR"))
```

Demon functions are closures which provide access to the corresponding visualization object. The gauges for assertion objects (broker, etc.) use indirect values to access assertion objects. The objects representing labels for constraints are the same as in the class browser example.

8.3 Method Demons

Slot demons offer an elegant way of defining slot accesses as interesting events and hence updating corresponding visualization objects. Not only slot accesses are subject to updating a visualization. Every method might define an event of interest. Slot accesses are only special cases. General *method demons* can be implemented using the meta-object protocol of CLOS. The idea¹³ is to wrap a method with a so-called *wrapper method* which has slots to refer to both the demon functions and the original method (see Figure 13). When all demons are removed the wrapper method itself is removed, too. In this case there is no overhead as with a metaclass which provides own methods for slots accesses that overwrite the standard slot accessor methods (e.g. for indirect values).

A major disadvantage of this wrapping slot accessor is that demons are evaluated for all instances, i.e. they are slot but not instance-specific. Method demons do not solve the problem of compound slot accesses, either.

¹³ An implementation proposal for CLOS was originally outlined by Gregor Kiczales.

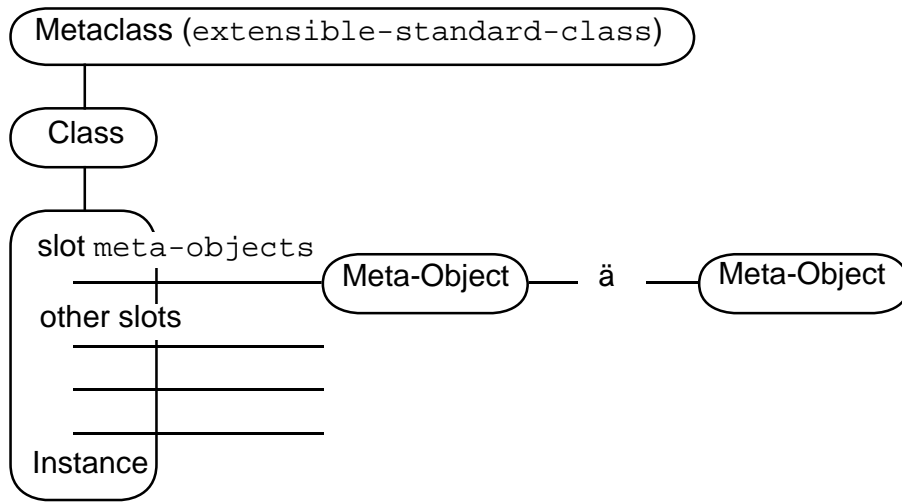


Figure 14: Meta-objects for instances with metaclass `extensible-standard-class`.

8.4 Instance-Specific Meta-Objects

The CLOS meta-object system assigns to metaclasses the responsibility for both structure (implementation) and behavior of instances. There are other meta-level systems which distinguish between structural and computational meta-objects [Ferber 89]. In this section we present ideas to provide some kind of dynamic meta-level influence in CLOS [Cunis 90]. We implement meta-objects as instances of the class `standard-meta-object` which is not a CLOS metaclass. The standard slot access protocol which uses the method `slot-value-using-class` is analogously extended for these “simple” meta-objects.

```

(defclass standard-meta-object ()
  ())

(defmethod slot-value-using-meta-object ((mobj standard-meta-object)
                                         object slot-name)
  (call-next-meta-method))

(defmethod (setf slot-value-using-meta-object) ((mobj standard-meta-object)
                                                object slot-name)
  (call-next-meta-method))

```

Meta-objects can be assigned to instances with metaclass `extensible-standard-class`. This metaclass describes classes with instances that have one additional or implicit slot called `meta-objects` (see Figure 14). A set of meta-objects can be assigned to this slot.

An example method handling slot accesses is defined as follows.

```

(defmethod slot-value-using-class ((class extensible-standard-class)
                                   object slot-name)
  (if (eq slot-name 'meta-objects)
      (call-next-method)
      (let ((*meta-objects* (slot-value object 'meta-objects)))
        (call-next-method)))

```

```

(*meta-class-generic-function* #'(lambda () (call-next-method)))
(*meta-object-generic-function*
 #'(lambda (meta-object) (slot-value-using-meta-object meta-object
                                                         object
                                                         slot-name))))

(declare (special *meta-objects*
                 *meta-object-generic-function*
                 *meta-class-generic-function*))
(if (null *meta-objects*)
    (call-next-method)
    (call-next-meta-method))))

```

The function `call-next-meta-method` is comparable to the function `call-next-method`. It evaluates `slot-value-using-meta-object` for the “next” meta-object in the list of meta-objects (see Figure 14). The default behavior of `slot-value-using-meta-object` is to evaluate `call-next-meta-method` again (s.a.). This default behavior may be augmented or overwritten by subclasses of `standard-meta-object` (s.b.). When there are no meta-objects (left), `call-next-meta-method` invokes the “normal” slot access functionality of `standard-class`. There is also some additional code needed to enable passing of different parameters to the next metamethod just as with `call-next-method`.

```

(defun call-next-meta-method ()
  (declare (special *meta-objects*
                  *meta-object-generic-function*
                  *meta-class-generic-function*))
  (if (endp *meta-objects*)
      (funcall *meta-class-generic-function*)
      (funcall *meta-object-generic-function* (pop *meta-objects*))))

```

We use these meta-level techniques to extend our constraint example. Using the protocol described above visualizations of particular instances can be provided with little programming effort. For instance, a meta-object could be defined by the class `visualizer-meta-object`. This class combines a visualization object with a list of interesting slots. Every writing access to these slots is followed by calling the instance-specific visualization object.

```

(defclass visualizer-meta-object (standard-meta-object)
  ((visualizer :initarg :visualizer
              :accessor visualizer
              :initform #'(lambda (&rest ignore) nil))
   (interesting-slots :initarg :interesting-slots
                     :reader interesting-slots)
   (:default-initargs :interesting-slots nil))
  (defmethod (setf slot-value-using-meta-object) :after (new-value
                                                         (mobj visualizer-meta-object)
                                                         object slot-name)
    (if (member slot-name (interesting-slots mobj))
        (funcall (visualizer mobj) object slot-name)))

```

We define a new metaclass for assertion objects which combines slot demons and meta-objects.


```
(defclass extensible-standard-class-with-slot-demons (extensible-standard-class
                                                     demon-slots-class)
  ())
```

Be `your-opinion` the assertion object of our constraint example. We add only to this object a corresponding meta-object which prints `your-opinion`'s decision about buying stocks. This behavior can be easily reverted by removing this meta-object from the implicit slot meta-objects (see Figure 14).

```
(add-meta-object your-opinion
  (make-instance 'visualizer-meta-object
                 :interesting-slots '(lower-bound upper-bound)
                 :visualizer
                 #'(lambda (assertion slot-name)
                     (if (> (assertion-lower-bound assertion) 0.75)
                         (print 'buy) ; or any other visual feedback
                         (print 'donot-buy))))))
```

Another behavior might be to temporarily modify a reading access to a slot value. After adding a meta-object of class `buying-indicator-meta-object` to the object `your-opinion` each reading access to the slot `lower-bound` of `your-opinion` returns the slot value and a buying indicator.

```
(defclass buying-indicator-meta-object (standard-meta-object)
  ())

(defmethod slot-value-using-meta-object ((mobj buying-indicator-meta-object)
                                         object slot-name)
  (if (eq slot-name 'lower-bound)
      (let ((slot-value (call-next-meta-method)))
        (if (> slot-value 0.75)
            (values slot-value 'buy)
            (values slot-value 'donot-buy)))
      (call-next-meta-method)))

(add-meta-object your-opinion (make-instance 'buying-indicator-meta-object))
```

One may of course argue that this implementation is a little impure because of using different mechanisms: metaclasses and meta-objects. Moreover not all meta-objects may be compatible. There remains also some overhead although no meta-objects are attached to an instance.

9 Related Work

The SymbolicsTM programming environment GeneraTM offers also means for specifying layout of windows [Symbolics 88]. Subwindows (panes) can be arranged in a frame in columns or rows. Window sizes can be determined as absolute (fixed), relative, or with respect to objects being allocated. These features can be compared with our box model. Filler specifications are also supported. Size specifications can be constrained by minimal and maximal distances. Genera only supports layouts for panes, but our layout algorithms can be applied to every object

conforming to the underlying abstract protocol. Genera offers the notion of presentation types which can be compared with our view item classes. Presentation types are also associated with handling user input. In contrast to Genera our approach offers a more uniform and orthogonal layout scheme combined with a compact and elegant \TeX -notation.

Recently, two other approaches were proposed which use \TeX -like layout schemes for user interfaces. They also use constructs such as boxes and fillers for expressing window layout. The InterViews System [Linton et al. 89] is a user-interface toolkit based on X windows and implemented in C++. Fillers and boxes are implemented as objects. With respect to our Lisp environment we found the representation of boxes as a combination of ordinary lists and macros more efficient for manipulation and pattern matching. Our layout scheme is in several respects more powerful than InterViews' scheme. A filler-like size specification of boxes is not possible in InterViews. Important and useful notions such as a frame box which constrains the size of a its box element or a general box which invokes user-defined parsers for layout specifications are not available.

The FormsVBT system [Avrahami et al. 89] offers a two-view approach to designing user interfaces. The layout of a dialog window can be specified using both a \TeX -like textual and a direct-manipulative graphical representation. Changes made in either representation are immediately updated in the other representation. FormsVBT is implemented in a dialect of Modula-2. Its specification language supports no macros and offers no support for new box types and layout schemes. Furthermore, we see the problem that the functionality of the textual specification notation has to conform with the graphical user interface. Mostly, this requires to reduce the functionality of the textual notation.

10 Summary and Future Work

This report presented a framework for visualizing object-oriented systems. It consists of a compact, flexible notation for specifying layout of graphical objects. This notation is fully integrated into a Lisp environment based on CLOS. Advantages of this \TeX -like notation are its expressiveness, user-predictable layouts, and efficient implementation schemes. The CLOS meta-level architecture is used to associate visualization and application objects. Supported techniques are indirect values, slot and method demons, and instance-specific meta-objects. These visualization techniques require no modifications to the systems which are selected for visualization.

Next steps might be to combine the advantages of \TeX -style notations with the general flexibility of constraint systems. Another useful extensions might be to support local variables in box specifications (see section 2.3) and naming of references (see Section 4.3). We also plan to address the problem of interpreting several different generic functions as a single interesting event. One solution might be to define higher-level demons combining demons of different methods. Furthermore, research is necessary to extend this approach to 2-1/2 or 3-D layout.

Acknowledgements

We are grateful for thoughtful comments from Roman Cunis and Ken Kahn. The first author has been partly supported by a DAAD scholarship granted by the NATO science committee.

References

- Avrahami et al. 89:** A Two-View Approach to Constructing User Interfaces, G. Avrahami, K.P. Brooks, M.H. Brown, *ACM Computer Graphics* **23**, 3 (July 1989), 137–146.
- Bobrow et al. 88:** Common Lisp Object System Specification, D.G. Bobrow, L.G. DeMichiel, R.P. Gabriel, S.E. Keene, G. Kiczales, D.A. Moon, *ACM Sigplan Notices* **23**, 9 (Sept. 1988).
- Bobrow & Stefik 83:** The LOOPS Manual, D.G. Bobrow, M. Stefik, Xerox Corporation, December 1983.
- Borning 81:** The Programming Language Aspects of ThingLab, a Constraint-Oriented Simulation Laboratory, A. Borning, *ACM Transactions on Programming Languages and Systems* **3** (1981), 353–387.
- Brown 88:** Algorithm Animation, M.H. Brown, ACM Distinguished Dissertations Series, MIT Press, 1988.
- Cox et al. 89:** Prograph: a step towards liberating programming from textual conditioning, P.T. Cox, F.R. Giles, T. Pietrzykowski, In: Proceedings, *1989 IEEE Workshop on Visual Languages*, Rome (Italy), Oct. 4-6, IEEE Computer Society Press, 1989, pp. 150–156.
- Cunis 90:** Some Ideas on Integrating Reflective Aspects into CLOS-type Object Systems, R. Cunis, Internal Report, University of Hamburg, Computer Science Department, 1990.
- Cunningham & Beck 86:** A Diagram for Object-Oriented Programs, W. Cunningham, K. Beck, *ACM Sigplan Notices* **21**, 11 (1986), 361–367.
- Eisenstadt & Brayshaw 88:** The Transparent Prolog Machine (TPM): An Execution Model and Graphical Debugger for Logic Programming, M. Eisenstadt, M. Brayshaw, *Journal of Logical Programming* **5** (1988), 277–342.
- Ferber 89:** Computational Reflection in Class-Based Object-Oriented Languages, J. Ferber, *ACM Sigplan Notices* **24**, 10 (1989), 317–335.
- Goldberg & Robson 83:** Smalltalk-80: The Language and its Implementation, A. Goldberg, D. Robson, Addison-Wesley, Reading, Mass., 1983.
- Graube 89:** Metaclass Compatibility, N. Graube, *ACM Sigplan Notices* **24**, 10 (1989), 305–315.
- Haarslev & Möller 90:** VIPEX: Visual Programming of Experimental Systems, V. Haarslev, R. Möller, In: *Visual Languages and Visual Programming*, S.K. Chang (ed.), Plenum Press, New York and London, 1990, pp. 185–212.
- Kahn & Saraswat 90:** Complete Visualizations of Concurrent Programs and their Executions, K.M. Kahn, V.A. Saraswat, In: *Visual Languages 90*, pp. 7–14. See also *Technical Report SSL-90-38* [P90-00099], Xerox Palo Alto Research Center, 1990.
- Keene 89:** Object-Oriented Programming in CLOS - A Programmer's Guide to CLOS, S. Keene, Addison-Wesley, 1989.
- Kimbrough & LaMotte 89:** Common Lisp User Interface Environment, K. Kimbrough, O. LaMotte, Preprint, Texas Instruments Inc., July 1989.
- Kleyn & Gingrich 88:** GraphTrace - Understanding Object-Oriented Systems Using Concurrently Animated Views, M.F. Kleyn, P.C. Gingrich, *ACM Sigplan Notices* **23**, 11 (1988), 191–205.
- Knuth 79:** T_EX and Metafont – New Directions in Typesetting, D.E. Knuth, Digital Press, 1979.
- Linton et al. 89:** Composing User Interfaces with InterViews, M.A. Linton, J.M. Vlissides, P.R. Calder, *IEEE Computer* **22**, 2 (1989), 8–22.

- Maloney et al. 89:** Constraint Technology for User-Interface Construction in ThingLab II, J.H. Maloney, A. Borning, B.N. Freeman-Benson, *ACM Sigplan Notices* **24**, 10 (1989), 381–388.
- Matwin & Pietrzykowski 85:** PROGRAPH: A Preliminary Report, S. Matwin, T. Pietrzykowski, *Computer Languages* **10**, 2 (1985), 91–126.
- Möller 90:** AI-Based Visualization Tools in Object-Oriented Systems (in German), R. Möller, *Technical Report FBI-HH-B-149/90*, University of Hamburg, Computer Science Department, 1990.
- Roberts & Goldstein 77:** The FRL Manual, R.E. Roberts, I.P. Goldstein, AI Memo 409 Edition, MIT Lab., 1977.
- Symbolics 88:** Handbooks of Symbolics Programming Environment, 7A, Programming the User Interface – Concepts, Symbolics Inc., 1988.
- Szekely & Myers 89:** A User Interface Toolkit Based on Graphical Objects and Constraints, P.A. Szekely, B.A. Myers, *ACM Sigplan Notices* **24**, 10 (1989), 36–45.
- Visual Languages 90:** Proceedings, *1990 IEEE Workshop on Visual Languages*, Skokie, Illinois, Oct. 4-6, IEEE Computer Society Press, 1990.
- Winston & Horn 89:** LISP, P.H. Winston, B.K.P. Horn, 3rd edition, Addison-Wesley, 1989.
- Wright et al. 85:** pluribus: A Visual Programming Environment for Education and Research, S. Wright, W. Feuerzeig, J. Richards, in: 1988 IEEE Workshop on Languages for Automation, Symbiotic and Intelligent Robotics, Univ. of Maryland, College Park, Maryland, Aug. 29–31, IEEE Soc. Press, 1985, pp. 29–31.