

# Inferring Constraints from Multiple Snapshots

DAVID KURLANDER

Microsoft Research

and

STEVEN FEINER

Columbia University

---

Many graphic tasks, such as the manipulation of graphical objects and the construction of user-interface widgets, can be facilitated by geometric constraints. However, the difficulty of specifying constraints by traditional methods forms a barrier to their widespread use. In order to make constraints easier to declare, we have developed a method of specifying constraints implicitly, through multiple examples. Snapshots are taken of an initial scene configuration, and one or more additional snapshots are taken after the scene has been edited into other valid configurations. The constraints that are satisfied in all of the snapshots are then applied to the scene objects. We discuss an efficient algorithm for inferring constraints from multiple snapshots. The algorithm has been incorporated into the Chimera editor, and several examples of its use are discussed.

Categories and Subject Descriptors: D.2.2 [Software Engineering]: Tools and Techniques—*user interfaces*; I.2.6 [Artificial Intelligence]: Learning—*concept learning*; I.3.6 [Computer Graphics]: Methodology and Techniques—*interaction techniques*

General Terms: Algorithms, Design, Experimentation, Theory

Additional Key Words and Phrases: Constraints, empirical learning, graphical editing

---

## 1. INTRODUCTION

Geometric constraints are used extensively in computer graphics in the specification of relationships between graphical objects [Sutherland 1963a; Borning 1979; Myers 1988; Olsen and Allan 1990]. They are useful during object construction for positioning components relative to one another precisely, as well as during subsequent manipulation of the components. Several graphical techniques, such as grids, snap-dragging [Bier and Stone 1986], and automatic beautification [Pavlidis and Van Wyk 1985], were developed to

---

This research was performed while David Kurlander was a doctoral student at Columbia University. Initial development of Chimera was facilitated by an equipment grant from Hewlett-Packard. David Kurlander's research at Columbia was funded by a grant from IBM.

Authors' addresses: D. Kurlander, Microsoft Research, 1 Microsoft Way, Redmond, WA 98052; S. Feiner, Department of Computer Science, Columbia University, New York, NY 10027.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1993 ACM 0730-0301/93/1000-0277 \$01.50

make the initial construction phase easier, since specifying constraints explicitly can be a complex task. However, when objects are to be manipulated frequently, permanent constraints have an advantage over these other techniques in that they need not be reapplied. Permanent constraints can be particularly useful when subsequent editing of a scene is required, in constructing parameterized shapes that can be added to a library, in specifying how the components of a window should change when the window is resized, or in building user-interface widgets by demonstration.

We introduce a technique for inferring geometric constraints from multiple examples, replacing the traditional constraint specification process with another that is often simpler and more intuitive. Initially, the designer draws a configuration in which all constraints are satisfied and presses a button to take a snapshot. A large number of possible constraints are inferred automatically. Subsequently, if the scene is modified and other snapshots are taken, previously inferred constraints are generalized or eliminated so that each snapshot is a valid solution of the constraint system. For example, we can define two objects to be squares, constrained to maintain the same proportional sizes, by taking a snapshot of two squares, scaling them by equal amounts, and taking another snapshot. Then, if the length of one of the square's sides is changed, the lengths of its other sides and the sides of the second square are updated automatically. The designer need not have a mental model of all of the constraints that must hold, and can test the results by manipulating the scene objects.

Furthermore, the designer may take snapshots at any time. If after one or more snapshots a set of graphical objects does not transform as expected or if the constraint solver cannot reconcile all inferred constraints simultaneously, the graphical primitives can be manipulated into a new configuration with constraints turned off, and a new snapshot can be taken. The incorrectly inferred constraint set is automatically modified so that the new snapshot is a valid constraint solution.

There are a number of problems with traditional constraint specification that this new technique attempts to address:

- Often many constraints must be specified.* Complex geometric scenes contain many degrees of freedom, and often most of these need to be constrained. It can be tedious to express large numbers of constraints explicitly.
- Geometric constraints can be difficult to determine and articulate.* Using constraints requires specialized geometric skills and the ability to articulate about geometric relationships. For example, people who are asked to define a square often describe it as a rectangle with four equal sides. This definition is incomplete, since it neglects the 90-degree-angle constraint. Yet ask them to draw a square, and they typically get it right. Traditional constraint-based drawing systems may not use the appropriate language or abstractions for expressing geometric relationships.
- Debugging, editing, and refining constraint networks are complex tasks.* When incorrect or contradictory constraints are specified, the designer

needs to debug the constraint network, which can be a cumbersome process. To support the debugging task, a visual representation is usually provided for constraints. WYSIWYG editors need a special mode for displaying constraints or need support for multiple views. When constraints and graphical objects are presented together, the scene becomes cluttered if more than a few constraints are displayed simultaneously.

Many approaches have been taken to solve these limitations. The first problem was addressed by Lee, who built a system to construct a set of constraint equations automatically for a database of geometric shapes [Lee 1983]. In doing so, he worked with a restricted class of mutually orthogonal constraints and required that the geometric shapes be aligned with the coordinate axes. Lee's problem domain and assumptions restricted the set of constraints such that there was never any ambiguity about which to select. In our domain the initial ambiguity is unavoidable, and we rely on multiple examples to converge to the desired constraint set.

Systems like Sketchpad [Sutherland 1963a; 1963b] and ThingLab [Borning 1979] make it easier to add large numbers of constraints to a scene, by allowing users to define new *classes* of objects that include the constraints that operate on them as part of the definition. When users create instances of a new class, the system automatically generates the associated constraints. However, people defining a new object class must still instantiate all of the constraints to include in their class definition or prototype. Constraints from multiple snapshots can help with this task.

One of the innovations of Myers's Peridot [Myers and Buxton 1986; Myers 1988] is a component that infers constraints automatically as objects are added to the scene. A rule base determines which relationships are sought, and when a match is found, the user is asked to confirm or deny the constraint explicitly. This reduces much of the difficulty inherent in choosing constraints; the designer is prompted with likely choices. Peridot's geometric inferencing component is limited to objects that can be represented geometrically as boxes aligned with the coordinate axes. The Chimera editor contains a constraint-based search and replace component, that infers general geometric constraints from a static scene according to user-defined rules [Kurlander and Feiner 1992]. However, a single example often contains insufficient information to infer all desired constraints. This paper describes another component of Chimera that uses multiple examples to support the constraint inferencing process.

Maulsby's Metamouse [Maulsby et al. 1989] induces graphical procedures by example and infers constraints to be solved at every program step. To make the task more tractable, he considers only touch constraints in the vicinity of an iconic turtle that the user teaches to perform the desired task. These constraints are treated as postconditions for learned procedural steps, and not as permanent scene constraints. Complex relationships between scene objects can be expressed through procedural constructions, but the relationships between objects in these constructions tend to be unidirectional, and procedures for every dependency need to be demonstrated.

The difficulty inherent in understanding interactions among multiple constraints and debugging large constraint networks has been addressed by the snap-dragging interaction technique [Bier and Stone 1986; Bier 1988] and by an automatic illustration beautifier [Pavlidis and Van Wyk 1985]. In snap-dragging, individual constraint solutions are isolated temporally from one another, so that their interaction cannot confuse the artist. The automatic beautifier infers a set of constraints sufficient to neaten a drawing, but the constraints are solved once and discarded; they are isolated temporally from subsequent user interaction. In the approach described here, constraints can interfere with one another when a new solution is computed. However, the conflicting constraints can be removed by taking additional snapshots.

A number of systems provide visual representations of constraints to facilitate debugging. Sutherland's Sketchpad connected constrained vertices together with lines accompanied by a symbol indicating the constraint. Nelson's Juno, a two-view graphical editor, provided a program view of constraints [Nelson 1985]. Peridot communicated constraints as English language fragments during confirmation, and Metamouse used buttons for confirming and prioritizing constraints. The OPUS interface editor represented constraints between interface components as arrows connecting hierarchical frames or drafting lines [Hudson and Mohamed 1990]. Our technique never requires that its users work with individual, low-level constraints. In both the specification and debugging stages, they can think entirely in terms of acceptable configurations of the illustration. The inferred constraints can be tested by manipulating scene objects, and the constraint set can be refined through additional snapshots. For those who prefer a more direct interface for verifying the inferred constraint set, we provide a browser that displays constraints in a Sketchpad-like fashion. Because our technique is particularly useful in heavily constrained systems, we allow constraints in the browser to be filtered by type or object reference.

One of Borning's ThingLab implementations allowed new types of constraints to be defined and viewed graphically [Borning 1986]. Several systems permit users to define new classes of constraints by filling in cell equations in a spreadsheet [Lewis 1990; Hudson 1990; Myers 1991]. The technique introduced here infers constraints from a fixed set of classes that have proved useful for graphical editing. The inference mechanism determines constants in the constraint equations, but does not synthesize new classes of equations.

Our technique is an application of learning from multiple examples, also known as *empirical learning*. Several empirical learning systems are discussed in [Cohen and Feigenbaum 1982]. In contrast, generalizing from a single example is called *explanation-based learning* and is surveyed in [Ellman 1989]. Explanation-based learning requires a potentially large amount of domain knowledge to determine why one explanation is particularly likely. As we illustrate in subsequent examples, there are often few or no contextual clues in a static picture indicating that one set of constraints is more likely than the next, so we felt the empirical approach was warranted. Empirical learning algorithms have been extensively studied by the artificial intelligence community, but we developed our own to take advantage of

certain features of the problem domain and to make learning from multiple examples a feasible approach to geometric constraint specification.

We have implemented this technique as part of Chimera, a multimodal editor with support for editing graphics, interfaces, and text [Kurlander 1993]. Constraints can be inferred on both graphical and interface primitives. Our initial experience suggests that the snapshot approach, like declarative constraint specification, has its own set of strengths and weaknesses. These will be discussed later in the paper.

In Section 2 we illustrate the user's view of constraint specification with a number of examples. We provide a detailed description of our algorithm in Section 3. In Section 4 we discuss implementation details. Finally, we mention limitations of the approach, present our conclusions, and discuss future work in Section 5.

## 2. EXAMPLES

In this section we show three examples of how constraints are inferred from multiple examples within the graphics and interface editing modes of Chimera. To facilitate the initial construction of the scene, Chimera provides both grids and snap-dragging alignment lines. Chimera has fixed square grids that can be turned off if they interfere with the drawing process. Alignment lines facilitate the establishment of geometric relationships that cannot be expressed with these grids. All figures in this paper were generated directly from Chimera's PostScript output.

### 2.1 Rhombus and Line

Suppose that we would like to add permanent constraints to the rhombus in Figure 1a, so that during subsequent graphical editing it will remain a rhombus, its horizontally aligned vertices will be fixed in space, and the nearby line will remain horizontal, of fixed length, to the right, and at the same  $Y$  position as the bottom vertex of the rhombus. After the initial scene is constructed in Figure 1a, the user presses the *snapshot* button in the editor's control panel. Next, the user translates the top and bottom vertices of the rhombus to make it taller, and translates the horizontal line to the same  $Y$  coordinate as the rhombus's bottom vertex, but to a different  $X$  so that its position will not be absolutely constrained in  $X$  with respect to the bottom vertex. The user presses the snapshot button once more. Figure 1b shows the second snapshot.

Initially, constraints were turned off. Now when the user turns them on from the control panel and edits the scene, the constraints inferred from the snapshot are maintained by the editor. In Figure 2a the user has selected the horizontal line and moved it upward. The top and bottom vertices of the rhombus automatically move so that the demonstrated constraints are maintained. When, in Figure 2b, the top joint of the rhombus is selected and translated to a higher grid location, the bottom rhombus vertex and the horizontal line both move appropriately.

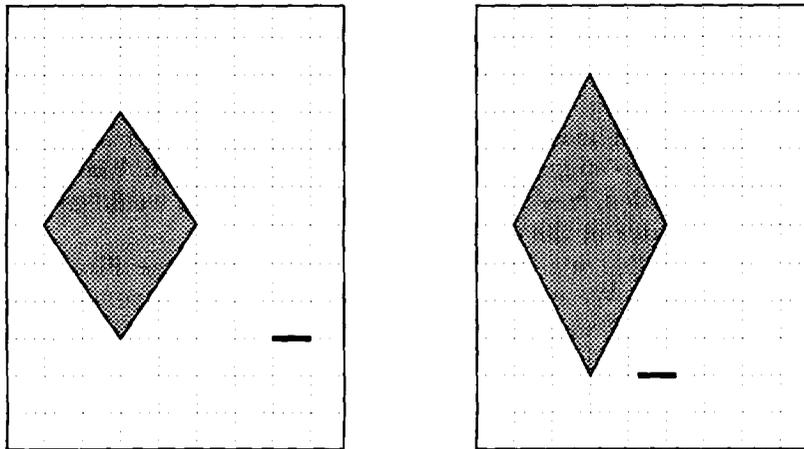


Fig. 1. Two snapshots of a rhombus and a line.

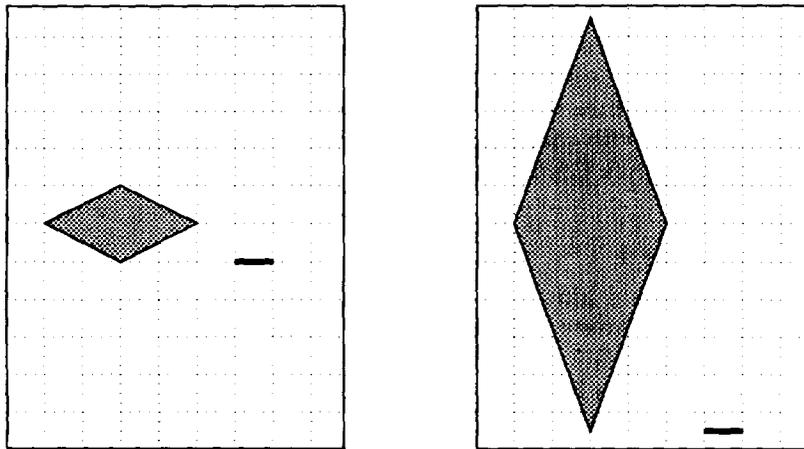


Fig. 2. Two constrained solutions to the snapshots in Figure 1.

The above-mentioned constraints are all specified implicitly, without the user having to express intent in low-level geometric terms. Inferring this information from a single example would be problematic, since it is not clear how to distinguish between those parameters that should be fixed (such as two of the rhombus's vertices, and the length and slope of the horizontal line) and those that should be allowed to vary (such as the length of the rhombus's sides and the locations of the horizontal line's vertices).

One might expect that a sophisticated understanding of the inferencing mechanism is needed to provide the right set of snapshots, but this is untrue. In the second snapshot, the user foresaw the need to move the horizontal line in  $X$ , relative to the bottom vertex of the rhombus, to allow it to move this

way during subsequent interaction. However, if the user neglected to think of this, the constraint solver would disallow such configurations during later manipulations of the scene. The user could then turn off constraints and provide as an extra example the configuration he or she tried to achieve, but could not. This third snapshot would automatically remove constraints that originally prohibited this configuration, without the user explicitly naming them.

In part because this is a highly constrained illustration, few editor operations are necessary to establish the necessary constraints using snapshots. In traditional constraint specification, the user starts with a clean slate and adds all of the intended constraints to the illustration. The snapshots technique takes a very different approach. It initially assumes that all constraints are present in the initial snapshot, and with additional snapshots the user prunes away undesired constraints. This approach is subtractive rather than additive and works best for heavily constrained scenes in which few constraints must be removed. In contrast, traditional declarative specification typically becomes more difficult as more constraints must be added to a scene. The two techniques complement one another. When only a few constraints must be instantiated, it is typically easier to use the traditional declarative approach. Having both forms of specification available allows each technique to be used in cases where it works best, so Chimera's interface supports both.

## 2.2 Resizing a Window

Constraints are useful in constructing user interfaces because they allow the attributes of one interface object to be defined in terms of the attributes of others. For example, when a window is resized, the position and size of the contents may change. Figure 3a shows a window that we have constructed in Chimera's interface editing facility, containing an application canvas (the darkly shaded rectangle), a scrollbar, and three buttons that invoke menus. After positioning these widgets within the parent window, the user presses the snapshot button. The components of the window are then shifted into another configuration, shown in Figure 3b, and a second snapshot is taken. Precise positioning in these snapshots was achieved by using a combination of grids and snap-dragging.

The user's intentions are that the buttons be a fixed distance above the bottom of the window, that the left side of the **Basics** button be a constant distance from the window's left, that the right side of the **File Ops** button be a constant distance from the window's right, and that the **Transformations** button be evenly spaced between the inner sides of the two other buttons. The scrollbar's dimensions are intended to be fixed by the top and right sides of the window, the top of the buttons, and by its constant width. The application canvas should be fixed relative to the left and top of the window, the top of the buttons, and the left side of the scrollbar. Now, when we turn on constraints and select the upper right corner of the window (while the lower left corner is fixed), the window and its contents reshape, as shown in Figure 3c.

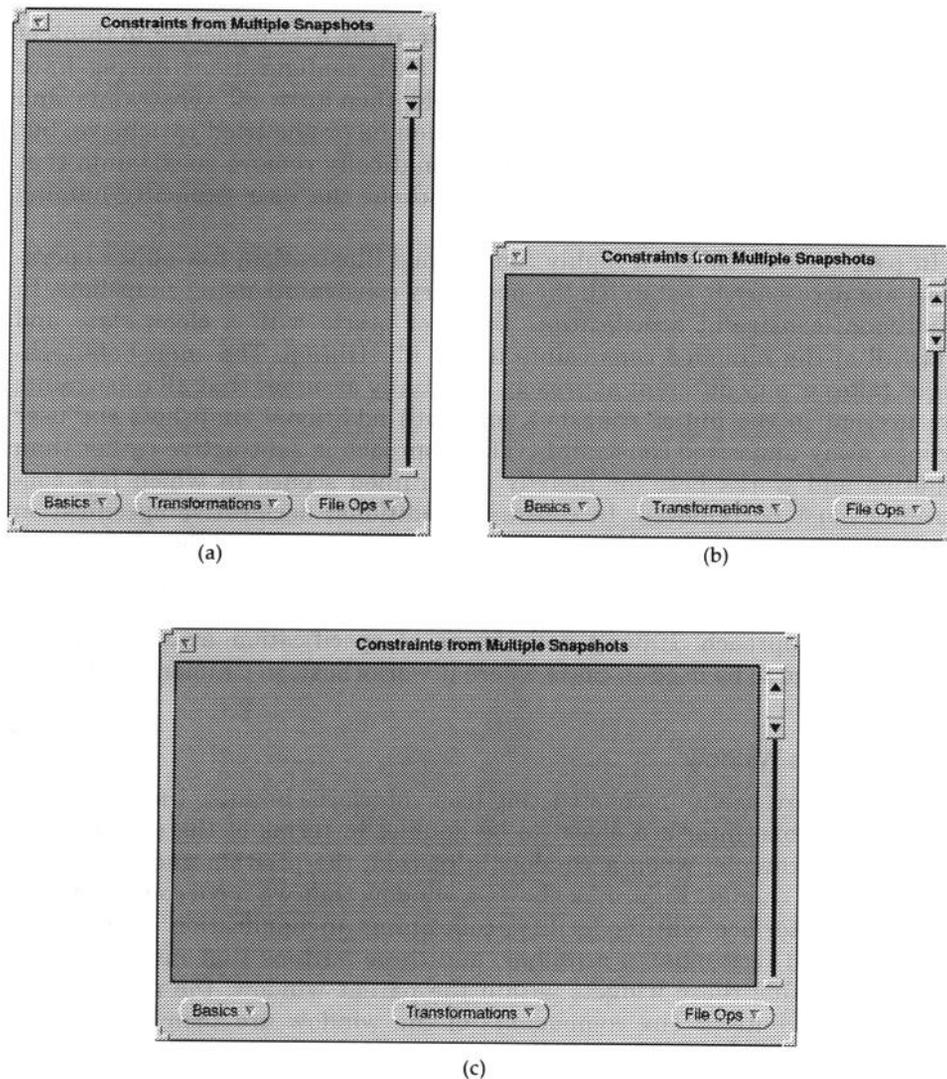


Fig. 3. Specifying window resizing constraints: (a) and (b) are the two snapshots; (c) was produced by dragging the upper right corner of the window.

### 2.3 Constraining a Luxo™ Lamp

This final example applies to both graphical editing and user-interface construction. We would like to constrain a 2-D illustration of a Luxo lamp, so that it behaves like a Luxo lamp. In particular, we want the various pieces to remain connected, the base to be fixed at its initial location, and the arms of the lamp to remain a constant length. Other constraints are important as well, but instead of determining which are significant ourselves, we would prefer to edit the lamp into a number of valid configurations and take snapshots. To control the direction of the lamp's beam, we have built a simple

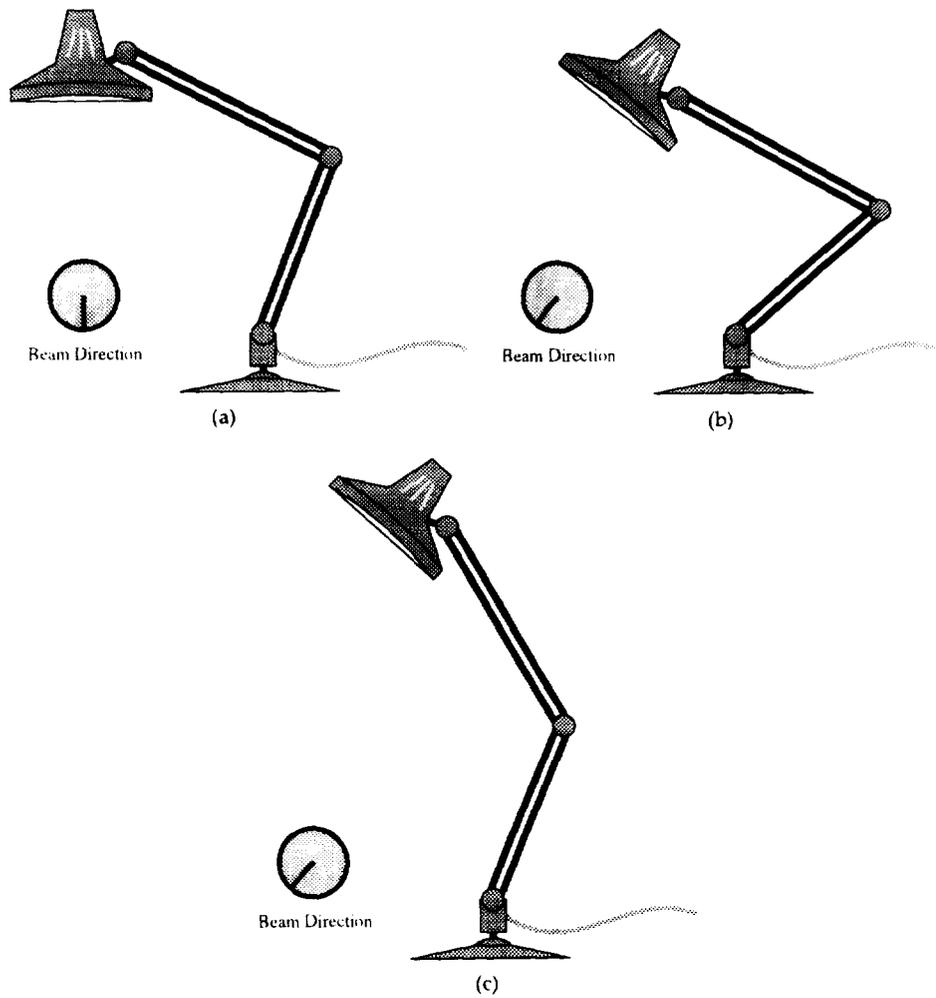


Fig. 4. Teaching Luxo constraints: Three snapshots of valid configurations, provided as input.

dial widget out of a circle and line, and we specify the behavior of the dial relative to the Luxo lamp by demonstration as well. Figures 4a and b show the initial two snapshots of valid configurations of our illustration. Note that the constraints inferred from these two snapshots are independent of the particular editing operations chosen, as explained in Section 3.

After taking the first two snapshots, we turn on constraints and try to manipulate the Luxo lamp, but the constraint solver indicates that it cannot solve the system. The source of the problem is an *incidental constraint*, that is, a constraint that was evident in the first two illustrations, yet was not an intended relationship. When incidental constraints interfere with a desired configuration, they can be removed by manipulating the scene into the new configuration with constraints turned off and by taking an additional snapshot. We could determine which incidental constraint(s) occur in the scene by

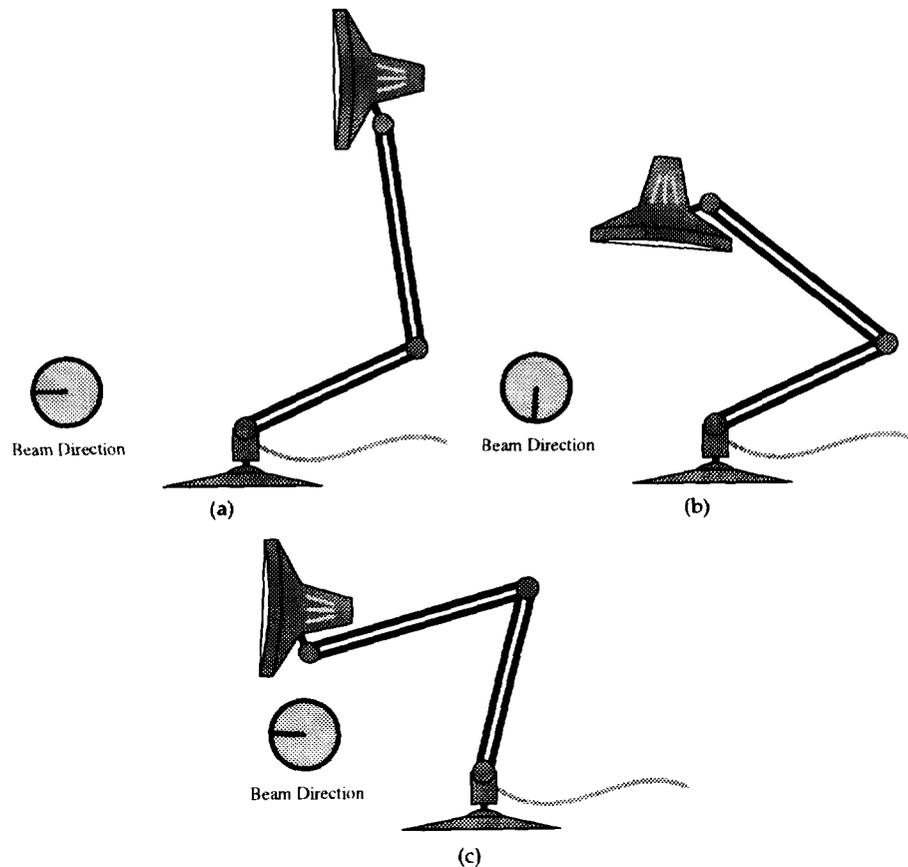


Fig. 5. Luxo on his own: Configurations created by manipulating the uppermost joint and the beam direction dial.

cycling through the visual representation of all constraints and explicitly deleting the undesired ones. However, this can be time consuming when a scene contains many constraints, and it requires that the end user understand the constraint composition of the scene. Fortunately, there is never a need to identify and cull unwanted relationships. While manipulating the scene, if users find that unwanted constraints prohibit a desired, valid configuration, they can turn off constraints, set up this configuration by hand, and take another snapshot. This additional snapshot removes all constraints prohibiting the new configuration. One need not be clever about conveying only the desired constraints in the first two snapshots. Refining a constraint set using snapshots, as with declarative specification, can be an incremental process.

Without identifying the incidental constraint, we set up a configuration that this constraint forbids. This additional snapshot appears in Figure 4c. Now the various components of the lamp move as we had intended. In Figure 5 we manipulate the lamp into three configurations by moving its top joint

and adjusting the beam direction dial. Note that these two controls are not independent: when the dial is rotated, the arms of the lamp can move during the solution of the constraint equations, since the dial's orientation does not uniquely determine a lamp configuration. We can temporarily place a declarative constraint on the joint if we want to change only the beam direction while keeping the arm fixed.

### 3. THE ALGORITHM

In this section we discuss the set of constraints that our system infers. Then we present an efficient algorithm for inferring these constraints and demonstrate the algorithm on a simple example. Next we analyze the algorithm and discuss how parameters can be inferred.

#### 3.1 The Constraint Set

All objects in the Chimera editor are defined geometrically in terms of vertices, and constraints fix the relationships between these vertices. Based on a finite set of example scenes, an infinite number of arbitrary constraints can be inferred. Hence, we have chosen to infer a fixed set of geometric relationships that have proved particularly useful in graphical editors.

Our system infers both absolute and relative geometric constraints. *Absolute constraints* fix geometric relations to constant values. *Relative constraints* associate geometric relations with one another. For example, an absolute constraint might fix a vertex to be at a particular location or a distance to be a constant scalar. A relative constraint might fix two distances or slopes to be the same. Figure 6 lists constraints supported by the Chimera editor. Chimera can infer these constraints from multiple snapshots, or users can specify them directly. Each relative constraint on the right corresponds to an absolute constraint on the left. The dots represent vertices, and C's in the equations represent arbitrary constants.

The relative slope constraint fixes one slope to be a constant offset from another (when represented in terms of degrees, not  $y/x$  ratio). Each of the relative distance constraints fixes two distances to be proportional to one another. Two of the above constraints subsume two others: The absolute distance constraint between vertices subsumes the coincident vertices constraint, and the relative slope constraint subsumes the absolute angle constraint. Our constraint solver does not explicitly support the subsumed constraints, since it handles the more general relationships. Similarly, the inference component has no support for coincident vertex constraints, though it does track absolute angle relationships since the algorithm uses these to find equal angle relationships. Chimera's declarative constraint interface differentiates between all of the constraints in Figure 6, since specifying the more general relationships requires additional input parameters.

Parallel and orthogonal vector relationships are largely captured by the relative slope constraint (which e.g., in the former case would not only fix the vectors between two pairs of two vertices as parallel, but would also constrain their relative directions). Similarly, the relative slope relation captures

**Absolute Constraints**

**Relative Constraints**

1. Fixed vertex location



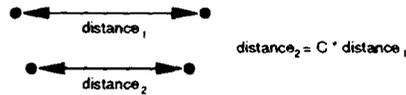
6. Coincident vertices



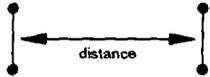
2. Distance between two vertices



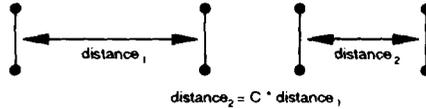
7. Relative distance between two pairs of vertices



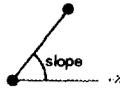
3. Distance between parallel lines



8. Relative distance between two pairs of parallel lines



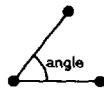
4. Slope between two vertices



9. Relative slope between two pairs of two vertices



5. Angle defined by three vertices



10. Equality between two angles, each defined by three vertices



Fig. 6. Constraints in Chimera.

collinearity, with an additional ordering on the vertices. The algorithm discussed in this section finds all of the constraints in Figure 6 that hold over a sequence of snapshots. Many higher-level constraints can be formed by the composition of these lower-level constraints and, thus, are also inferred by the algorithm. For example, the constraint that one box be centered within another is captured by two relative distance constraints between parallel lines.

**3.2 Algorithm Description**

In the rest of this section, we describe the algorithm that infers these constraints. An overview of its steps is given in Figure 7. It may be helpful to refer back to this figure during the subsequent discussion.

With the first snapshot, the scene is entirely constrained, and each subsequent snapshot acts to reduce or generalize the constraints on the system. If we were to represent explicitly each constraint that could hold at any one time, the space and time costs would be prohibitive. Instead, we economically represent similar constraints over sets of vertices as *groups*. For example,

```

IF first snapshot THEN
  add vertices to initial transformational group
ELSE BEGIN
  split transformational groups to form new child groups
  identify intragroup constraints of child transformational groups
  identify intergroup group-to-group constraints due to splitting transformational groups
  identify intergroup vertex-to-vertex constraints
  break previously instantiated constraints that have been violated
  form delta-value groups from broken absolute constraints
  make a copy of the constraints with redundancies filtered out for the solver
END;
```

Fig. 7. Steps of the inferencing algorithm.

after the initial snapshot, all vertices are constrained to a set location, and the distance and slope between each pair of vertices are fixed, as is the angle between each set of three vertices. Although we could instantiate each of these constraints explicitly, it is far more efficient to represent the vertices as a group with a tag indicating the relationships that hold among all of its members. As will be discussed later, groups can also accelerate the process of determining which constraints hold over a series of snapshots, and can ultimately reduce the number of constraint equations that are passed to the solver.

**3.2.1 Transformational Groups.** The most important type of group in our inferencing mechanism is the *transformational group*. A transformational group contains a set of vertices that have always been transformed together since the first snapshot. At the first snapshot, the algorithm places all of the vertices into a fixed location transformational group, since their positions are initially constrained to be fixed. As vertices are transformed, our undo mechanism keeps track of the sets of vertices selected and the transformation applied, and this information is used by the inferencing mechanism to fragment existing transformational groups into smaller ones. The transformations that can be applied in our system currently include translations, rotations, and isotropic scales, although we plan to extend this algorithm to work with any affine transformation.

**3.2.2 Intragroup Constraints.** We can efficiently determine *intragroup constraints*, that is, constraints that hold within a given transformational group. Figure 8 shows various affine transformations and the geometric relationships that they preserve.<sup>1</sup> The transformation listed in the top half of each box maintains the relationships listed in the lower half of the box and those relationships in the boxes above it. For example, if a transformational group has only been scaled and translated, the slopes, angles, ratios of distances, and parallel relationships are all maintained. By tracking the

<sup>1</sup> Note that *scale* in this diagram refers to isotropic scale, and the vector relationship is the combination of slope and distance constraints.

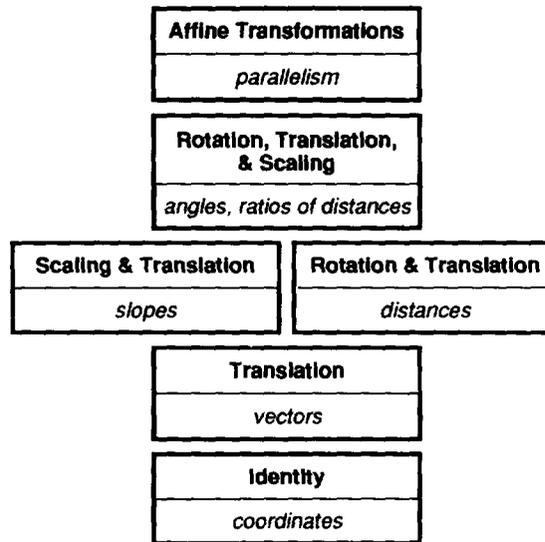


Fig. 8. Transformations and the geometric relationships that they maintain. Reprinted with permission from Bier and Stone [1986].

transformations that have been applied to a transformational group, we determine which constraints must hold within the group without examining its individual vertices.

We next determine which constraints *cannot* hold within the transformational group. Again, this is easily done by examining the transformations that have been applied to the group. If a group has been translated, all of its fixed location constraints are broken. Fixed location constraints are also broken among vertices during rotations and scales if the vertices are not at the center of the transformation. Scales break all fixed distance relationships within a transformational group, and rotations break all fixed slope relationships within a transformational group.

After determining which relationships *must* hold within a group and which *cannot* hold, we must consider the relationships that *might* hold. For each of these constraint relationships, we must examine the vertices in the group, looking for invariant relationships. Fortunately, this expensive task need not be done for the most common transformations, translations, rotations, and isotropic scales, since all relationships in our constraint set can immediately be classified as either definitely present or definitely not present. If we were to extend this algorithm to other less common affine transformations, then the vertices would need to be examined explicitly.

For every snapshot after the initial one, the first step fragments existing transformational groups into new ones, accounting for the transformations that have occurred since the last snapshot. Each new child group has all the constraints of its parent, except those broken by the transformations performed since the last snapshot. Since we are only interested in effective

transformations at the snapshot granularity level, we factor the composition of transformations applied since the last snapshot into scale, rotation, and translation components, and use these, as described above, in determining which intragroup constraints were broken. This allows us to ignore transformations that have been undone by subsequent operations between the two snapshots. For example, if a set of vertices is translated away from its original location and then back again between snapshots, then those translations are effectively ignored.

To illustrate transformational groups and several other algorithmic details discussed later, consider the two simple snapshots given in Figure 9. Two boxes were captured in the first snapshot (Figure 9a). Initially, all vertices were in the same transformational group and were constrained to have fixed locations. After this, but prior to the next snapshot (Figure 9b), the boxes were both scaled by a factor of 2, and the right box was translated to the left, one large grid unit from the left box. Taking the second snapshot caused the system to split the original transformational group into two children, each containing the vertices of one of the boxes. The second snapshot broke the fixed location constraints for all vertices except the bottom left of the left rectangle, since this vertex's effective transformation had no translational component and its location was at the center of the scale. The intragroup absolute distance constraints were broken for each group because there was a net scale, but isotropic scales maintain proportional distances, so an implicit relative distance constraint was added to each group. It is important to note that transformational groups are dependent on the transformations performed, but they have no impact on the constraint set that will eventually be inferred. They accelerate the search process by pruning the search space.

**3.2.3 Intergroup Constraints.** The next step is to compute *intergroup constraints*: constraints between different transformational groups or their vertices. These constraints are generated in several ways. They can be formed from a relative intragroup constraint when a transformational group is split by a transformation that preserves the relation. Consider a transformational group with a relative slope constraint among all of its vertices. If the group is split in two by a translation or scale, then we must add a relative slope constraint between the two groups, relating the slopes contained within one group to the slopes within the other. Similarly, if a transformational group has a relative distance constraint among all of its vertices and the group is split by a translation or rotation, then we need to add a relative distance constraint between the two groups, specifying that the distances within one group will remain proportional to the distances in the other.

We have just described *group-to-group* intergroup constraints: constraints that make entire groups rotate or scale with another. There are also *vertex-to-vertex* intergroup constraints, which express relationships between a small number of vertices. Finding these is the most costly step in our algorithm, but the cost is reduced since we only need to compute intergroup constraints between a child transformational group, its parent, and its siblings (other child groups of the same parent spawned during the same snapshot). Inter-

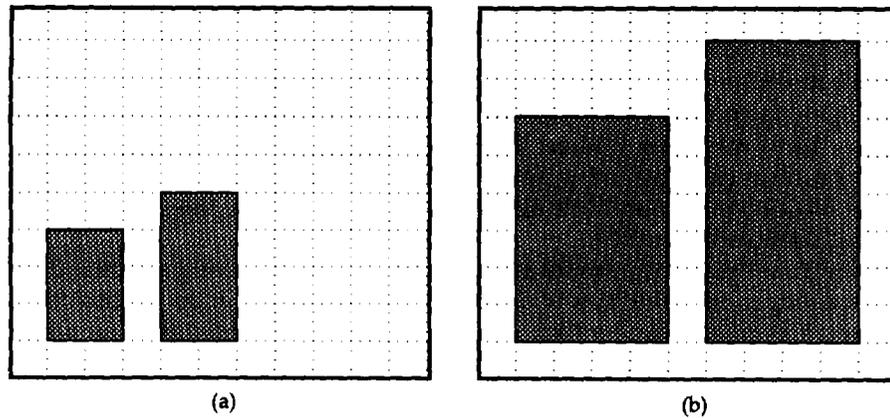


Fig. 9. Two snapshots of a simple scene.

group constraints between the child and other groups were already formed when their ancestors were split.

For small sets of vertices containing members both from the newly created child group and from its parent or siblings, we look for relationships that have not changed and generate absolute constraints for these when found. For example, we compute the slope and distance between such pairs of vertices in the current snapshot and the previous snapshot. If either of these values are unchanged, we create an absolute constraint between the two vertices. Similarly, for each pair of lines constrained to have the same slope, which were contained in a single transformational group during the last snapshot but are now split among groups, we identify absolute parallel distance constraints.

In our rectangle example, an intergroup vertex-to-vertex constraint inferred at the second snapshot declares that the inner segments be one large grid unit apart. This constraint was implicit after the first snapshot, when both rectangles were members of the same transformational group, but must be made explicit after the second snapshot since the relationship still holds after the transformational group was split.

**3.2.4 Delta-Value Groups.** Existing constraints between groups or vertices transformed since the last snapshot are now considered, and those that no longer hold are broken. Broken *relative* constraints, constraints relating geometric measures (such as slope) of more than one object, are split if possible into constraints that are still satisfied among fewer objects. Absolute constraints that have been broken during the current snapshot are matched, as is now described, to form new relative constraints.

We have already described how absolute intergroup constraints are found by locating relationships that do not change. One type of relative intergroup constraint is found by locating relationships that change together. If two pairs of vertices are constrained to have constant slopes, then there is no

need for a relative constraint between the two, since the individual values are fixed. However, if these slopes now change by the same amount, it becomes necessary to create a relative constraint between them. Collections of relations that were absolutely constrained in a previous snapshot, but have broken by equal amounts in the current snapshot, are bundled together into *delta-value groups*.

Delta-value groups, like transformational groups, allow us to represent similar constraints among many objects compactly, but otherwise, they are unrelated. Delta-value groups are simply relative constraints between arbitrary numbers of relations. For example, a delta-value group might constrain  $n$  distances to be proportional. However, when passing the delta-value group to the solver, it need only be expanded to  $n - 1$  binary constraints when solving the system (relating the first element to each subsequent element) rather than to  $n^2$  constraints relating each pair of elements.

Every absolute constraint broken in the current snapshot must be considered for inclusion in a delta-value group. There are three steps in our algorithm where broken absolute constraints are identified:

- (1) during the fragmentation of transformational groups,
- (2) during the identification of vertex-to-vertex intergroup constraints, and
- (3) during the breaking of constraints instantiated during previous snapshots.

We place together in delta-value groups distance relations that change by the same proportion, and angle and slope relations that change by the same number of degrees. Since typically many absolute constraints break during the same snapshot, it is important to find matches efficiently. We employ hashing to match constraints that break by similar amounts, so this step is performed in linear time with respect to the number of broken constraints identified.

Returning to the example in Figure 9, the two rectangles are in separate transformational groups after the second snapshot. This snapshot broke absolute distance constraints for both of these groups, since they were scaled differently than their parent, which had an implicit absolute distance constraint among all of its vertices. Both of these absolute distance constraints broke by a factor of two. As a result, they were added to the same delta-value group, maintaining that distances in the two groups be proportional.

**3.2.5 Redundant Constraints.** We have now computed all of the constraints that are invariant among snapshots. When objects are transformed with constraints turned on, the inferred constraints are passed to our solver. Typically, our constraint set contains a large number of redundant constraints, that is, constraints derivable from others through geometric tautologies. The algorithm finds all constraints from our set that occur in the snapshots, not just the minimal set, though some of the constraints are represented implicitly and efficiently in groups. To accelerate the process of finding a solution to the constraint set, we try to remove redundant constraints. There are two ways that this can be done, both of which involve

looking for simple geometric relationships. The first looks for these relationships as a postprocess after the inferencing has been performed, and filters out extra constraints known to hold in those circumstances. This is the only method currently implemented in Chimera for filtering redundant constraints, and it works well for those relationships that generate a constant number of redundant constraints. However, certain relationships yield a polynomial number of such constraints, and it would be more efficient never to generate them.

These redundancies could be avoided by building additional kinds of groups during the inferencing process. As discussed earlier, transformational groups and delta-value groups allow large numbers of graphical relationships to be represented tersely. By identifying relationships that lead to redundant constraints and classifying them as special groups, we could pass only the essential constraints to the solver. Figure 10 illustrates two relationships that would be particularly useful to express as groups since they are common and yield many redundant constraints if fully expanded. In Figure 10a vertices  $p$  and  $q$  are constrained to be coincident. If each other vertex  $r_i$  were part of a separate transformational group, our algorithm would instantiate the constraints  $\text{distance}(p, r_i) = \text{distance}(q, r_i)$  and  $\text{slope}(p, r_i) = \text{slope}(q, r_i)$  for all  $r_i$ . These redundant constraints could be avoided by building *coincident vertex groups* for sets of vertices currently constrained to be coincident. These groups could be used in lieu of their actual vertices while computing intergroup vertex-to-vertex constraints. If vertices in the group are not coincident in a subsequent snapshot, the group would be broken, and the formerly redundant constraints that still hold would then be instantiated.

Another common geometric relationship leading to redundant constraints is shown in Figure 10b. Here, snapshots have resulted in a set of collinear vertices,  $s_i$ , such that each vertex is in a separate transformational group and the slope between each pair of vertices is fixed. Here, only  $n - 1$  constraints are necessary to represent the slope constraints between the  $n$  vertices, but the algorithm identifies constant slope constraints between each pair of vertices,  $s_i$  and  $s_j$ , such that  $i < j$ . By identifying this relationship as a group during the inferencing process, we could avoid generating these redundant constraints.

Currently, we look for only a few classes of redundant constraints, which we filter as a postprocess, and often a large number eludes us. We are working on improving this component of our system.

**3.2.6 Solving the Constraint System.** When constraints are turned on and constrained objects transformed, we compute the effects on other objects in the scene. The constraint system can be viewed as a graph, with the nodes being vertices of scene objects and the arcs being constraints between the vertices. Changes to one disconnected subgraph cannot affect another, since there are no constraints linking them. We find the disconnected subgraphs containing the vertices actively being transformed, by performing a simple graph traversal beginning at these vertices. Constraints that are not a part of any of these subgraphs cannot affect our solution and can be safely ignored.

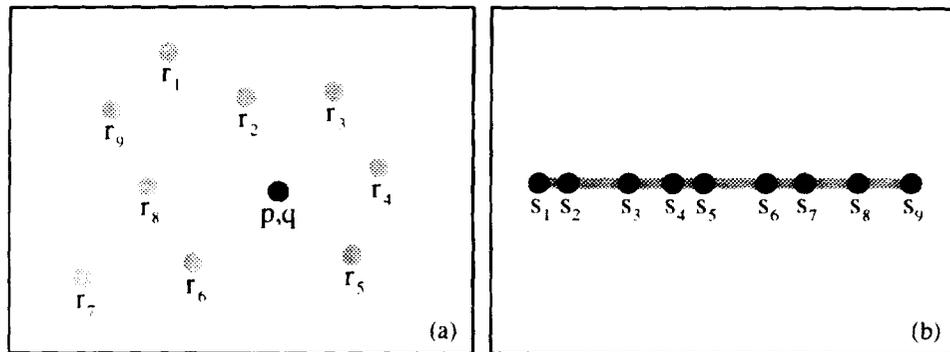


Fig. 10. Two geometric relationships that lead to redundant constraints.

Also, since the constraints of different subgraphs are mutually independent, they are solved independently, thereby reducing the cost of the solution.

We also reduce the solution cost by using a simple generalization of the technique many constraint-based systems use to solve for rigid bodies efficiently. All vertices that are part of the same transformational group are constrained to transform together under a restricted class of transformations, and often we can use this information to avoid passing certain constraints and vertices to the solver. A transformational group that has only been translated has absolute slope and distance constraints between each pair of vertices, and these same constraints ensure that all vertices in the group will translate together. If some vertices in the transformational group participate only in these constraints, then, instead of passing them to the solver, we can explicitly apply to these vertices the translation that the solver finds for other vertices in the group. Similar approaches can be taken for isotropic scales, rotations, and compositions of these transformation classes.

As an example, consider Figure 11. The snapshots in Figures 11a and b constrain the hand to scale so that the lower left vertex of the wrist is fixed and the rightmost vertex of the index finger aligns with the arrow. All vertices of the hand are part of the same scale transformational group, and those of the arrow are part of the same translation transformational group. Figure 11c shows all of the vertices in the system that participate in the constraint solution. However, only a few of these 134 vertices must be passed to the solver.

In Figure 11d, we choose to translate the lower right vertex of the arrow. We begin traversing the constraint graph at this vertex to determine which constraints and vertices must be passed to the solver. This vertex will be passed to the solver, since it is being manipulated directly by the user. The top vertex of the arrow must also be passed to the solver, since it participates in an intergroup slope constraint. These two vertices are bound together by absolute slope and distance constraints because the arrow is a single translation transformational group. The displacement of all of the other vertices in

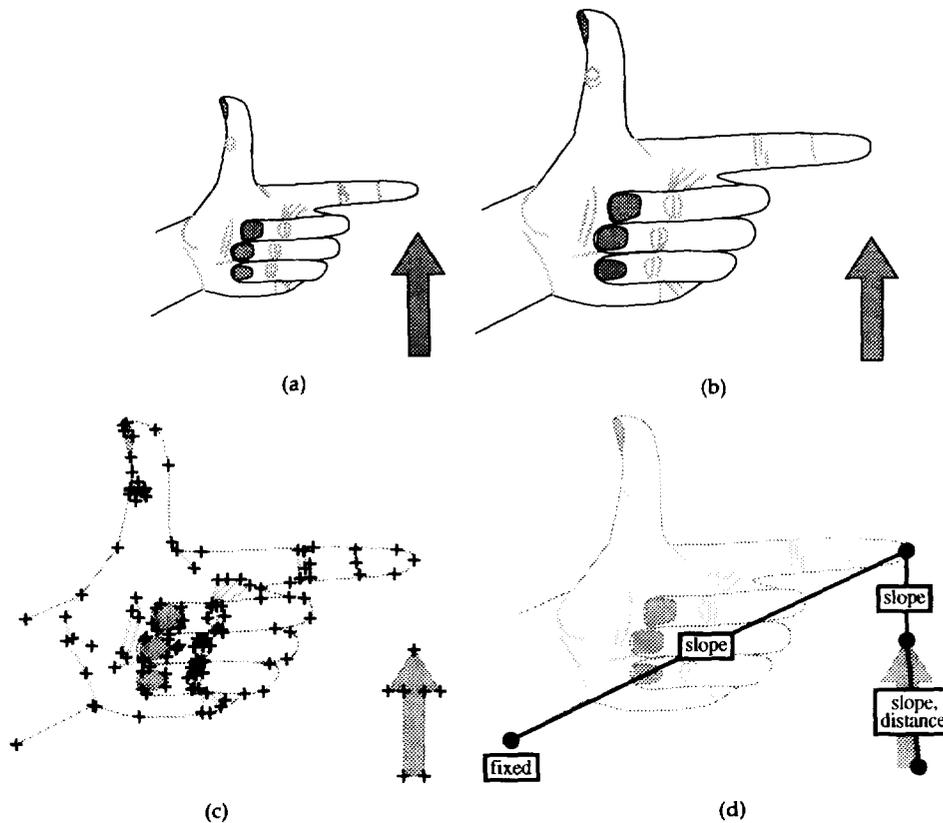


Fig. 11. Efficient constraint formulations for transformational groups: Snapshots (a) and (b) constrain the scene. A naive approach solves constraints for all vertices marked in (c). A more efficient method solves only constraints shown in (d).

this group will be determined by calculating the displacement vector that the solver finds for these points.

Similarly, the vertex at the tip of the index finger is passed to the solver, since it participates in an intergroup slope constraint with the point of the arrow. The lower left vertex of the wrist must also be passed to the solver, since it has a fixed location constraint. These two vertices of the hand are connected by an absolute slope constraint, because they are part of the same scale transformational group. The positions of all of the other vertices in the hand are easily determined by the scale transformation that maps these two vertices to their new positions.

### 3.3 Analyzing the Algorithm

This section summarizes the inferencing algorithm and presents informal arguments for its correctness. The technique described in this paper finds all relationships of the classes listed in Figure 6 that are present in a sequence

of snapshots and instantiates these into constraints. A brute-force algorithm would consider each relationship applied in turn to every collection of vertices of the appropriated size and then determine whether the relationship in fact changes over the course of the snapshots. This would be computationally expensive, so our algorithm takes a different approach. However, to show that it works, it suffices to explain how it finds the same constraint set as the brute-force algorithm.

The snapshot process partitions all vertices in the scene into transformational groups. The set of translations, rotations, and isotropic scales applied to a transformational group since the first snapshot automatically determines which of the relationships in our constraint set hold among its vertices and which do not. This provides the first savings over the brute-force algorithm: It is not necessary to search through collections of vertices in the same transformational group for invariant relationships, since these relationships are completely determined by the group's transformations. However, our algorithm, like the brute-force algorithm, must consider constraints that span multiple transformational groups (intergroup constraints) as well as these constraints that lie in a single group (intragroup constraints). Together, intergroup and intragroup constraints constitute all possible constraints in a scene.

Finding intergroup constraints is more difficult. The constraints of interest to us include *absolute constraints* and *relative constraints*. Absolute constraints express a single geometric relationship to be constant, whereas relative constraints compare multiple geometric relationships. To find *absolute* intergroup constraints, our algorithm does the same thing as the brute-force approach: It considers all collections of vertices of the appropriate number, spanning multiple groups, and looks for absolute relationships unchanged over all of the snapshots. It does this incrementally, as transformational groups are split from their parents, but the effect is the same as seeking these relationships after all of the snapshots are given.

Performing an exhaustive search for *relative* intergroup constraints would be more costly, since they typically involve larger numbers of vertices than absolute constraints. Fortunately, they can be found without resorting to the brute-force approach. Relative relationships are merely pairs of absolute relationships that always change the same way, for example, two distances that always remain proportional. There is no need to create a relative constraint before its absolute components change for the first time, since the absolute components already capture the relative relationships. For example, if two distances are fixed absolutely, then their proportion is implicitly defined. All of the absolute constraints in Figure 6 capture the relative relationships to their right until a snapshot breaks the absolute constraints. Then, since relative relationships are pairs of absolute relationships that always change together and in the same way, the algorithm finds relative intergroup constraints merely by identifying absolute constraints that always change identically. So, instead of performing a costly exhaustive search for relative intergroup constraints, this algorithm monitors all absolute constraints, both absolute intergroup constraints and absolute constraints on

entire groups, and matches those that always change by equal factors (in the case of distances) or degrees (in the case of angles and slopes). In this way, the algorithm finds all inter- and intragroup relationships in the scene, and finds an equivalent set to the brute-force approach with less computation.

The algorithm discussed so far is *not* heuristic; it finds all constraints of the classes in Figure 6 obeyed by a sequence of snapshots. As discussed later, considering all of these relationships often results in a number of incidental constraints, that is, relationships in a snapshot sequence that the user did not intend. To combat this problem, we also experimented with a simple modification of the algorithm that instantiates relative intergroup slope and distance constraints only between two pairs of connected vertices, and absolute and relative intergroup angle constraints only on angles formed by three connected vertices. This removes from consideration some relationships that are usually not significant, yet often yield a large number of incidental constraints.

The complexity of the inferencing algorithm depends on whether we restrict the above constraints to connected vertices. If not, the most expensive step is finding intergroup absolute angle constraints, which is  $O(n^3)$  with the number of vertices. Since vertices in our system connect no more than two lines, the task of searching for these constraints between *connected* vertices is  $O(n)$ . But then the cost of finding intergroup absolute distance and slope constraints between arbitrary vertices is still  $O(n^2)$ , and this becomes the bottleneck of the inferencing component. At this time, our system only removes a few classes of redundant constraints, and we do not know the cost of implementing a good, general redundancy filter.

#### 3.4 Parameterizing an Illustration

Often it is convenient to be able to parameterize graphical illustrations. A slight modification to the algorithm described above allows simple relationships between scene objects and numeric text fields to be inferred during the snapshot process. We provide an Arguments window in which scalar values can be typed as the illustration is edited into new configurations. These values are interpreted by the algorithm as though they were distances, slopes, or angles between vertices. If one of the changing geometric relationships in the scene matches a changing numeric argument, a relative constraint is created between the two values.

In Figure 12 we have drawn a scrollbar in the Chimera editor, and we would like to equate the percentage typed in the Argument 1 field of the Arguments window to the height of the scrollbar's slider. We constrain the scene by providing the two snapshots depicted in Figures 12a and b, but in addition to presenting two valid versions of the scene's geometry, we type corresponding values in the Argument 1 text field. As shown in Figure 12c, after turning constraints on, we can adjust the scrollbar's slider by editing the value in this same text field and pressing the Apply-Arguments button. Alternatively, we can adjust the slider, and the value of Argument 1 changes accordingly.

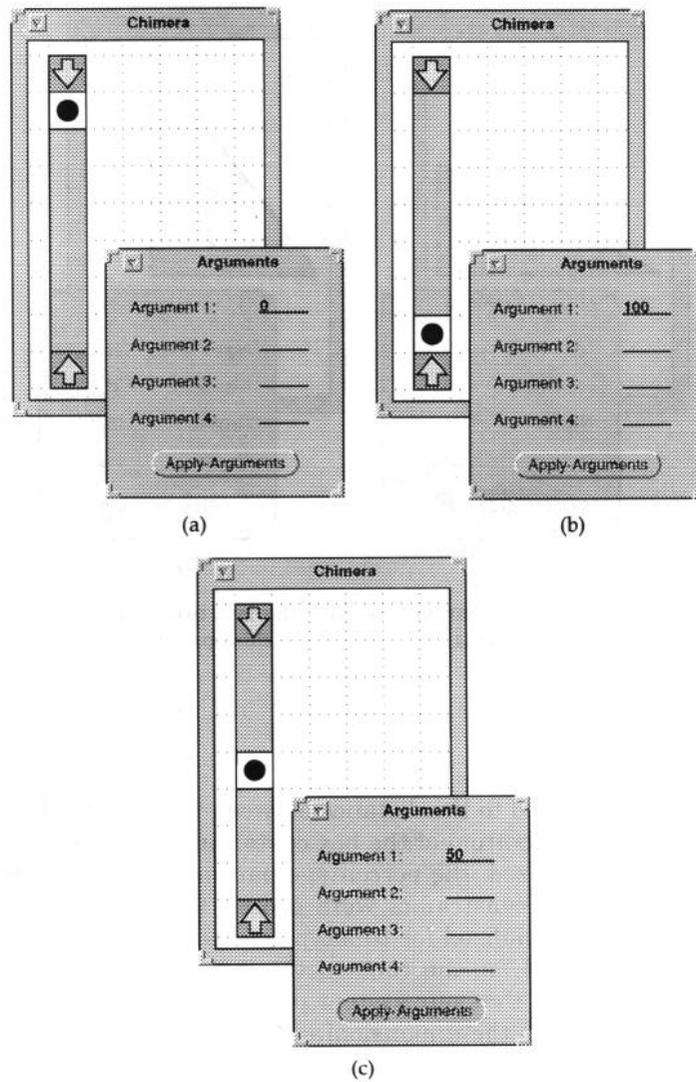


Fig. 12. Dimensioning the height of a scrollbar: Initially two snapshots, (a) and (b), are specified. A new value for Argument 1 is entered in (c), and the scrollbar adjusts automatically.

Myers presents a similar example of parameterizing scrollbar behavior [Myers 1988]. His method linearly interpolates between two different constrained configurations, which is a more powerful abstraction, particularly for defining the behavior of widgets. For example, in Peridot the slider height can be parameterized with respect to the bottom and top of the scrollbar. This cannot currently be done in our system. In our example, Argument 1 is

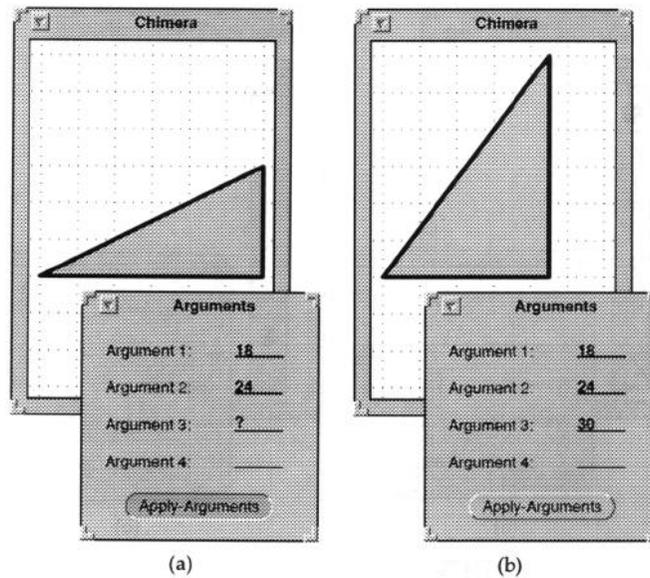


Fig. 13. Specifying a subset of the parameters: Only the first two parameters are specified in (a). The triangle resizes, and a value is computed for the third parameter in (b).

interpreted as proportional to the distance between two parallel lines, the top of the slider box and the bottom of the box containing the upper scroll arrow. So, if the scrollbar is resized, the percentage parameter will no longer range from 0 to 100. Peridot's constraints were chosen for the domain of widget construction and are specialized for this type of task. Our system provides a lower-level constraint set for the construction of general illustrations. The type of parameterization that our system provides is useful for many basic illustration tasks, such as the dimensioning of distances, slopes, and angles.

Since parameters of the illustration can be mutually dependent, the values of a subset may determine the rest. Sometimes, the user may care to set only a few of the available parameters. For these reasons, we allow parameters to be either specified or unspecified. Specified parameters are constrained to their current value during the constraint solution, but unspecified parameters are allowed to vary. Figure 13a shows a Chimera editor scene containing a single triangle. Two previous snapshots (which are not shown) have constrained it to be a right triangle, with a fixed lower left corner and horizontal base. They also have constrained Argument 1 to be proportional to the length of the base, Argument 2 to be proportional to the length of the vertical segment, and Argument 3 to be proportional to the hypotenuse's length. In Figure 13a we type the desired lengths of each of the sides but the hypotenuse into the Arguments window. The question mark entered for

Argument 3 requests that it be chosen by the constraint solver. After the Apply-Arguments button is pressed, the triangle resizes subject to its constraints, and Argument 3 is filled in with a suitable value.

#### 4. IMPLEMENTATION

The Chimera editor is implemented mainly in Lucid Common LISP and CLOS (the Common LISP Object System), with some C code as well. Our constraint solver is implemented in C, but the inferencing mechanism is in LISP. The code runs on Sun workstations under OpenWindows.

We use Levenberg-Marquadt iteration [Press et al. 1988] to solve the constraint systems. This method uses gradient descent when far from a solution, but switches to the inverse Hessian method to converge quadratically when a solution is near. Levenberg-Marquadt is a least-squares method. Each constraint is implemented as an error function, and the algorithm finds the best solution to a set of error functions according to a least-squares evaluation, provided it does not fall into a local minimum. The functions are not limited to be linear or even algebraic. If the constraint solver cannot find an acceptable solution, the user is notified of this, and he or she then has the option of undoing the operation or trying to coax the system out of a local minimum by further manipulating the graphical objects. We would eventually like to add multiple constraint solvers, so that when one fails to find a solution another can be invoked. In systems with multiple correct solutions, which occur occasionally, the iterative solver tends to find the solution closest to its inputs. By manipulating objects in the scene, users can cause the solver to choose a particular solution.

Part of the Levenberg-Marquadt method requires solving a system of equations to determine how the current solution estimate should change. If the error functions are not mutually independent (which is the case with redundant constraints), the system cannot be solved using Gaussian elimination. Instead, we use singular value decomposition [Press et al. 1988] to find a solution at this step.

In looking for absolute and relative relations in the scene, it is important to build tolerances into the matching process. We use small, fixed, empirically derived tolerances, just large enough to account for floating-point inaccuracies during the construction and editing of the scene. If the tolerances were large, the number of incidental constraints would increase. Our small tolerances are on the order of fractions of degrees and millimeters. These small tolerances require that the snapshots be drawn accurately, so Chimera provides both grids and snap-dragging for this purpose.

Both the inferencing algorithm and constraint solver typically run at interactive speeds, on a 1.4 MFLOP Sun SparcStation 1+, for systems of the size presented in this paper. The slowest snapshot (that of Figure 4b) took about 3 seconds. Constraint solutions were obtained in under a second in all cases but the window-resizing example. This took somewhat longer because a large number of redundant constraints were passed to the solver by the

inferencer. Further work on the inferencer should reduce the number of redundant constraints and speed up constraint solutions.

## 5. CONCLUSIONS AND FUTURE WORK

Snapshots appear to be a very intuitive way of specifying constraints and often allow complex constraint systems to be specified with relatively few operations. However, we will not know until performing user trials whether, and under what conditions, people prefer the technique to traditional declarative specification. Our personal experience with snapshot constraint specification in Chimera suggests that it is not a panacea. There are certain tasks for which it appears to be a simpler, more natural method of constraint specification, but for others, traditional declarative specification remains easier. The snapshot approach works well for highly constrained scenes, particularly those that easily can be manipulated into example configurations. Explicit constraint declaration is an additive technique, rather than a subtractive one, and often seems preferable for weakly constrained scenes and those for which setting up snapshots would be difficult. There are a number of problems using snapshots that the traditional method does not have:

- Certain pictures can be difficult to edit into new configurations. In some of these cases, it may be easier to specify constraints explicitly.
- Incidental, unintended relationships often occur in large scenes, necessitating extra snapshots.
- Redundant constraints are commonly passed to the solver, increasing solution costs.

When it is easier to specify constraints declaratively than by example, then the declarative technique should be used. We have built a traditional declarative constraint interface for our editor that is useful in these cases and that will allow us to compare the two methods better.

Incidental constraints can be reduced by restricting the classes of constraints that can be inferred. In our initial implementation, we inferred intergroup relative distance and slope constraints between any two pairs of vertices. This resulted in too many incidental constraints, so we restricted these constraints to pairs of two connected vertices (although there need not be a connection between the pairs). We still infer absolute distance and slope constraints between any two vertices, and intragroup relative distance and slope constraints between any two pairs of vertices. We are looking for additional restrictions that will not significantly impair the utility of the system.

Another way to reduce incidental constraints is to have the user select a set of objects prior to the beginning of a snapshot sequence and to have the inferencer look for constraints only among these objects. Partitioning the scene has the additional benefit of accelerating snapshots. In traditional constraint specification, constraints are also often added in partitions, to speed up solution. Currently, we do not allow inferencing to be restricted to a

subset of the scene, but this option is important for large scenes, and we plan to include it in the future.

One approach to reducing redundant constraints might involve using algorithms similar to those Chyz developed for maintaining complete and consistent constraint systems [Chyz 1985]. When a new constraint is added to the network, his algorithms determine which constraint must be eliminated to avoid overconstraining the system. These methods may allow us to reduce the set of constraints passed to the solver. However, we would not filter out most redundancies from our master constraint set, since after subsequent snapshots they may no longer be redundant.

There are a number of other interesting topics for future work. We would like to extend our system to handle constraints between nongeometric properties such as color or font. Animating the constrained systems would provide an intuitive display of the set of constraints inferred, in the same visual language as the snapshot specification. We would like to provide an audit trail of snapshots by incorporating them into our graphical edit history representation [Kurlander and Feiner 1990]. This will allow individual snapshots to be eliminated and the constraint network recalculated.

It would be helpful to infer a few additional geometric relationships, such as the distance between a vertex and a line. This constraint could be easily added to the inferencing algorithm. Currently, we infer constraints only among vertices in the initial drawing. There are cases when we would also like to infer relationships among implied objects, such as the center of a rotation or the bounding box of an object. We also plan to allow constrained shapes inferred by our technique to be parameterized in more complex ways and included as part of macros.

#### ACKNOWLEDGMENTS

We thank Terry Boulton for valuable advice on numerical techniques, and Larry Koved and Dan Ling for several helpful discussions. Eric Bier suggested very good background material. Eric Bier, Michael Elhadad, Dan Olsen, Ken Perlin, and a crew of anonymous reviewers provided useful comments on earlier drafts.

#### REFERENCES

- BIER, E. A. 1988. Snap-Dragging: Interactive geometric design in two and three dimensions. Ph.D. thesis, Electrical Engineering and Computer Science Dept., Univ. of California, Berkeley, Apr.
- BIER, E. A., AND STONE, M. C. 1986. Snap-Dragging. In *Proceedings of SIGGRAPH '86. Comput. Graph.* 20, 4 (Aug.), 233-240.
- BORNING, A. 1979. ThingLab: A constraint-oriented simulation laboratory. Tech. Rep. SSL-79-3, Xerox Palo Alto Research Center, Palo Alto, Calif. (Revised version of Stanford Ph.D. thesis, July 1979).
- BORNING, A. 1986. Graphically defining new building blocks in ThingLab. *Hum. Comput. Interaction* 2, 4, 269-295. (Reprinted in *Visual Programming Environments: Paradigms and Systems*, E. Glinert, Ed. IEEE Computer Society Press, Los Alamitos, Calif. 1990, pp. 450-469.)

- CHYZ, G. W. 1985. Constraint management for constructive geometry. Master's thesis, Mechanical Engineering Dept., Massachusetts Institute of Technology, Cambridge, Mass., June.
- COHEN, P. R., AND FEIGENBAUM, E. A. 1982. *The Handbook of Artificial Intelligence*. Vol. 3. Kaufmann, Los Altos, Calif.
- ELLMAN, T. 1989. Explanation-based learning: A survey of programs and perspectives. *Comput. Surv. (ACM)* 21, 2 (June), 163-221.
- HUDSON, S. E. 1990. An enhanced spreadsheet model for user interface specification. Tech. Rep., TR 90-33, Dept. of Computer Science, Univ. of Arizona, Tucson, Oct.
- HUDSON, S. E., AND MOHAMED, S. P. 1990. Interactive specification of flexible user interface displays. *ACM Trans. Inf. Syst.* 8, 3 (July), 269-288.
- KURLANDER, D. 1993. Graphical editing by example. Ph.D. thesis, Computer Science Dept., Columbia Univ., New York, July.
- KURLANDER, D., AND FEINER, S. A. 1990. A visual language for browsing, undoing, and redoing graphical interface commands. In *Visual Languages and Visual Programming*, S. K. Chang, Ed. Plenum Press, New York, pp. 257-275.
- KURLANDER, D., AND FEINER, S. 1992. Interactive constraint-based search and replace. In *CHI '92 Conference Proceedings* (Monterey, Calif. May 3-7). ACM, New York, pp. 609-618.
- LEE, K. 1983. Shape optimization of assemblies using geometric properties. Ph.D. thesis, Mechanical Engineering Dept., Massachusetts Institute of Technology, Cambridge, Mass., Dec.
- LEWIS, C. 1990. NoPumpG: Creating interactive graphics with spreadsheet machinery. In *Visual Programming Environments: Paradigms and Systems*, E. Glinert, Ed. IEEE Computer Society Press, Los Alamitos, Calif., pp. 526-546.
- MAULSBY, D. L., WITTEN, I. H., AND KITTLITZ, K. A. 1989. Metamouse: Specifying graphical procedures by example. In *Proceedings of SIGGRAPH '89*. *Comput. Graph.* 23, 4 (July), 127-136.
- MYERS, B. A. 1988. *Creating User Interfaces by Demonstration*. Academic Press, Boston, Mass.
- MYERS, B. A. 1991. Graphical techniques in a spreadsheet for specifying user interfaces. In *CHI '91 Conference Proceedings* (New Orleans, La., Apr. 27-May 2). ACM, New York, pp. 243-249.
- MYERS, B. A., AND BUXTON, W. 1986. Creating highly interactive and graphical user interfaces by demonstration. In *Proceedings of SIGGRAPH '86*. *Comput. Graph.* 20, 4 (Aug.), 249-268.
- NELSON, G. 1985. Juno, a constraint-based graphics system. In *Proceedings of SIGGRAPH '85*. *Comput. Graph.* 19, 3 (July), 235-243.
- OLSEN, D. R., JR., AND ALLAN, K. 1990. Creating interactive techniques by symbolically solving geometric constraints. In *Proceedings of UIST '90* (Snowbird, Utah, Oct. 3-5). pp. 102-107.
- PAVLIDIS, T., AND VAN WYK, C. J. 1985. An automatic beautifier for drawings and illustrations. In *Proceedings of SIGGRAPH '85*. *Comput. Graph.* 19, 3 (July), 225-234.
- PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A., AND VETTERLING, W. T. 1988. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge.
- SUTHERLAND, I. E. 1963. Sketchpad, A man-machine graphical communication system. Ph.D. thesis, Electrical Engineering Dept., MIT, Cambridge, Mass., Jan.
- SUTHERLAND, I. E. 1963. Sketchpad, A man-machine graphical communication system. In *AFIPS Conference Proceedings, Spring Joint Computer Conference*. AFIPS, Arlington, Va., pp. 329-346.

Received May 1991; revised and accepted March 1993