

# Tioga-2: A Direct Manipulation Database Visualization Environment

Alexander Aiken      Jolly Chen      Michael Stonebraker  
Allison Woodruff

Department of Electrical Engineering and Computer Sciences  
University of California, Berkeley\*  
email: tioga@postgres.berkeley.edu

## Abstract

This paper reports on user experience with Tioga, a DBMS-centric visualization tool developed at Berkeley. Based on this experience, we have designed Tioga-2 as a direct manipulation system that is more powerful and much easier to program. A detailed design of the revised system is presented, together with an extensive example of its application.

## 1 Introduction

Database system performance—as measured by either processing speed or the quantity of data that can be managed—has grown by an order of magnitude in recent years, making increasingly sophisticated applications feasible on ever-larger data sets. However, database query languages have changed relatively little and are difficult for non-experts to use. The vast majority of database users are unable to customize applications to their own needs, let alone develop their own custom applications. Thus, at present the expanding capabilities of database systems can be exploited fully only by expert programmers. Making databases easier to use and program, and thereby more accessible, is an important issue today and will become more important as database technology becomes faster, cheaper, and more powerful [SAD<sup>+</sup>93].

This paper reports on the design of Tioga-2, a new database visualization environment. We use the term “visualization environment” rather than “programming environment” to emphasize that most programming operations in Tioga-2 are performed by manipulating graphical representations of either programs or data. Tioga-2 is based on a small set of primitive operations for transforming data and its visualization. These primitives have been chosen carefully to have clear, simple semantics and to be composable. Thus, Tioga-2 users can build sophisticated applications—or modify existing applications—by successive composition of the primitives. We believe that by providing a small set of general “building blocks”, minimum language syntax, and immediate feedback on the effect of incremental program modifications, Tioga-2 will make it much easier for database users to develop database applications.

Tioga-2 has not been designed in a vacuum. Previously, we reported on the design and implementation of Tioga, a visualization system that is coupled closely with the POSTGRES DBMS [SCN<sup>+</sup>93]. The design of Tioga-2 has been influenced heavily by what we learned from user experiences with Tioga and a companion commercial product, Illustra Object-Knowledge, based on the same ideas. In the rest of this introduction, we first discuss the problems and lessons from Tioga and then outline our solution to those problems in Tioga-2.

---

\*This research was sponsored by NSF under grant FD94-00773 and CNRI under grant M1717.

## 1.1 Tioga

Tioga adopts the “boxes-and-arrows” programming paradigm popularized by AVS [U<sup>+</sup>89], Data Explorer [LAC<sup>+</sup>92], and Khoros [RY92]. Every box is a user-defined function, which has been registered with POSTGRES. A programmer constructs a Tioga program using a drag and drop editor to move and connect boxes on the screen. The editor performs appropriate type checking to ensure that each box corresponds to a POSTGRES function and that the data type of the output of every box corresponds to the input data type expected by its successor box.

Every Tioga program has a designated *viewer* connected to the output of a specified box. The viewer provides the user with a two-dimensional *canvas* onto which the programmer places renderable objects. In addition, the viewer provides a *pan* feature whereby the user can “fly over” the canvas viewing areas of interest. Furthermore, the user can *zoom* into areas of the canvas to see more detail. Zoom is a powerful construct, as it supports so-called *drill-down*—the ability to change the visual representation of data. For example, a state map of the United States could become a county map upon suitable zooming. In addition, we specified but never implemented the features of multiple viewers, viewers within viewers, cloning of viewers, slaving of viewers, and wormholes [WWT<sup>+</sup>94].

Experience with Tioga and Illustra Object-Knowledge can be summarized as follows:

### 1. Programmer model

Tioga is based on the idea that an expert programmer constructs POSTGRES user-defined functions (boxes) and that a second programmer uses an editor to “wire up” visualizations. In this way, Tioga implements a “big programmer / little programmer” environment.

It has been sufficiently hard to construct boxes-and-arrows programs that the little programmer must, in fact, be a big programmer. The key problem is that simplifying the specification of control logic through a boxes-and-arrows notation does not simplify programming sufficiently. For example, to construct Tioga applications, the little programmer must understand locating objects on a canvas and turning objects into graphical representations. It turns out that even expert programmers find these tasks tedious, and the little programmer has great difficulty with them. As a result, little programmers have not been able to program in Tioga because it is not nearly easy enough to use.

### 2. Programming environment

Tioga has the familiar notion of building a program, compiling it, and then running the compiled result. Novices have difficulty learning how to program effectively in this paradigm. For example, if nothing appears on the screen, then there is a “bug” in the program. Bugs are hard for the programmer to find because Tioga provides a viewer only for the final result; it is not possible to place a viewer on any edge in a diagram to visualize the data that is flowing along that edge.

### 3. Expressive power

As a result of trying to provide a simple programming model, Tioga is in some ways oversimplified. To select only a single example, because every box must be a user-defined function, a box has a single output, which must be of a specific type. This makes it difficult to implement functionality of the form:

**if** condition **then** deliver data to box *i* **else** deliver data to box *j*

## 1.2 Tioga-2: Guiding Principles

Based on our experiences described in Section 1.1, we are redesigning Tioga completely from scratch, and the result is Tioga-2, described in this paper. We begin with the principles that have guided the redesign.

Much of the problem with the original Tioga system is that there is no way to specify some aspects of a visualization except via ordinary statement- and expression-oriented programming. Learning to write procedural code is a high hurdle for many non-programmers, and some visualization aspects—such as writing functions to position data in a multi-dimensional space—are difficult even for expert programmers.

There is an alternative way to specify data visualizations. Even non-programmers intuitively understand how to specify desired computations “by example”—by manipulating sample data. Instead of writing in a standard programming syntax, the programmer begins with very simple displays of data and composes them directly on the screen to construct elaborate visualizations. In moving from the boxes-and-arrows notation of Tioga to the direct manipulation programming paradigm of Tioga-2, we have identified a number of principles we believe to be important to a usable, flexible, and powerful direct manipulation visualization system:

1. Every result of a user action has a valid visual representation.

All data types constructible by Tioga-2 programs have a well-defined screen representation. As such, the programmer obtains immediate visual feedback of the effect of any change to a Tioga-2 program and can visually inspect intermediate results. This principle facilitates debugging activities and solves problem (2) noted above.

2. Programming is incremental.

We are committed to a system with a small number of simple primitives. Visualizations are constructed incrementally by successive composition of the primitives. Combined with the ability to visualize results of incremental changes immediately, we believe that we can empower the little programmer to construct Tioga-2 programs. In Tioga-2, there is no distinction between constructing a program, modifying an existing program, and using an existing program.

3. To the extent possible, programming is specified visually by direct manipulation of visualized data and graphical representations.

A boxes-and-arrows representation of the user’s program is available and must be used for certain operations. However, we expect considerable programming to be done by direct manipulation of the screen without reference to this data structure.

4. Every operation has a clear, well-specified semantics.

In contrast to many previous direct manipulation systems there is no complex inference procedure to synthesize a program from a user’s examples [Cyp93]. Instead, every Tioga-2 operation has a straightforward, unambiguous meaning as a step in a program.

5. Retention of the “big programmer/little programmer” model.

We recognize that there are computations that cannot be specified in Tioga-2. Thus, it is expected that big programmers will still construct additional Tioga-2 boxes as in the original Tioga system. To solve problem (3) above, we generalize boxes-and-arrows diagrams in several ways. Boxes may have multiple outputs, which allows control flow to be introduced into a Tioga-2 program and substantially enhances the expressiveness of the system. We also provide facilities for combining

multiple boxes into a single logical box and for specifying boxes with “pieces missing”. These last two facilities provide graphical analogs of the standard notions of procedures and macros.

The remainder of this paper is organized as follows. We begin in Section 2 with a quick tour of the structure of Tioga-2. This section introduces terminology and notation used throughout the paper. Section 3 presents the user’s view of Tioga-2, the user interface. The description of Tioga-2 programming begins in Section 4 with the primitive operations for editing boxes-and-arrows diagrams and performing standard database operations. Section 5 presents primitives for defining visualizations of database relations. Section 6 describes three sets of primitives for defining alternative views of data and connections between related data: (a) *drill down*, in which a user moves from a coarse visualization (e.g., a state map) to a more refined visualization of the same data (e.g., a county map), (b) *wormholes*, in which a user can move from a visualization of one data set to a visualization of a different data set, and (c) *rear view mirrors*, which allow users to keep track of “where they came from” (i.e., wormholes through which they have travelled). Section 7 continues with mechanisms to link multiple visualizations together. Lastly, Section 8 discusses database updates, Section 9 covers related work, and Section 10 presents a few conclusions.

## 2 The Model

Before presenting the Tioga-2 system in detail, we define some basic terminology and concepts. The reader may wish to skim this section on a first reading and use it primarily for reference.

Tioga-2 programs are represented by dataflow graphs with boxes and arrows. A *box* is a primitive procedure with some number of *inputs* and *outputs*. The output of one box may be connected to the input of another box by an *edge* (also called an *arrow*). Box inputs and outputs are typed and edges connect outputs to inputs of compatible types. Any attempt to connect an output to an input of incompatible type is a type error. The semantics of Tioga-2 programs is similar to the semantics of programs in dataflow languages. When data is present on all of a box’s inputs, the box can “fire”, producing results on one or more outputs. Execution is lazy, evaluating only what is required to produce the demanded visualization.

A box input or output may be a scalar value (e.g., a runtime parameter supplied by the user) or a *displayable*, described below. Displayables define visualizations. Most Tioga-2 boxes compute displayable outputs from one or more displayable inputs. Tioga-2 has three displayable types: *extended relations*, *composites*, and *groups*. In the rest of this section we discuss each displayable type in turn.

The first displayable type is an extended database relation  $R$ . In Tioga-2, the visualization of  $R$  is defined by  $R$ ’s attributes. Intuitively,  $R$  “knows” how to display itself. We assume an object-relational DBMS in which a relation has stored attributes as well as methods defining additional attributes. A displayable relation has methods to compute *location attributes* (floating point numbers) and *display attributes* (graphics). For each tuple  $t$  of  $R$ , the location attributes define the position of  $t$  on the screen and a display attribute defines the screen representation of  $t$ . Thus, Tioga-2 visualizations are constructed “tuple-wise”—the visualization of a relation  $R$  is the sum of the visualizations of each tuple of  $R$ .

Every visualization has at least two dimensions, the horizontal and vertical dimensions of the screen, and a representation for every tuple. Therefore, a displayable relation has at least  $x$  and  $y$  location attributes, corresponding to the two screen dimensions, and an attribute *display*. A displayable relation may have more than two location attributes. The number of location attributes corresponds to the dimensionality of the visualization space; the *dimension* of  $R$  is the number of  $R$ ’s location attributes.

A relation may also have multiple display attributes defining multiple, alternative representations of the data.

We adopt a uniform notation and write  $t.l$  to denote attribute  $l$  of tuple  $t$ , whether  $l$  is a stored or computed attribute. We stress that the location and display attributes used to define visualizations are computed attributes and are not stored in the database.

Displayable types are translated into screen output by *viewer* boxes. If an  $n$ -dimensional relation  $R$  is the input to a viewer, then the viewer has an  $n + 1$ -dimensional *position* specifying the location of the viewer for each of the  $n$  dimensions and the *elevation*. The user controls the position by panning in the  $n$  viewing dimensions and by zooming, which changes the elevation, moving the user “closer to” or “further from” the data. A viewer displays the  $x$  and  $y$  dimensions of  $R$  on the 2-D canvas; the remaining  $n - 2$  dimensions are available as sliders. If  $R$  has location attributes  $x, y, l_1, \dots, l_{n-2}$  each tuple  $t$  of  $R$  is rendered by drawing  $t.display$  at position  $\langle t.x, t.y, t.l_1, \dots, t.l_{n-2} \rangle$  in  $n$ -space. Because a visualization space may be much larger than the canvas, the viewer filters tuples to the ranges specified by the sliders for dimensions  $l_1, \dots, l_{n-2}$ , filters tuples to the visible real estate on the screen for dimensions  $x$  and  $y$ , and then renders the tuples’ *display* attribute to the screen.

The second displayable type is a composite of relations  $C = Composite(R_1, \dots, R_n)$ . A composite visualization is the overlay of each of the composite’s components—the visualizations are simply superimposed. Thus, composites provide a way to combine visualizations of different relations in the same viewing space. All of the constituent relations  $R_i$  have the same dimension  $n$ , which is also the dimension of the composite. If the input to a viewer is a composite  $Composite(R_1, \dots, R_n)$ , the viewer renders each of the relations  $R_i$  in order on the canvas; thus, the order of the relations specifies the drawing order. Since all of the  $R_i$  have the same dimension, a viewer for a composite has a single set of panning and zooming controls.

The third type of displayable is a group of composites  $G = Group(C_1, \dots, C_n)$ . A group visualization is just the visualization of each of the composites  $C_i$  arranged either side-by-side, top-to-bottom, or in a tabular fashion according to the user’s specification. Groups allow visualizations of different viewing spaces to be combined. To render  $Group(C_1, \dots, C_n)$ , a viewer displays each of the  $C_i$  in the specified layout. The viewer has a position for each of the  $n$  displayables—the user may independently pan and zoom in each of the grouped visualizations.

In summary, there are three kinds of displayable types, described as follows:

$$\begin{aligned} G &= Group(C_1, \dots, C_n) \\ C &= Composite(R_1, \dots, R_n) \\ R &= \text{relations with attributes } x, y, display \end{aligned}$$

Many Tioga-2 operations presented in subsequent sections are defined only for  $R$  or  $C$  inputs. To make programming easier, Tioga-2 extends such operations to work on “higher” types. For example, the **Restrict** operation filters a relation; it is a box that takes an  $R$  input and produces an  $R$  output. Given a group  $G$  input to **Restrict**, Tioga-2 asks the user for the composite within the group, and the relation within that composite, to which the **Restrict** applies. After applying the **Restrict** to the selected relation, Tioga-2 reassembles the composite and the group in the obvious way. This is all done graphically with point-and-click operations, so that the user need not be aware explicitly of how **Restrict** is overloaded to work on group and composite displayables. Similarly, an operation defined on composite types (such as **Overlay**) is extended to work on group displayables by having the user first specify which component of the group is to be the operation’s input.

Finally, Tioga-2 provides the obvious type equivalences  $R = Composite(R)$  and  $C = Group(C)$ , so

that a relation can be used wherever a composite is expected, and a composite can be used wherever a group is expected.

### 3 User Interface

The Tioga-2 user interface contains several main windows. All may be visible on the screen or iconified. There is a single user interface both for building and for using programs, but a user browsing a previously constructed visualization will not require all of the windows available. A screen dump of the interface is shown in Figure 1.<sup>1</sup> The user interface windows are:

- a *program* window, containing a boxes-and-arrows representation of a Tioga-2 program,
- a *canvas* window for each viewer in the current program,
- a *menu bar* containing the pull-down menus to invoke primitive operations.

A canvas window shows data visible in a viewer at the current position. In addition, each canvas window includes

- a *rear view mirror*,
- zero or more *slider* bars,
- an *elevation map*, and
- an *elevation control* (a dashed line through the elevation map).

The menu bar includes

- a menu of all operations available,
- a menu of all tables available,
- a menu of all boxes available,
- an *undo* button to undo the last operation performed, and
- a *help* button.

A Tioga-2 program is constructed incrementally by applying program editing operations to the program window (thereby modifying the boxes-and-arrows diagram) and rendering and/or drill down operations to a canvas window (thereby making modifications via direct manipulation). At any stage in the construction of a program the current result is displayed on all non-iconified canvases.

Since a canvas may be much larger than the available screen real estate, we allow the user to change the viewer's position, altering the area visible in the viewer. Scroll bars control panning in the screen dimensions  $x$  and  $y$ ; canvas slider bars control panning in any remaining dimensions. The elevation control allows the user to drill down into data displayed on the screen. Elevation maps are an interface for programming drill down (Section 6).

---

<sup>1</sup>The program and visualization in this figure are discussed in Section 4.

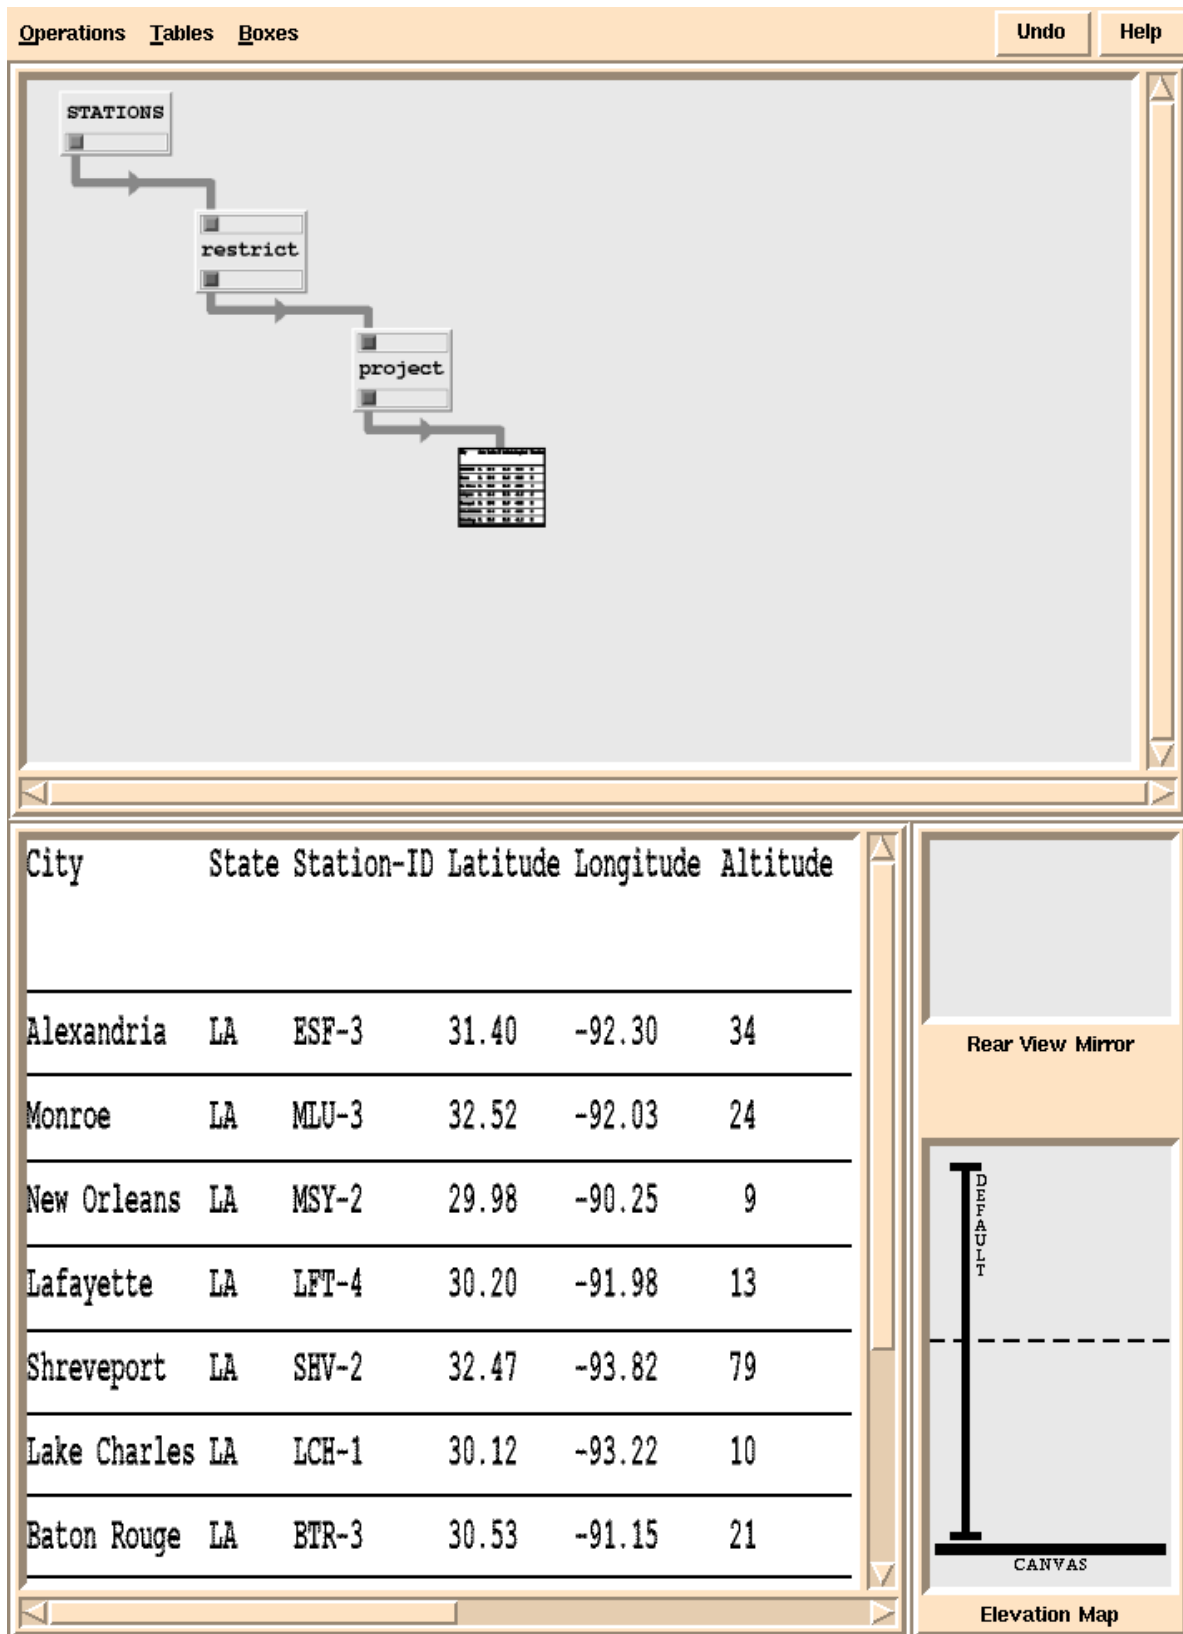


Figure 1: Weather stations in Louisiana.

## Program and Data Management Operation

This section discusses the operations available in the program window and Tioga-2's database operations. These operations allow the incremental construction of a boxes-and-arrows program specifying data the user wishes to visualize. Operations for constructing visualizations themselves are discussed beginning in Section 5.

We use the following example to illustrate Tioga-2 programming. An agricultural specialist wishes to construct a visualization of temperature and precipitation data for various sites in Louisiana. The data is stored in two relations: *Stations*, which contains a tuple describing each weather station, and *Observations*, which contains all observations (e.g., data, time, conditions) from all stations. The data covers all of North America and contains a great deal of information besides temperature and precipitation.

As a first step toward constructing a temperature and precipitation visualization for Louisiana, the user limits the *Stations* relation to the stations of interest. For every relation known to the Tioga-2 system there is a box of the same name that takes no inputs and produces as output the tuples of the relation. Beginning with the *Stations* box, the user incrementally adds boxes to perform standard database operations such as restricting the data to tuples satisfying a predicate (e.g., stations in Louisiana) and projecting out unneeded fields (e.g., date of construction). A boxes-and-arrows diagram and canvas are shown in Figure 1. The last box in Figure 1 is a viewer, which in this case displays data using a default two-dimensional table format. The user can also inspect any of the partial results. If the user discovers that any step produces unexpected results, he can inspect, delete, and replace boxes as necessary to fix the program.

For convenience, the operations in this section are subdivided into operations that manipulate program structure and database operations.

### 4.1 Program Operations

This group of primitives permits the initialization, loading, and saving of programs, as well as the deletion, insertion, and connection of boxes into an existing program. There are also primitives that provide familiar language abstractions analogous to procedures and macros. The operations are listed in Figure 2; we briefly discuss the most interesting.

If the user clicks on one or more edges in the current program, **Apply Box** gives the user a menu of all boxes whose inputs match the types of the selected edges. This is a shorthand way to identify those boxes in the database that could possibly take the indicated edges as input.

A design principle of Tioga-2 is that every operation preserves a visual representation. Deleting boxes from a program is dangerous, because inputs of other boxes may be left dangling and, therefore, their results unavailable for visualization. To preserve the property that "everything is always visualizable", arbitrary box deletions are not allowed in Tioga-2. A box may be deleted if (1) it has no outputs connected to other boxes (in which case no box inputs are left dangling), or (2) it has a single input and output of the same type (in which case the system connects the deleted box's predecessor to its successor). A box may also be **Replaced** by another box with compatible types.

A **T** box simply passes its input unchanged to both outputs, and allows another box, for example a viewer, to be connected to the **T**.

**Encapsulate** permits the user to define new boxes. The user specifies a portion of the program to be encapsulated by drawing a closed curve around a region of the program. Edges cut by the curve are the inputs and outputs of the new box. The new box may be used like any other primitive box.

Encapsulated boxes may also be parameterized to create something akin to a macro or (more accurately) a higher-order function. The user draws additional closed areas within the program region to be



<i>Operation</i>	<i>Effect</i>
<b>New Program</b>	Erase the program canvas.
<b>Add Program</b>	Add a named program to the program canvas.
<b>Load Program</b>	Shorthand for <b>New Program</b> followed by <b>Add Program</b> .
<b>Save Program</b>	Save the current program in the database.
<b>Apply Box</b>	<i>see discussion</i>
<b>Delete Box</b>	<i>see discussion</i>
<b>Replace Box</b>	Replace one box by a different box with compatible types.
<b>T</b>	Add a T-node to a designated edge.
<b>Encapsulate</b>	<i>see discussion</i>

Figure 2: Operations that manipulate the boxes-and-arrows diagram.

<i>Operation</i>	<i>Box Type</i>	<i>Effect</i>
<b>Add Table</b>	$\emptyset \rightarrow R$	Add the box producing a specified relation as output.
<b>Project</b>	$R \rightarrow R'$	Standard database projection; user is prompted for fields.
<b>Restrict</b>	$R \rightarrow R$	Filter a relation to tuples satisfying a predicate.
<b>Sample</b>	$R \rightarrow R$	Randomly sample a relation.
<b>Join</b>	$R \times R' \rightarrow R''$	Standard join of two relations; user is prompted for join predicate.

Figure 3: Operations on relations.

encapsulated. These areas become “holes”—they are not included in the encapsulated box, and edges cut by a hole are unconnected. To use an encapsulated box with holes, the user must specify a box—with compatible types—that can be plugged into each hole.

## 4.2 Database Operations

The primitives in this group provide traditional (and a couple of non-traditional) database operations, which are listed in Figure 3. Each operation adds a new box to the program. The type of the introduced box is indicated in Figure 3. Note that all input/output types are  $R$ . As discussed in Section 2, these operations are extended to apply to composite ( $C$ ) and group ( $G$ ) types as well.

As mentioned above, the **Add Table** operation adds a new “source” box to the current program. The box is named for a table in the database and has a single output edge. The parameters of many Tioga-2 operations can be specified in several ways; usually there is at least one textual and one graphical method. For example, the user may specify the table to add to the program by either typing the name or selecting it from a menu of available relations. Note that **Add Table** is a special case of **Apply Box** with zero inputs.

A **Restrict** box filters its input, retaining only tuples that satisfy a restriction predicate. The user is prompted for the predicate to be applied. A **Sample** box produces a random subset of an input relation on its output. Each input is retained with a user-specified probability. **Sample** is useful for improving interactive response by reducing the size of data sets to be processed.

The result of applying these operations is to iteratively build up a boxes-and-arrows program in the program window. We now turn to the visualization of the result of such programs.

## 5 Rendering Operation

The previous section has indicated how a Tioga-2 program can be built to retrieve complex computations (relations) from the database. Now we must deal with two additional questions:

- How are tuples positioned on the canvas?
- How are tuples rendered as screen pixels?

As discussed in Section 2, these questions are addressed by location attributes specifying the position of tuples in  $n$ -space and display attributes that specify tuples' screen representations. This section describes location and display attributes, default displays, and their associated operations.

### 5.1 Location and Display Attributes

Figure 4 shows a visualization of the Louisiana weather station data produced by the diagram shown in Figure 1. Each station in the state is represented by one tuple in the relation. The visualization shows a circle and the name of each station at its (*longitude*, *latitude*) coordinate. To position representations of tuples on the screen, relations have location attributes. Every relation must have  $x$  and  $y$  location attributes to specify the  $x$  and  $y$  dimensions of a 2-D canvas; in Figure 4, the  $x$  dimension is longitude, the  $y$  dimension is latitude. There may be additional location attributes, which specify slider dimensions. In Figure 4, there is a slider dimension *Altitude*. By setting the range of altitude values that are visible using the slider, the user can see any appropriate subset of the stations. Location attributes are represented by floating point numbers.

Tioga-2 requires that every relation have at least one *display* attribute. A display attribute is a list of primitive *drawable* objects. Intuitively, a viewer renders a tuple by simply painting each drawable in its *display* attribute on the screen. The list order specifies the drawing order. In Figure 4, the display attribute is a list containing the text of the name of the station and a circle. There may be additional display attributes to provide alternative visualizations of the data.

The primitive drawables include: point, line, rectangle, circle, polygon, text, and viewer. Each primitive drawable has an *offset*, a *color*, and a *style*. The offset gives a position relative to the location attributes of the tuple; thus, multiple drawables need not be stacked directly one atop the other. In Figure 4, the name is positioned below the circle. Of the primitives listed above, all but viewers are standard primitives for graphics hardware. Viewers are used to implement wormholes (Section 6). The list of primitive drawables is preliminary and more may be added in the future.

In Tioga-2, every relation is augmented with location and display attributes. Actually computing the values of these attributes should be avoided except where necessary. As discussed in Section 2, display and location attributes, along with any other “extra” attributes, are specified by functions of the base tuple.

### 5.2 Defaults

To guarantee that boxes produce relations with initial valid displays, Tioga-2 provides default location and display attributes. There is a default display for each atomic type (i.e., each type of a column of a relation). The default display for a relation renders each field in the tuple, side by side, using the default display for each column type to produce a screen representation. The default space has two dimensions: the  $x$ -location is 0 and the  $y$ -location is the sequence number of the tuple. Typically, the default attributes

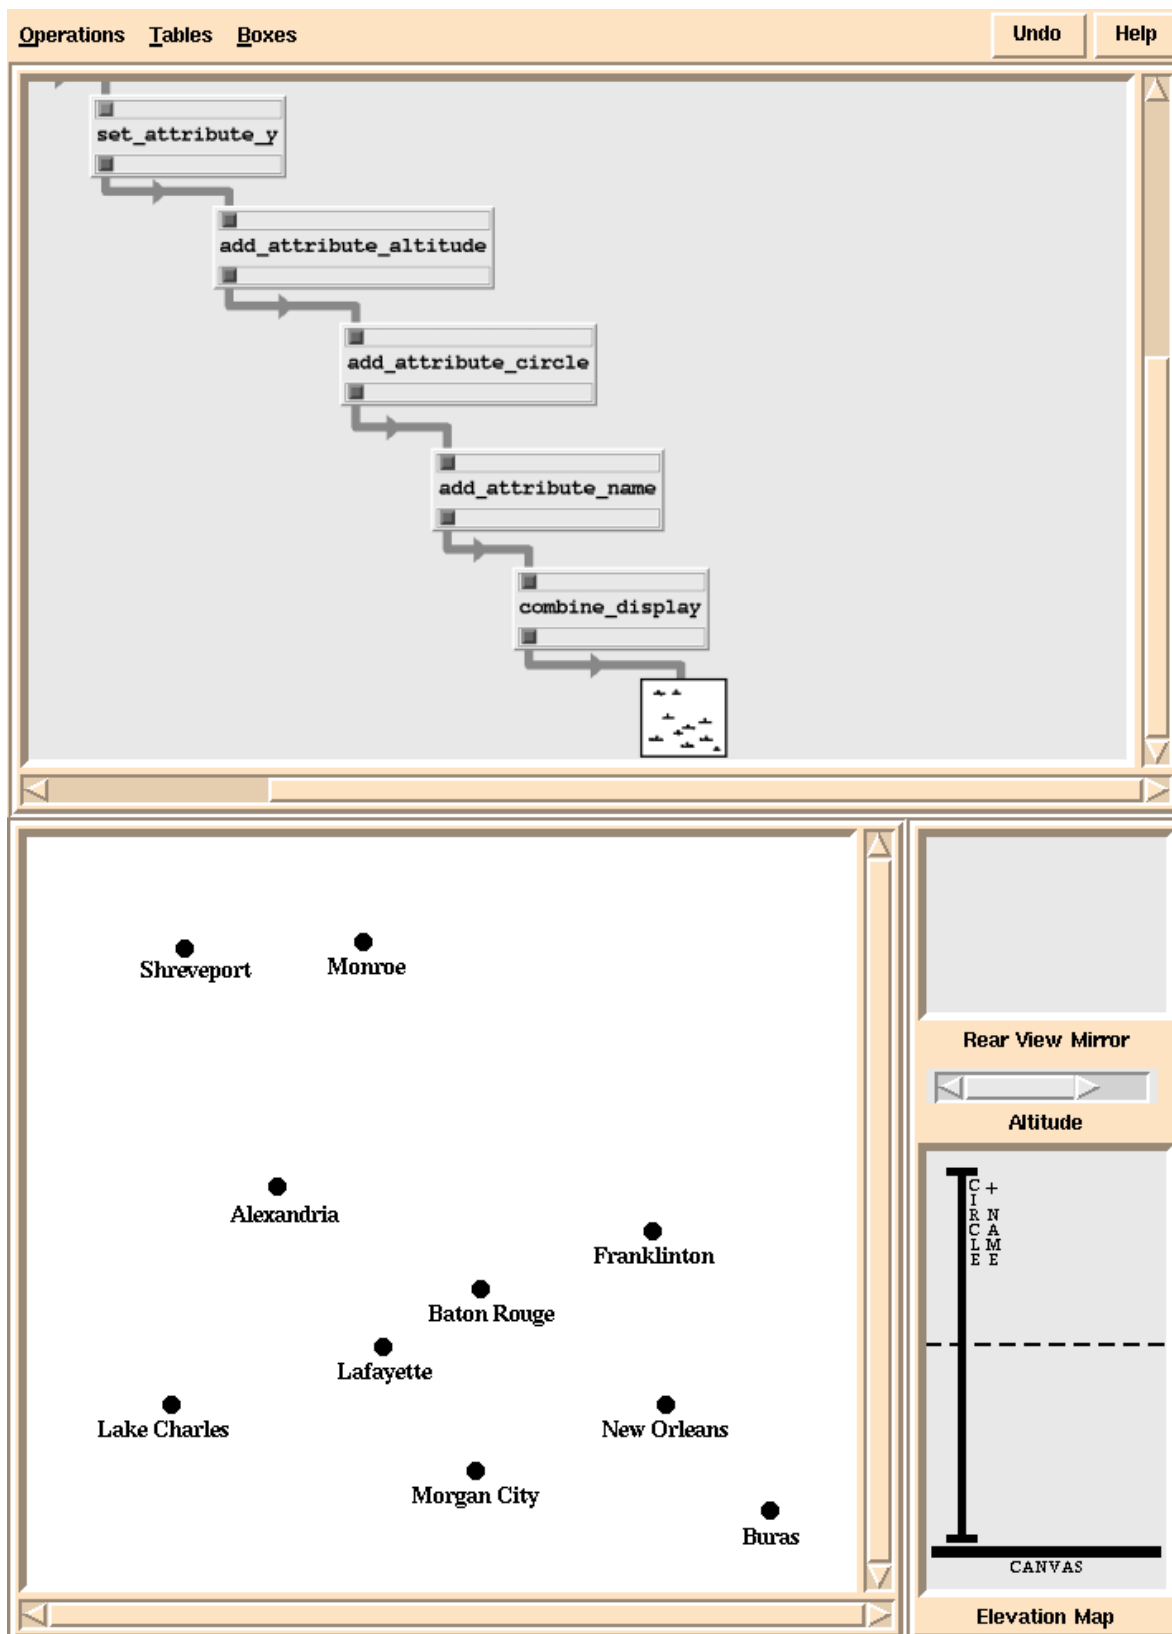


Figure 4: A visualization of weather station locations.

<i>Operation</i>	<i>Box Type</i>	<i>Effect</i>
<b>Add Attribute</b>	$R \rightarrow R'$	Add an attribute to a relation; user is prompted for definition.
<b>Remove Attribute</b>	$R \rightarrow R'$	Remove an attribute; cannot remove attributes $x, y$ , or <i>display</i> .
<b>Set Attribute</b>	$R \rightarrow R'$	Change the value of an existing attribute.
<b>Swap Attributes</b>	$R \rightarrow R'$	Interchange two attributes of the same type.
<b>Scale Attribute</b>	$R \rightarrow R'$	Multiply numerical attribute by a number.
<b>Translate Attribute</b>	$R \rightarrow R'$	Add a number to a numerical attribute.
<b>Combine Displays</b>	$R \rightarrow R'$	Combine two display attributes.

Figure 5: Location and display operations.

define a display consisting of a sequence of tuples in ASCII. The major relational DBMS vendors all have so-called *terminal monitors*, which produce a display of this form for the result of any possible query.

Just as the user may incrementally modify the data management operations to change the data to be visualized, so may the user incrementally modify the location and display attributes of a relation to change the visualization. Initially, every **Add Table** operation introduces a box that produces a relation with the default display and location. The user may then incrementally modify the defaults, or replace them altogether, by adding boxes to the diagram or by manipulating data on the canvas. In Figure 4, the default viewer of Figure 1 has been changed by modifying location functions (to associate  $(longitude, latitude)$  with  $(x, y)$  canvas coordinates) and the display function (changed to the combination of station name and a circle).

### 5.3 Operations

In the remainder of this section we discuss the operations for modifying location and display attributes listed in Figure 5. Most of these operations apply to all attributes, not just location or display attributes.

The user may add new attributes, including new location and display attributes. Adding a location attribute adds a new dimension to the visualization. Adding a display attribute creates an alternative visualization of the data. **Add Attribute** prompts for the type and definition of the new attribute; the definition may depend only on other attributes of the relation. **Set Attribute** changes the type and definition of an existing attribute.

In both **Add** and **Set Attribute**, an attribute's definition may be given in a general query language. However, the preferred method is to begin with a very simple definition (e.g., a copy of another field, or a single primitive drawable) and then refine it using the other operations.

**Swap Attributes** is handy for interchanging two dimensions (two location attributes), thereby “rotating” the canvas, or interchanging the display attribute with one of the alternative displays, thereby changing the visualization of the data.

**Scale** and **Translate Attribute** are defined only for numeric fields. These operations are convenient shorthands for more complex **Set Attribute** commands. **Scale** and **Translate** are useful for changing location attributes, thereby scaling or translating dimensions of a visualization.

**Combine Display** is the mechanism for combining primitive drawables to form more complex displays. The user positions the displays on top of one another graphically to establish the relative position; alternatively, an explicit offset of one display to the other can be entered. The combined display becomes a new display attribute. The user may combine any of the display attributes of the relation. In Figure 4, a circle display has been combined with a text display showing the name of the station.

<i>Operation</i>	<i>Box Type</i>
<b>Set Range</b>	$R \rightarrow R$
<b>Overlay</b>	$Composite(R_1, \dots, R_n) \times Composite(R_{n+1}, \dots, R_m) \rightarrow Composite(R_1, \dots, R_m)$
<b>Shuffle</b>	$Composite(R_1, \dots, R_{i-1}, R_i, R_{i+1}, \dots) \rightarrow Composite(R_i, R_1, \dots, R_{i-1}, R_{i+1}, \dots)$

Figure 6: Primitives for drill down.

## 6 Drill Down

Drill down allows users to see more details in data of interest. There are two distinct, useful notions of drill down. The first provides a more refined view of the same data in the same visualization space (e.g., switching from a state to a county map). The second allows movement between one space and a different, but semantically related, space (e.g., after finding a weather station, switch to look at its temperature/precipitation data).

Two mechanisms provide drill down in Tioga-2. First, the user can specify that additional detail about screen objects becomes available as the user zooms in. Second, we have a notion of *wormholes*, by which a user can move from one canvas to another canvas.

### 6.1 Additional Detail

The first form of drill down is defined as operations on relations  $R$  and composites  $C$ . There are three operations:

- **Set Range**

This operation specifies the maximum and minimum elevations at which a relation’s display is defined. Outside of this range, the relation contributes nothing to the canvas.

- **Overlay**

Two composites may be overlaid. Since a relation  $R$  is also a trivial composite  $Composite(R)$  (Section 2), relations may also be overlaid. The relative position of one overlay to another may be given either by an explicit  $n$ -dimensional offset, or by dragging one canvas over the other. If the component displays are defined with different elevation ranges, then it is possible to program drill down by having the displayable at the lower elevation provide a specialization of the displayable at the higher elevation.

- **Shuffle**

It may be desirable to change the drawing order of the relations within a composite. The **Shuffle** command moves a relation to the “top” of the drawing order.

Figure 7 illustrates overlay and setting ranges. Weather stations are now shown together with a map of Louisiana; this is achieved by overlaying the map (derived from a relation of lines defining the map) with the result of Figure 4. In addition, a third display is overlaid to give less detail at higher elevations. This display shows only a circle at the station’s location. The ranges of the two weather station displays are set so that station names disappear at high elevations, where they would be illegible.

There is a small difficulty with the overlay in Figure 7. The visualization of the state map of Louisiana has no *Altitude* dimension, and such a dimension makes no sense for a flat map. However, the composite has an *Altitude* slider; how are changes in Altitude to be interpreted for the Louisiana map? If the user attempts to overlay relations with different dimensions, Tioga-2 warns about the mismatch. If the user

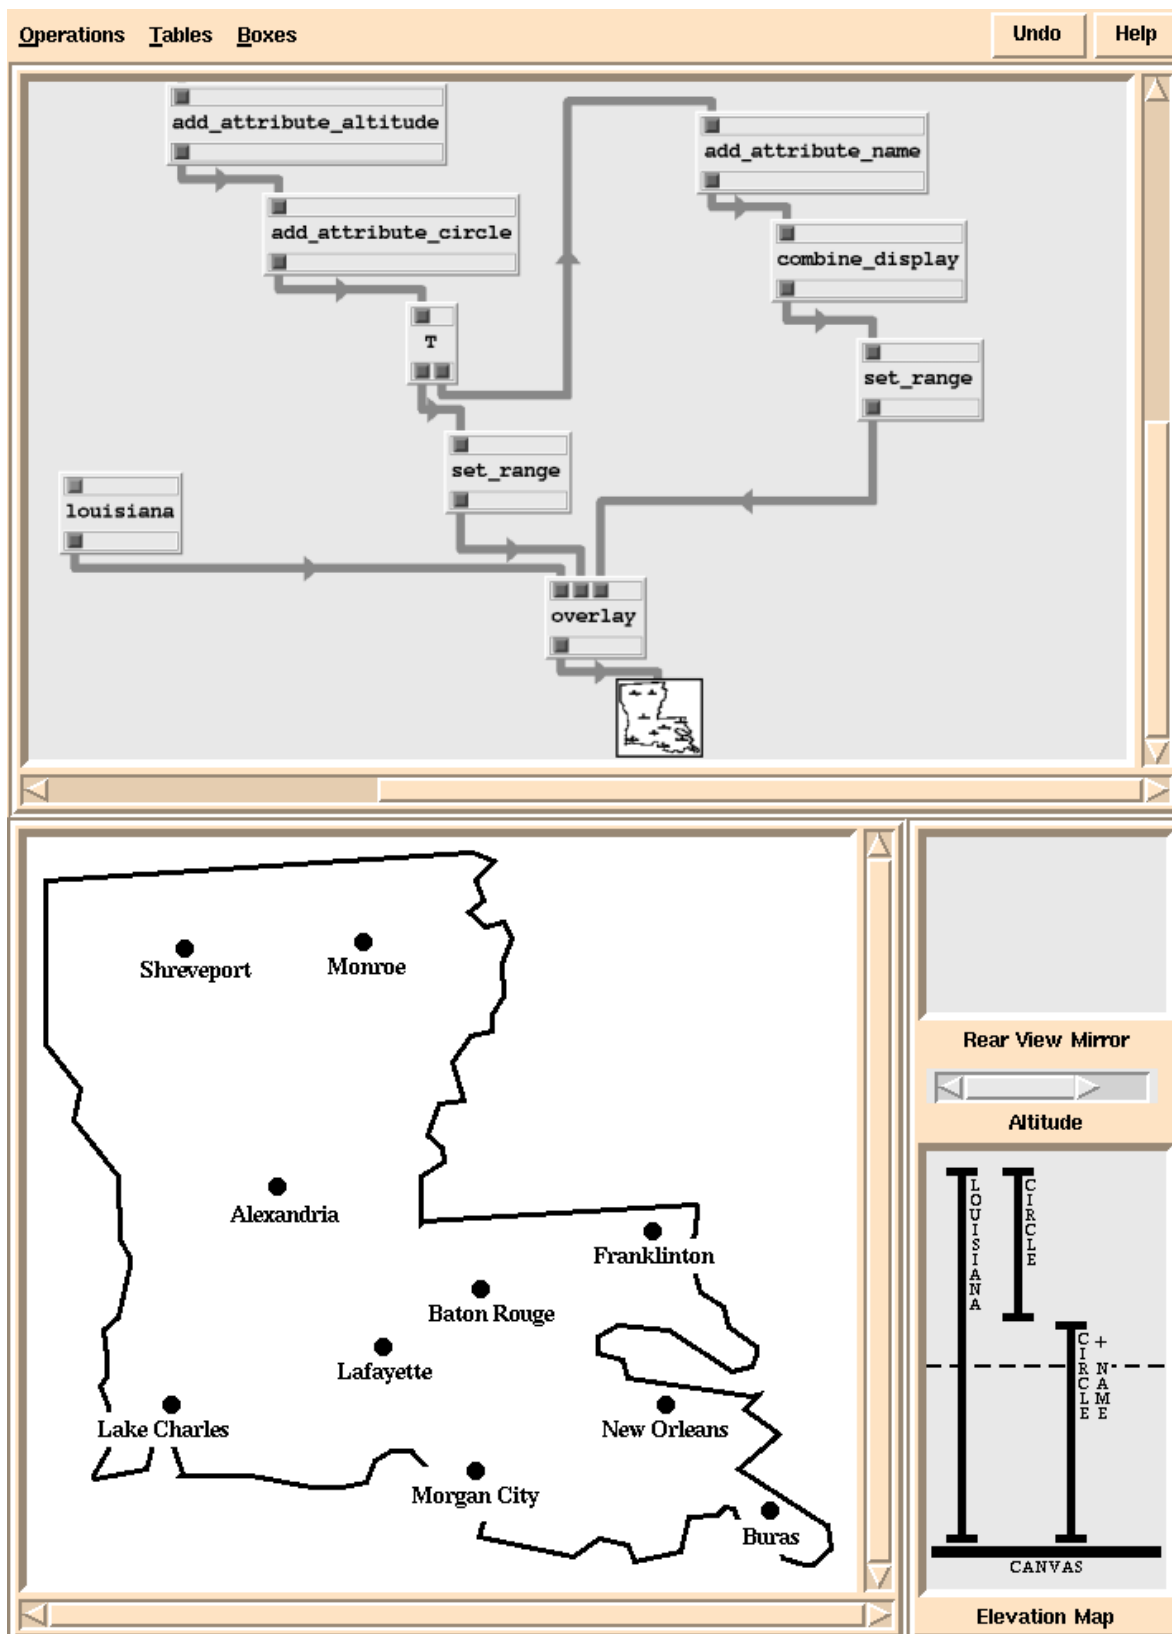


Figure 7: Overlaid displays with restricted ranges.

wishes, the underlying relations are treated as invariant in the “extra” dimensions. This achieves the desired effect in Figure 7: the user can change the *Altitude* slider to see different subsets of the stations, but the Louisiana map remains in place for reference.

The *elevation map* is a bar-chart display of the maximum/minimum elevations and drawing order of all elements of a composite on the current canvas (see Figure 7). The elevation map can be manipulated directly by the user to adjust the ranges and drawing order of overlaid relations. For a group displayable, a viewer shows an elevation map for only one member of the group at a time. In this case, the user can explicitly cycle through all of the elevation maps.

## 6.2 Wormholes

It is often desirable to associate objects in one visualization space directly with objects in a different visualization space. A *wormhole* is a viewer onto another canvas, i.e., what is visible inside a wormhole is a point on another canvas from some elevation. Figure 8 shows an example application of wormholes. Upon zooming into an individual station  $s$ , a wormhole appears (achieved by a combination of modifying display functions and overlaying and setting ranges) that takes the user to a canvas displaying temperature data for each station as a function of time. The user is initially positioned viewing the data for station  $s$ .

Providing wormholes is technically straightforward. Viewers are primitive drawable objects; thus, Tioga-2 programs may produce displays containing viewers (wormholes). A viewer drawable requires several parameters, including the size for the viewer, a destination canvas, the elevation from which the canvas is viewed, and the initial location; the user defines these values as part of the display attribute. As with any drawables, wormholes can be overlaid with other drawables. In Figure 8, the axes labels are the result of overlaying text at an offset from the wormhole (for brevity, these boxes are not shown).

When a user zooms in on a wormhole and reaches zero elevation he passes through the wormhole and moves from his original canvas to the destination canvas. Needless to say, the user can pan and zoom on this second canvas, as well as move to a third canvas. After changing canvases several times, there is a definite possibility the user will get lost. For this reason, we introduce the notion of a rear view mirror.

## 6.3 Rear View Mirrors

For each canvas, we introduce an additional window called a *rear view mirror*. This window shows the “bottom side” of the canvas through which the user last moved. Hence, immediately after going through a wormhole, the user is looking down at a new canvas from some specific elevation and is at negative ground level for the canvas he just left. As he descends toward the new canvas, he increases the distance from the previous canvas. In Figure 8, the rear view mirror shows that the user came through a wormhole at New Orleans in Figure 7.

Every Tioga-2 displayable has a minimum and maximum elevation. If both are positive, then the viewer only shows objects on the top side of the canvas. If the minimum and maximum elevations are both negative, then the viewer places objects only on the underside of the canvas, and they are visible only in the rear view mirror after the user proceeds through a wormhole. If the minimum elevation is negative and the maximum is positive, then the objects can be seen on both sides of the canvas. Thus, the programmer can create overlays in such a way that the top side and the underside of the canvas both have meaning. One is visible from above in the viewer window and one is visible from below in the rear view mirror.

A natural use of the rear view mirror is to illuminate the wormholes back to the canvas from which the user came to this canvas. In this way, the user can “find his way home” if he gets lost. As such, the rear view mirror is a generalization of the notion of “back” in a hypertext system.




Figure 8: A visualization with wormholes.



## 7 Addition l Oper tion

This section discusses the remaining Tioga-2 features. *Slaving* constrains two viewers to move together. *Magnifying glasses* provide hierarchical viewers (viewers within viewers). As discussed below, magnifying glasses are quite different from wormholes. *Stitch* and *replicate* produce group displays. Slaving and magnifying glasses are operations on viewers, while stitch and replicate are operations on displayables.

### 7.1 Slaving

Two viewers may be *slaved* together, in which case the system maintains the relative offset between the two viewers. When a viewer is deleted, all of its slaving relationships are also deleted. Slaving relationships may be removed explicitly as well. Slaving is only defined for two viewers with the same dimensions.

### 7.2 Magnifying Glasses

A user may create a *magnifying glass* by placing a viewer inside of another viewer. Typically, a user will place a copy of the current viewer inside of itself; he will then zoom the inner viewer, so it magnifies what is in the outer viewer. Magnifying glasses must have the same dimension as their containing viewer. The inner and outer viewers may be slaved so that they move in unison. Magnifying glasses may also be deleted.

An simple technique for correlating temperature and precipitation is shown by the use of a magnifying glass in Figure 9. In this case, the user begins with a temperature vs. time display. The underlying relation that is being visualized has more information—in particular, the precipitation data—that is not being utilized. An alternative display attribute shows precipitation vs. time (the boxes defining the precipitation display are not shown). By creating a magnifying glass using this alternative display, the user can see the precipitation data for points underneath the magnifying glass. In Figure 9, the magnifying glass is realized by making the precipitation display the *display* attribute (done by the **Swap Attribute** box) and then viewing the resulting relation.

### 7.3 Stitch

Any number of composites can be *stitched* together to form a group displayable. Groups can be displayed side-by-side, arranged vertically, or laid out in a tabular fashion. If the user performs a window operation on one of the group members, such as moving the window on the screen or iconifying it, then the same operation is performed on the other members. Zooming and panning is defined for each of the constituent displays. That is, there is a separate focus for all components, as well as separate  $x$ ,  $y$ , slider, and zoom dimensions. Components may be slaved to one another.

In Figure 10, a display showing temperature vs. time is stitched to a display showing precipitation vs. time. The precipitation display is slaved to the temperature display, so that whenever the user changes the date range under temperature, the precipitation display changes to display the same date range.

### 7.4 Replicate

A relation can be *replicated* by specifying a partition. Replicated displays for each partition are stitched together into a group. The user must specify the area to be given to each display and the initial point of focus.

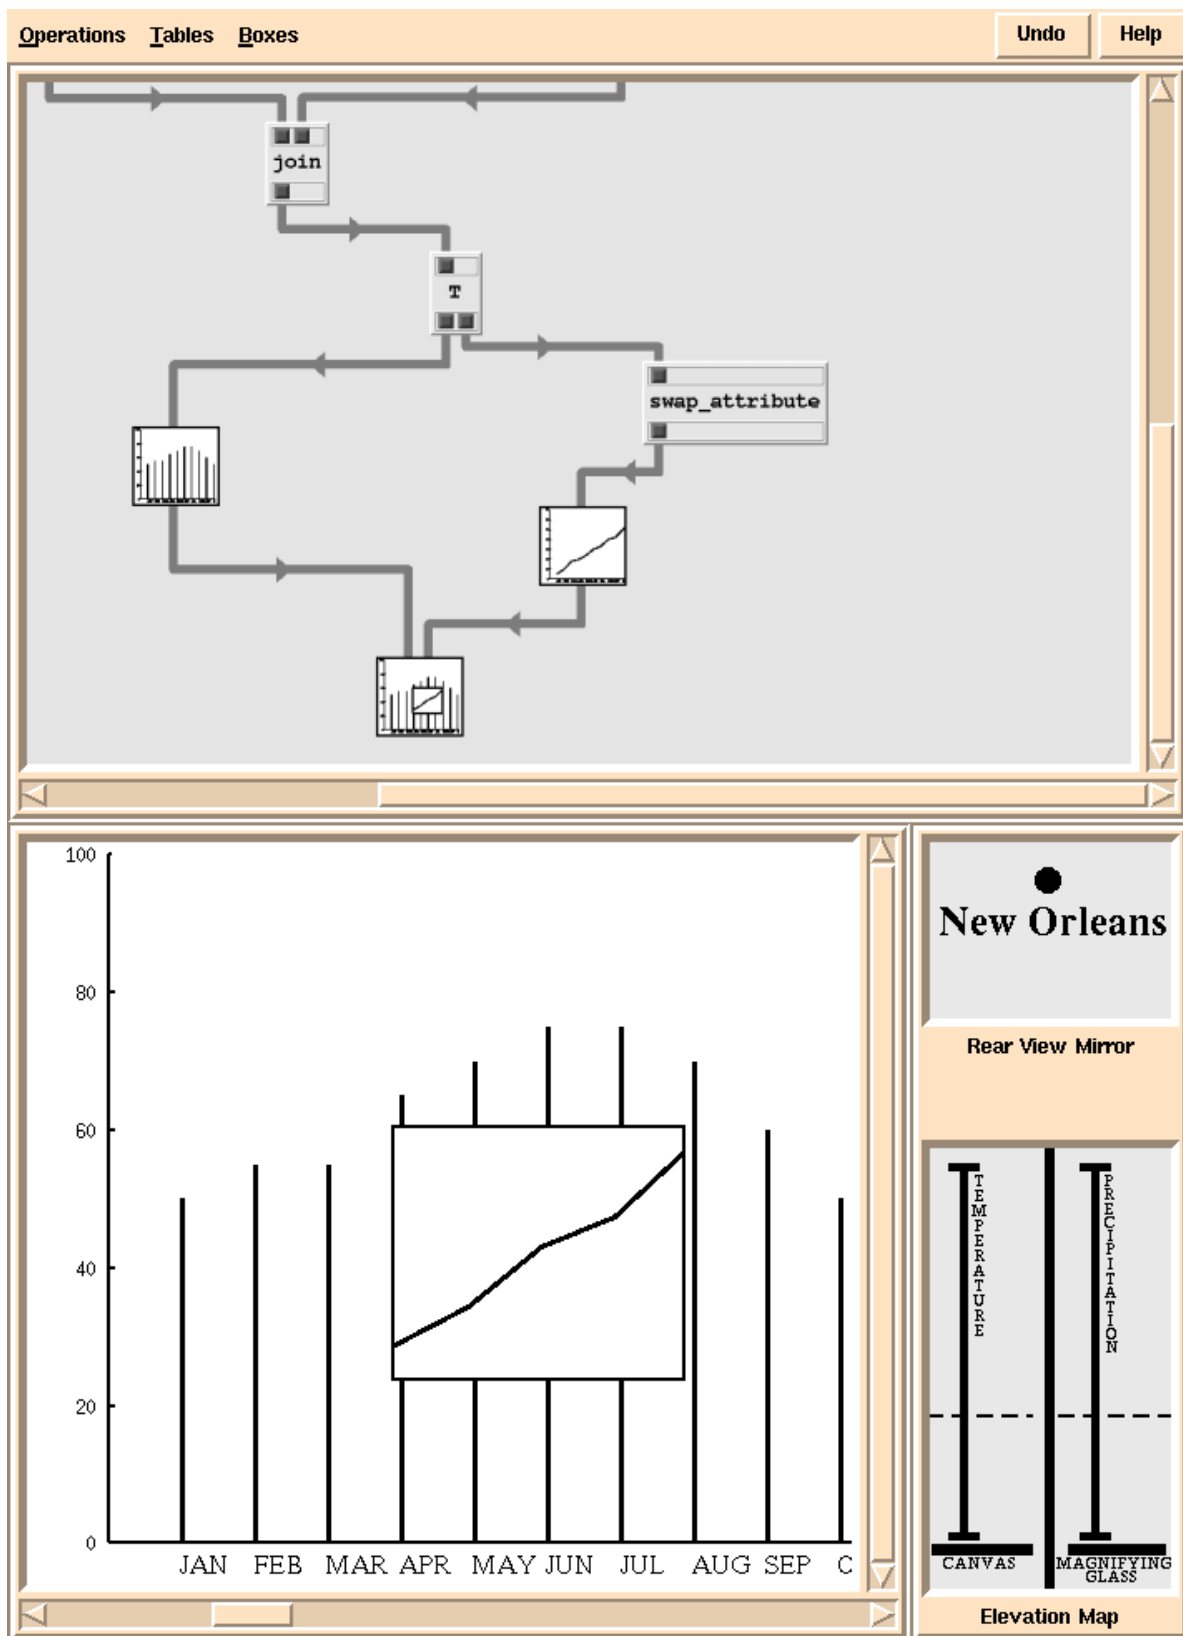


Figure 9: Using a magnifying glass.

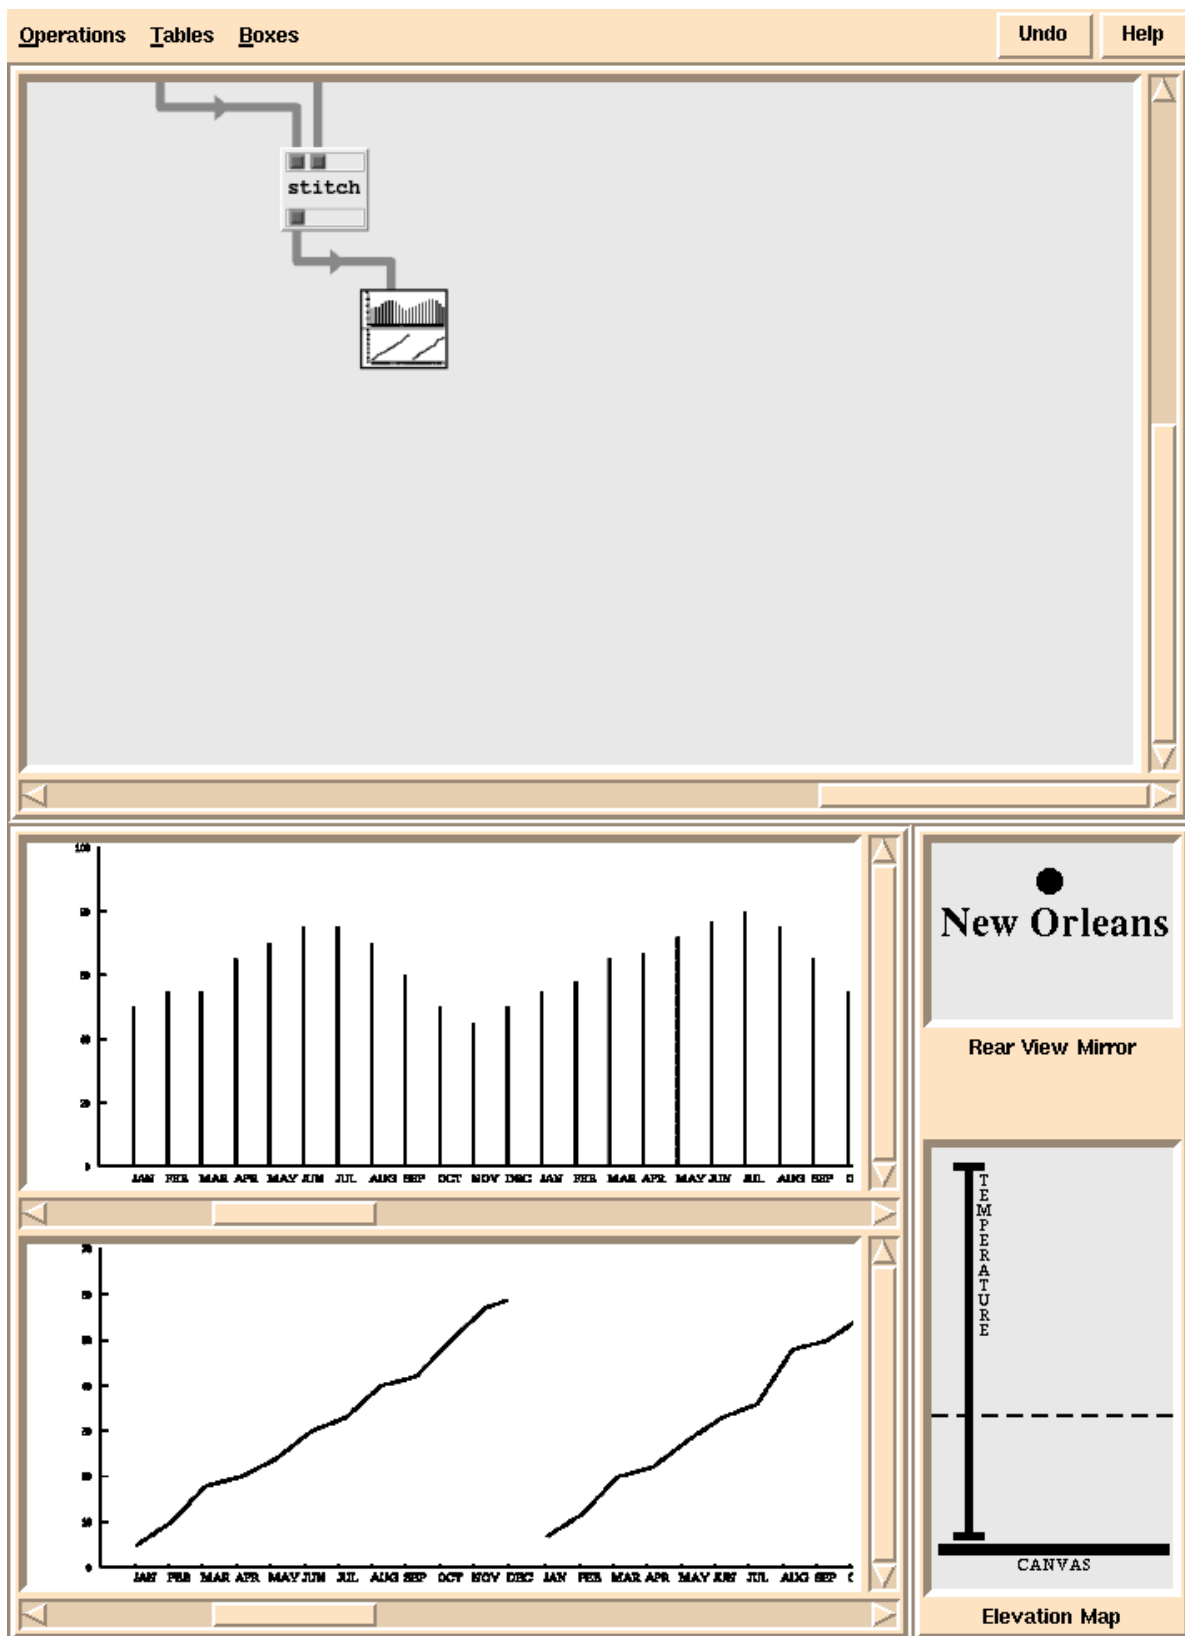


Figure 10: An example of stitched viewers.

The partitioning predicate is specified by giving a collection of predicates in the underlying query language or an enumerated type. For example, the specification may be that replication is tabular, with predicates `salary ≤ 5000` and `salary > 5000` in the horizontal dimension and the enumerated type `department` in the vertical dimension.

In Figure 11, a viewer showing temperature vs. time and precipitation vs. time has been replicated to show records for years prior to 1990 and after 1990 separately. This example motivates the need for the operator overloading discussed in Section 2. Because **Replicate** partitions a relation, it takes an  $R$  as input and produces multiple  $R$ 's as output. However, in this example the display is a  $G$  type (a group of two displays). Thus, before the replication can be performed, the user must specify the relation. When the user selects **Replicate**, the system prompts the user for the group component on which the replication is to be done.

## 8 Update in Tioga-2

Tioga-2 is oriented toward browsing a database. As such, we expect users to wander around a canvas and possibly notice things they wish to update. For example, the quantity on hand of specific items could appear on a canvas. The user would find an item of interest and then wish to order a certain number of the item, thereby decreasing the quantity on hand. The user could also notice data errors and simply wish to fix them. As a result, we focus on providing an update capability that allows specific screen objects to be updated in the database. We do not consider general SQL update statements in Tioga-2.

For each primitive type, the type definer is required to implement a default display function that is used by Tioga-2 to render tuples containing this type. Similarly, we require the type definer to write a second *update* function that enables Tioga-2 to provide updates for instances of the type that appear on the screen. When a user clicks on a screen object, the Tioga-2 run time system activates a generic update procedure, passing it the tuple corresponding to the screen object. The function engages a dialog with the user to construct a new tuple—using the primitive update functions for the fields—and then perform an SQL update to install the new value in the database. This machinery is all encapsulated within the update function itself.

When the user customizes a visualization, he can replace the default update command with one of his own choosing, if he so desires. In this way, he can make an update system with a desired “look and feel”.

## 9 Related Work

While developing browsers for exploring data is a relatively new research area, the literature is already substantial. This section surveys a cross-section of related work.

As discussed in Section 1, Tioga-2 retains the boxes-and-arrows notation for programs originally developed for dataflow languages and popularized for visualization by AVS [U<sup>+</sup>89], Data Explorer [LAC<sup>+</sup>92], and Khoros [RY92]. These systems are similar to Tioga in their reliance on simplifying programming by using dataflow graphs. Thus, these systems share Tioga's basic problem that boxes-and-arrows notation alone does not simplify programming sufficiently for novice programmers (see Section 1.1). *Weaves* is another boxes-and-arrows system [CGR91]. Weaves are intended to support visual programming, so the boxes-and-arrows program is itself the only visualization of interest. An extension of weaves supports a limited notion of drill down [GQ94].

Many browsing systems are based on a “paradigm”. A classic example is the Fisheye interface, which

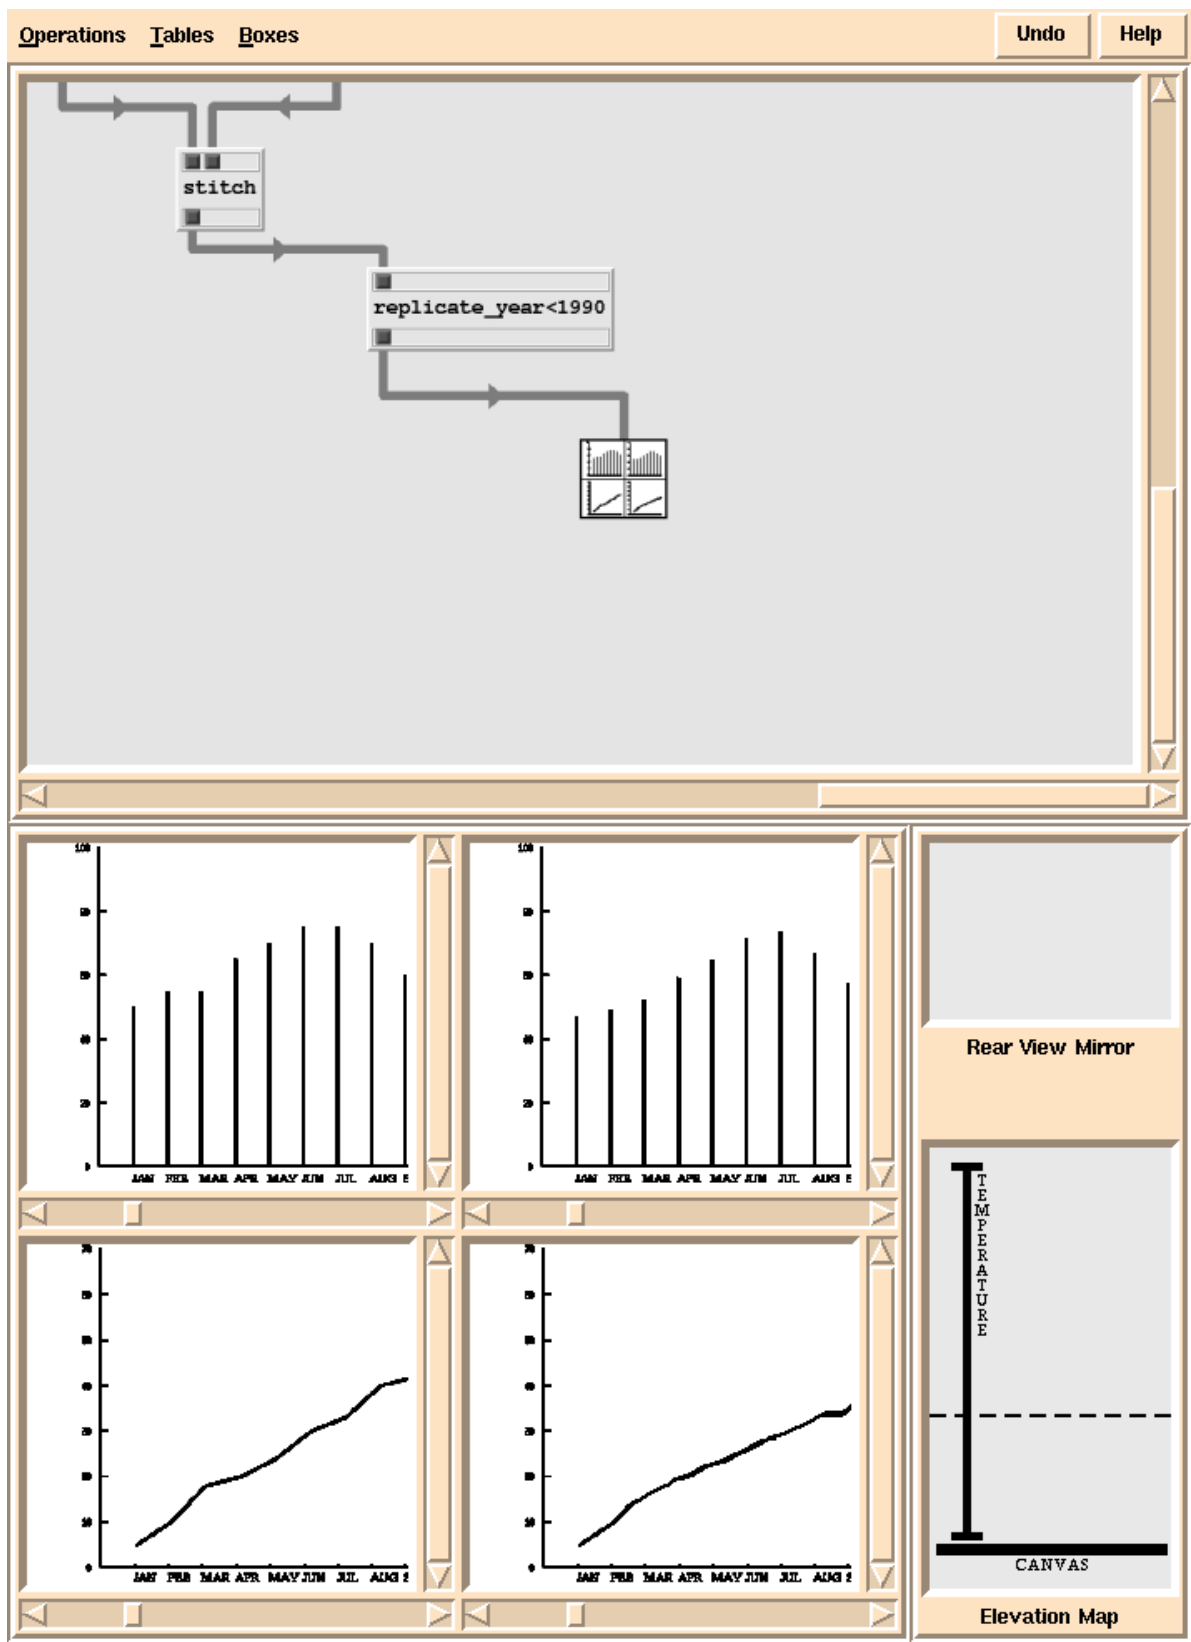


Figure 11: A replicated viewer.

magnifies data in the center of focus to a greater degree than data at the periphery [SS94]. Another example is Magic Lenses, which provides a set of primitive lenses (windows akin to our magnifying glasses) that can be placed over data and over each other to modify a visualization [BSP+93]. While we find paradigms appealing, we suspect a flaw in the assumption that the space of possible visualizations can or must be greatly restricted in advance.<sup>2</sup> In our experience, paradigms serve a class of users well and frustrate users with other applications. To be generally useful—as Tioga-2 aims to be—it is important that users be able to construct arbitrary *ad hoc* visualizations of their own, even inventing their own paradigms if necessary. In short, visualizations should be as programmable as possible.

A different approach has been taken by the ambitious Pad project [PF93]. In Pad, all data lives on a two-dimensional plane. As in our system, every entity (an object in Pad, a tuple in Tioga-2) has a position and “knows” how to draw itself. Pad also provides facilities for overlay and drill down that are in some ways richer than the facilities in Tioga-2. Pad allows a very large class of visualizations to be built. However, Pad is not end-user programmable; it is designed as a toolkit for expert programmers and provides a traditional programming interface.

Within the area of browsers for databases, the work of Krishnamurthy and Zloof on Rendering By Example (RBE) is closest to our own. In particular, RBE shares our view on the importance of a system that is both highly programmable and easy to program [KZ95]. RBE provides a more declarative programming interface than Tioga-2, but RBE can construct a much less general class of visualizations.

Finally, a database-centric visualization system raises the issue of how browsing queries are implemented with tolerable performance. This question is beyond the scope of this paper, which focusses on specification rather than implementation. The interested reader is referred to [Che95] for related work on the optimization and efficient implementation of browsing queries.

## 10 Conclusion

This paper has described the design of Tioga-2. It attempts to rectify the deficiencies of the original Tioga system in the following areas:

- Too hard to use.

Tioga-2 is a direct manipulation system whereby the user makes incremental changes via direct manipulation of data. Each change results in a new visualization on the canvas. As such, we expect it to be much more accessible to untrained programmers.

- Better programming environment.

Tioga-2 allows a viewer to be installed on any arc in a diagram. It is easy to instrument a program to understand how it is working and to see visually where it fails.

- Functionality

Tioga-2 has multiple output boxes, T’s, customizable viewers, and rear view mirrors, all of which are absent from Tioga. In addition, Tioga-2 provides a more systematic treatment of magnifying glasses, wormholes, and updates.

We are now hard at work implementing Tioga-2 and expect to have an initial system by the end of the year. We plan to systematically test the implementation on little programmers to ascertain whether it lives up to its goals.

---

<sup>2</sup>In fairness, Magic Lenses is not intended strictly as a browsing paradigm, but as a general user interface paradigm.

## Reference

- [BSP<sup>+</sup>93] E. Bier, M. Stone, K. Pier, W. Buxton, and T. DeRose. Toolglass and magic lenses: The see-through interface. In *Proceedings of SIGGRAPH 1993*, pages 73–80, Anaheim, CA, August 1993.
- [CGR91] P. Cox, M. Gorlick, and R. Razouk. Using weaves for software construction and analysis. In *Proceedings of the 13th International Conference on Software Engineering*, pages 23–34, Austin, TX, May 1991.
- [Che95] Jolly Chen. Optimizing interactive browsing queries. June 1995. Submitted for publication.
- [Cyp93] Allen Cypher. *Watch What I Do: Programming by Demonstration*. MIT Press, Cambridge, MA, 1993.
- [GQ94] M. Gorlick and A. Quilici. Visual programming-in-the-large versus visual programming-in-the-small. In *Proceedings of the IEEE Symposium on Visual Languages*, pages 137–144, St. Louis, MO, October 1994.
- [KZ95] R. Krishnamurthy and M. Zloof. RBE: Rendering by example. In *Proceedings of the 11th International Conference on Data Engineering*, pages 288–297, Taipei, Taiwan, March 1995.
- [LAC<sup>+</sup>92] B. Lucas, G.D. Abram, N.S. Collins, D.A. Epstein, et al. An architecture for a scientific visualization system. In *Proceedings of the IEEE Visualization Conference*, pages 107–14, Boston, MA, October 1992.
- [PF93] K. Perlin and D. Fox. Pad: An alternative approach to the computer interface. In *Proceedings of SIGGRAPH*, pages 57–64, Anaheim, CA, August 1993.
- [RY92] J. Rasure and M. Young. An open environment for image processing software development. In *Proceedings of the SPIE Symposium on Electronic Image Processing*, pages 300–310, San Jose, CA, February 1992.
- [SAD<sup>+</sup>93] M. Stonebraker, R. Agrawal, U. Dayal, E. Neuhold, and A. Reuter. DBMS research at a crossroads: The Vienna update. In *Proceedings of the 19th International Conference on Very Large Data Bases*, pages 688–692, Dublin, Ireland, August 1993.
- [SCN<sup>+</sup>93] M. Stonebraker, J. Chen, N. Nathan, C. Paxson, and J. Wu. Tioga: Providing data management support for scientific visualization applications. In *Proceedings of the 19th International Conference on Very Large Data Bases*, pages 25–38, Dublin, Ireland, August 1993.
- [SS94] M. Sarkar and M. Sarkar. Graphical fisheye views. *Communications of the ACM*, pages 73–84, December 1994.
- [U<sup>+</sup>89] C. Upson et al. The application visualization system. *IEEE Computer Graphics and Applications*, 9(4):30–42, July 1989.
- [WWT<sup>+</sup>94] A. Woodruff, P. Wisnovsky, C. Taylor, M. Stonebraker, C. Paxson, J. Chen, and A. Aiken. Zooming and tunneling in Tioga: Supporting navigation in multidimensional space. In *Proceedings of the IEEE Symposium on Visual Languages*, pages 191–193, St. Louis, MO, October 1994.