UIMSs, Toolkits, Interface Builders[1]

Brad A. Myers
Human Computer Interaction Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
bam@cs.cmu.edu
http://www.cs.cmu.edu/~bam
(412) 268-5150

May 24, 1996

1. Introduction

User interface software is often large, complex and difficult to implement, debug, and modify. A 1992 study found that an average of 48% of the code of applications is devoted to the user interface, and that about 50% of the implementation time is devoted to implementing the user interface portion [Myers 92a], and the numbers are probably much higher today. As interfaces become easier to use, they become harder to create [Myers 94]. Today, direct manipulation interfaces (also called ‘‘GUIs’’ for Graphical User Interfaces) are almost universal: one 1993 study found that 97% of all software development on Unix involved a GUI [XBusiness 94, p.80]. These interfaces require that the programmer deal with elaborate graphics, multiple ways for giving the same command, multiple asynchronous input devices (usually a keyboard and a pointing device such as a mouse), a ‘‘mode free’’ interface where the user can give any command at virtually any time, and rapid ‘‘semantic feedback’’ where determining the appropriate response to user actions requires specialized information about the objects in the program. Tomorrow’s user interfaces will provide speech and gesture recognition, intelligent agents and integrated multi-media, and will probably be even more difficult to create. Furthermore, because user interface *design* is so difficult, the only reliable way to get good interfaces is to iteratively re-design (and therefore re-implement) the interfaces after user-testing, which makes the implementation task even harder.

Fortunately, there has been significant progress in software tools to help with creating user interfaces, and today, virtually all user interface software is created using tools that make the implementation easier. For example, the MacApp system from Apple was reported to reduce development time by a factor of four or five [Wilson 90]. A study commissioned by NeXT claimed that the average application programmed using the NeXTStep environment wrote 83% fewer lines of code and took one-half the time compared to applications written using less advanced tools, and some applications were completed in

---

[1]This paper is revised from an earlier version which appeared as: Brad A. Myers. ‘‘User Interface Software Tools,’’ *ACM Transactions on Computer-Human Interaction*. vol. 2, no. 1, March, 1995. pp. 64-103.

one-tenth the time [Booz 92].

Furthermore, user interface tools are a major business. In the Unix market alone, over US$133 million of tools were sold in 1993, which is about 50,000 licenses [XBusiness 94]. Forrester Research claims that the total market for UI software tools on all platforms, including ‘‘vertical tools’’ which include database and user interface construction tools, was 130,000 developers generating US$400 million in revenue in 1993 [DePalma 93]. Whereas the ‘‘first generation’’ of commercial tools were not fully graphical or not sufficiently powerful, this is no longer true for today’s tools. Furthermore, prices for tools have dropped significantly, and fees for run-times have been mostly eliminated (so that designers do not have to pay the tool creator for products created using the tools). For the future, there is still a tremendous opportunity for good tools, especially in niche areas like multimedia, 3D, virtual reality, distributed systems, and geographical information systems.

This article surveys user interface software tools, and explains the different types and classifications. However, it is now impossible to discuss *all* user interface tools, since there are so many.[2] For example, there are over 100 commercial graphical user Interface Builders, and many new research tools are reported every year at conferences like the annual ACM User Interface Software and Technology Symposium (UIST)[3] and the ACM SIGCHI conference.[4] There are also about three PhD theses on user interface tools every year. Therefore, this article provides an overview of the most popular approaches, rather than an exhaustive survey.

2. Definitions

The **user interface** (UI) of a computer program is the part that handles the output to the display and the input from the person using the program. The rest of the program is called the **application** or the **application semantics**.

User interface tools have been called various names over the years, with the most popular being **User Interface Management Systems (UIMS)** [Olsen 92]. However, many people feel that the term UIMS should be used only for tools that handle the sequencing of operations (what happens after each event from the user), so other terms like **Toolkits**, **User Interface Development Environments**, **Interface Builders**, **Interface Development Tools**, and **Application Frameworks** have been used. This paper will

[2]A comprehensive list which is frequently updated is available through the World-Wide-Web as `http://www.cs.cmu.edu/~bam/toolnames.html`

[3]See http://www.acm.org/uist/

[4]For example, see http://www.acm.org/sigchi/chi96/

try to define these terms more specifically, and use the general term ‘‘user interface tool’’ for all software aimed to help create user interfaces. Note that the word ‘‘tool’’ is being used to include what are called ‘‘toolkits,’’ as well as higher-level tools, such as Interface Builders, that are *not* toolkits.

Four different classes of people are involved with user interface software, and it is important to have different names for them to avoid confusion. The first is the person using the resulting program, who is called the **end user** or just **user**. The next person creates the user interface of the program, and is called the **user interface designer** or just **designer**. Working with the user interface designer will be the person who writes the software for the rest of the application. This person is called the **application programmer**. The designer may use special user interface tools which are provided to help create user interfaces. These tools are created by the **tool creator**. Note that the designer will be a user of the software created by the tool creator, but we still do not use the term ‘‘user’’ here to avoid confusion with the end user. Although this classification discusses each role as a different person, in fact, there may be many people in each role or one person may perform multiple roles. The general term **programmer** is used for anyone who writes code, and may be a designer, application programmer, or tool creator.

3. Importance of User Interface Tools

There are many advantages to using user interface software tools. These can be classified into two main groups:

- **The quality of the interfaces might be higher.** This is because:
  - Designs can be rapidly prototyped and implemented, possibly even before the application code is written.
  - It is easier to incorporate changes discovered through user testing.
  - More effort can be expended on the tool than may be practical on any single user interface since the tool will be used with many different applications.
  - Different applications are more likely to have consistent user interfaces if they are created using the same user interface tool.
  - It will be easier for a variety of specialists to be involved in designing the user interface, rather than having the user interface created entirely by programmers. Graphic artists, cognitive psychologists, and human factors specialists may all be involved. In particular, professional user interface designers, who may not be programmers, can be in charge of the overall design.
  - Undo, Help and other features are more likely to be available since they might be supported by the tools.
- **The user interface code might be easier and more economical to create and maintain.** This is because:
  - Interface specifications can be represented, validated, and evaluated more easily.
  - There will be less code to write, because much is supplied by the tools.
  - There will be better modularization due to the separation of the user interface

component from the application. This should allow the user interface to change without affecting the application, and a large class of changes to the application (such as changing the internal algorithms) should be possible without affecting the user interface.

- The level of programming expertise of the interface designers and implementors can be lower, because the tools hide much of the complexities of the underlying system.
- The reliability of the user interface will be higher, since the code for the user interface is created automatically from a higher level specification.
- It will be easier to port an application to different hardware and software environments since the device dependencies are isolated in the user interface tool.

Based on these goals for user interface software tools, we can list a number of important functions that should be provided. This list can be used to evaluate the various tools to see how much they cover. Naturally, no tool will help with everything, and different user interface designers may put different emphasis on the different features.

In general, the tools might:

- help *design* the interface given a specification of the end users' tasks,
- help *implement* the interface given a specification of the design,
- help *evaluate* the interface after it is designed and propose improvements, or at least provide information to allow the designer to evaluate the interface,
- create easy-to-use interfaces,
- allow the designer to rapidly investigate different designs,
- allow non-programmers to design and implement user interfaces,
- allow the end user to customize the interface,
- provide portability, and
- be easy to use themselves.

This might be achieved by having the tools:

- automatically choose which user interface styles, input devices, widgets, etc. should be used,
- help with screen layout and graphic design,
- validate user inputs,
- handle user errors,
- handle aborting and undoing of operations,
- provide appropriate feedback to show that inputs have been received,
- provide help and prompts,
- update the screen display when application data changes,
- notify the application when the user modifies graphical objects,
- handle field scrolling and editing,
- help with the sequencing of operations,
- insulate the application from all device dependencies and the underlying software and hardware systems,
- provide customization facilities to end users, and
- evaluate the graphic design and layout, usability, and learnability of the interface.

4. Overview of User Interface Software Tools

Since user interface software is so difficult to create, it is not surprising that people have been working for a long time to create tools to help with it. Today, many of these tools and ideas have progressed from research into commercial systems, and their effectiveness has been amply demonstrated. Research systems also continue to evolve quickly, and the models that were popular five years ago have been made obsolete by more effective tools, changes in the computer market (e.g., the demise of OpenLook has taken with it a number of tools), and the emergence of new styles of user interfaces such as pen-based computing and multi-media.

4.1 Components of User Interface Software

As shown in Figure 1, user interface software may be divided into various layers: the windowing system, the toolkit and higher-level tools. Of course, many practical systems span multiple layers.

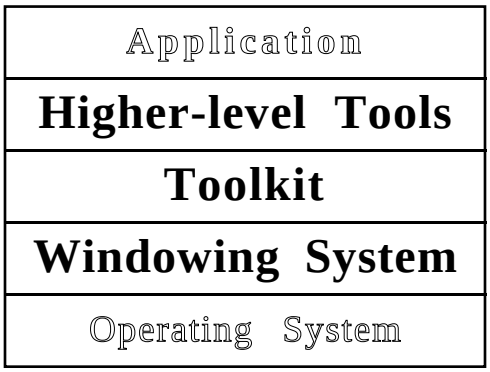

Figure 1: The components of user interface software discussed in this article.

The **windowing system** supports the separation of the screen into different (usually rectangular) regions, called *windows*. The X system [Scheifler 86] divides the window functionality into two layers: the window **system**, which is the functional or programming interface, and the window **manager** which is the user interface. Thus the ‘‘window system’’ provides procedures that allow the application to draw pictures on the screen and get input from the user, and the ‘‘window manager’’ allows the end user to move windows around, and is responsible for displaying the title lines, borders and icons around the windows. However, many people and systems use the name ‘‘window manager’’ to refer to both layers, since systems such as the Macintosh and Microsoft Windows do not separate them. This article will use the X terminology, and use the term ‘‘windowing system’’ when referring to both layers.

On top of the windowing system is the **toolkit**, which contains many commonly used **widgets** such as menus, buttons, scroll bars, and text input fields. On top of the toolkit might be **higher-level tools**, which

help the designer use the toolkit widgets. The following sections discuss each of these components in more detail.

5. Windowing Systems

A windowing system is a software package that helps the user monitor and control different contexts by separating them physically onto different parts of one or more display screens. A survey of various windowing systems was published earlier [Myers 88a]. Although most of today's systems provide toolkits on top of the windowing systems, as will be explained below, toolkits generally only address the drawing of widgets such as buttons, menus and scroll bars. Thus, when the programmer wants to draw application-specific parts of the interface and allow the user to manipulate these, the window system interface must be used directly. Therefore, the windowing system's programming interface has significant impact on most user interface programmers.

The first windowing systems were implemented as part of a single program or system. For example, the EMACs text editor [Stallman 79], and the Smalltalk [Tesler 81] and DLISP [Teitelman 79] programming environments had their own windowing systems. Later systems implemented the windowing system as an integral part of the operating system, such as Sapphire for PERQs [Myers 84], SunView for Suns, and the Macintosh and Microsoft Windows systems. In order to allow different windowing systems to operate on the same operating system, some windowing systems, such as X and Sun's NeWS, operate as a separate process, and use the operating system's inter-process communication mechanism to connect to applications.

5.1 Structure of Windowing Systems

A windowing system can be logically divided into two layers, each of which has two parts (see Figure 2). The **window system**, or **base layer**, implements the basic functionality of the windowing system. The two parts of this layer handle the display of graphics in windows (the **output model**) and the access to the various input devices (the **input model**), which usually includes a keyboard and a pointing device such as a mouse. The primary interface of the base layer is procedural, and is called the windowing system's **application programmer interface (API)**.

The other layer of windowing system is the **window manager** or **user interface**. This includes all aspects that are visible to the user. The two parts of the user interface layer are the **presentation**, which is comprised of the pictures that the window manager displays, and the **commands**, which are how the user manipulates the windows and their contents.

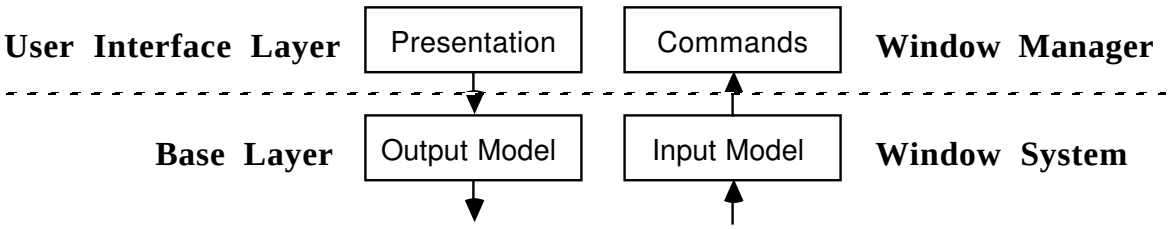

Figure 2: The windowing system can be divided into two layers, called the **base** or **window system** layer, and the **user interface** or **window manager** layer. Each of these can be divided into parts that handle output and input.

5.2 Base Layer

The base layer is the procedural interface to the windowing system. In the 1970s and early 1980s, there were a large number of different windowing systems, each with a different procedural interface (at least one for each hardware platform). People writing software found this to be unacceptable because they wanted to be able to run their software on different platforms, but they would have to rewrite significant amounts of code to convert from one window system to another. The X windowing system [Scheifler 86] was created to solve this problem by providing a hardware-independent interface to windows. X has been quite successful at this, and has driven virtually all other windowing systems out of the workstation hardware market. In the small computer market, the Macintosh runs its own window system, and IBM PC-class machines primarily run Microsoft Windows or IBM's Presentation Manager (part of OS/2).

5.2.1 Output Model

The output model is the set of procedures that an application can use to draw pictures on the screen. It is important that all output be directed through the window system so that the graphics primitives can be clipped to the window's borders. For example, if a program draws a line that would extend out of a window's borders, it must be clipped so that the contents of other, independent, windows are not overwritten. These operations can be much quicker, but are very dangerous and therefore should seldom be used. Most modern computers provide graphics hardware that is optimized to work efficiently with the window system.

In early windowing systems, such as Smalltalk [Tesler 81] and Sapphire [Myers 86], the primary output operation was BitBlt (also called ''RasterOp''). These systems primarily supported monochrome screens (each pixel is either black or white). BitBlt takes a rectangle of pixels from one part of the screen and copies it to another part. Various boolean operations can be specified for combining the pixel values of the source and destination rectangles. For example, the source rectangle can simply replace the

destination, or it might be XORed with the destination. BitBlt can be used to draw solid rectangles in either black or white, display text, scroll windows, and perform many other effects [Tesler 81]. The only additional drawing operation typically supported by these early systems was drawing straight lines.

Later windowing systems, such as the Macintosh and X, added a full set of drawing operations, such as filled and unfilled polygons, text, lines, arcs, etc. These cannot be implemented using the BitBlt operator. With the growing popularity of color screens and non-rectangular primitives (such as rounded rectangles), the use of BitBlt has significantly decreased. It is primarily used now for scrolling and copying off-screen pictures onto the screen (e.g., to implement double-buffering).

A few windowing systems allow the full Postscript imaging model to be used to create images on the screen. Postscript provides device-independent coordinate systems and arbitrary rotations and scaling for all objects, including text. Another advantage of using Postscript for the screen is that the same language can be used to print the windows on paper (since many printers accept Postscript). Sun created a version used in the NeWS windowing system, and then Adobe (the creator of Postscript) came out with an official version called ‘‘Display Postscript’’ which is used in the NeXT windowing system and is supplied as an extension to the X windowing system by a number of vendors, including DEC and IBM.

All of the standard output models only contain drawing operations for two dimensional objects. Two extensions to support 3-D objects are PEX and OpenGL. PEX [Gaskins 92] is an extension to the X windowing system that incorporates much of the PHIGS graphics standard. OpenGL [OpenGL 93] is based on the GL programming interface that has been used for many years on Silicon Graphics machines. OpenGL provides machine independence for 3-D since it is available for various X platforms (SGI, Sun, etc.) and is included as a standard part of Microsoft Windows NT.

As shown in Figure 3, the earlier windowing systems assumed that a graphics package would be implemented using the windowing system. For example, the CORE graphics package was implemented on top of the SunView windowing system. All newer systems, including the Macintosh, X, NeWS, NeXT, and Microsoft Windows, have implemented a sophisticated graphics system as *part* of the windowing system.

5.2.2 Input Model

The early graphics standards, such as CORE and PHIGS, provided an input model that does not support the modern, direct manipulation style of interfaces. In those standards, the programmer calls a routine to request the value of a ‘‘virtual device’’ such as a ‘‘locator’’ (pointing device position), ‘‘string’’ (edited text string), ‘‘choice’’ (selection from a menu), or ‘‘pick’’ (selection of a graphical object). The program would then pause waiting for the user to take action. This is clearly at odds with the direct manipulation

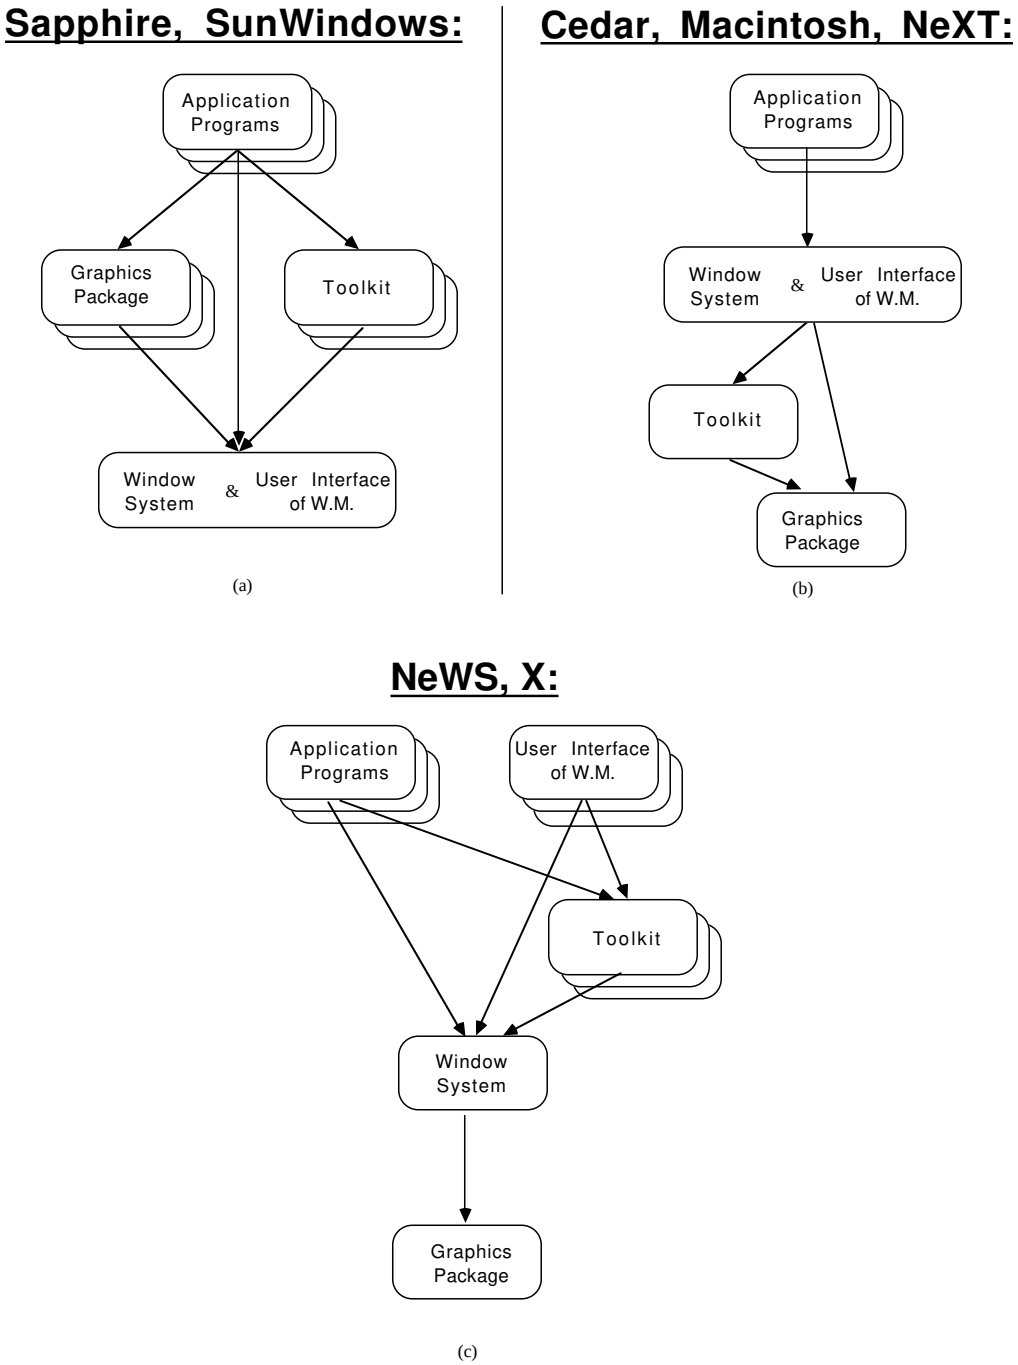

Figure 3: Various organizations that have been used by windowing systems. Boxes with extra borders represent systems that can be replaced by users. Early systems (a) tightly coupled the window manager and the window system, and assumed that sophisticated graphics and toolkits would be built on top. The next step in designs (b) was to incorporate into the windowing system the graphics and toolkits, so that the window manager itself could have a more sophisticated look and feel, and so applications would be more consistent. Other systems (c) allow different window managers and different toolkits, while still embedding sophisticated graphics packages.

‘‘mode-free’’ style, where the user can decide whether to make a menu choice, select an object, or type something.

With the advent of modern windowing systems, a new model was provided: a stream of event records is sent to the window which is currently accepting input. The user can select which window is getting events using various commands, described below. Each event record typically contains the type and value of the event (e.g., which key was pressed), the window to which the event was directed, a timestamp, and the x and y coordinates of the mouse. The windowing system queues keyboard events, mouse button events, and mouse movement events together (along with other special events) and programs must dequeue the events and process them. It is somewhat surprising that, although there has been substantial progress in the output model for windowing systems (from BitBlt to complex 2-D primitives to 3-D), input is still handled in essentially this same way today as in the original windowing systems, even though there are some well-known unsolved problems with this model:

- There is no provision for special stop-output (control-S) or abort (control-C, command-dot) events, so these will be queued with the other input events.
- The same event mechanism is used to pass special messages from the windowing system to the application. When a window gets larger or becomes uncovered, the application must usually be notified so it can adjust or redraw the picture in the window. Most window systems communicate this by enqueuing special events into the the event stream, which the program must then handle.
- The application must always be willing to accept events in order to process aborts and redrawing requests. If not, then long operations cannot be aborted, and the screen may have blank areas while they are being processed.
- The model is device dependent, since the event record has fixed fields for the expected incoming events. If a 3-D pointing device or one with more than the standard number of buttons is used instead of a mouse, then the standard event mechanism cannot handle it.
- Because the events are handled asynchronously, there are many race conditions that can cause programs to get out of synchronization with the window system. For example, in the X windowing system, if you press inside a window and release outside, under certain conditions the program will think that the mouse button is still depressed. Another example is that refresh requests from the windowing system specify a rectangle of the window that needs to be redrawn, but if the program is changing the contents of the window, the wrong area may be redrawn by the time the event is processed. This problem can occur when the window is scrolled.

Although these problems have been known for a long time, there has been little research on new input models (an exception is the Garnet Interactors model [Myers 90a]).

5.2.3 Communication

In the X windowing system and NeWS, all communication between applications and the window system uses inter-process communication through a network protocol. This means that the application program can be on a different computer from its windows. In all other windowing systems, operations are implemented by directly calling the window manager procedures or through special traps into the operating system. The primary advantage of the X mechanism is that it makes it easier for a person to utilize multiple machines with all their windows appearing on a single machine. Another advantage is that it is easier to provide interfaces for different programming languages: for example the C interface (called xlib) and the Lisp interface (called CLX) send the appropriate messages through the network protocol. The primary disadvantage is efficiency, since each window request will typically be encoded, passed to the transport layer, and then decoded, even when the computation and windows are on the same machine.

5.3 User Interface Layer

The user interface of the windowing system allows the user to control the windows. In X, the user can easily switch user interfaces, by killing one window manager and starting another. Popular window managers under X include uwm (which has no title lines and borders), twm, mwm (the Motif window manager), and olwm (the OpenLook window manager). There is a standard protocol through which programs and the base layer communicate to the window manager, so that all programs continue to run without change when the window manager is switched. It is possible, for example, to run applications that use Motif widgets inside the windows controlled by the OpenLook window manager.

A discussion of the options for the user interfaces of window managers was previously published [Myers 88a]. Also, the video *All the Widgets* [Myers 90b] has a 30 minute segment showing many different forms of window manager user interfaces.

Some parts of the user interface of a windowing system, which is sometimes called its ‘‘look and feel,’’ can apparently be copyrighted and patented. Which parts is a highly complex issue, and the status changes with decisions in various court cases [Samuelson 93].

5.3.1 Presentation

The **presentation** of the windows defines how the screen looks. One very important aspect of the presentation of windows is whether they can **overlap** or not. Overlapping windows, sometimes called **covered** windows, allow one window to be partially or totally on top of another window, as shown in Figure 4. This is also sometimes called the **desktop metaphor**, since windows can cover each other like

pieces of paper can cover each other on a desk.[5] The other alternative is called **tiled** windows, which means that windows are not allowed to cover each other. Figure 5 shows an example of tiled windows. Obviously, a window manager that supports covered windows can also allow them to be side-by-side, but not vice-versa. Therefore, a window manager is classified as ‘‘covered’’ if it allows windows to overlap. The tiled style was popular for a while, and was used by Cedar [Swinehart 86], and early versions of the Star [Smith 82], Andrew [Palay 88], and Microsoft Windows. A study even suggested that using tiled windows was more efficient for users [Bly 86]. However, today tiled windows are rarely seen on conventional window systems, because users generally prefer overlapping.

Modern ‘‘browsers’’ for the World-Wide Web, such as Mosaic, Netscape, and Microsoft’s Internet Explorer, provide a windowing environment inside the computer’s main windowing system. Newer versions of browsers support pages containing multiple ‘‘panes,’’ which are a form of tiled window. In addition, if an application written in Java is downloaded (see the section on Virtual Toolkits below), it can create multiple, overlapping windows like conventional GUI applications. Many people are claiming that browsers will be the desktop of the next generation of systems and replace today’s windowing systems.

Another important aspect of the presentation of windows is the use of **icons** (shown in Figure 4 and 5). These are small pictures that represent windows (or sometimes files). They are used because there would otherwise be too many windows to conveniently fit on the screen and manage. Other aspects of the presentation include whether the window has a title line or not, what the background (where there are no windows) looks like, and whether the title and borders have control areas for performing window operations.

5.3.2 Commands

Since computers typically have multiple windows and only one mouse and keyboard, there must be a way for the user to control which window is getting keyboard input. This window is called the **input** (or **keyboard**) **focus**. Another term is the **listener** since it is listening to the user’s typing. Some systems called the focus the ‘‘active window’’ or ‘‘current window,’’ but these are poor terms since in a multi-processing system, many windows can be actively outputting information at the same time. Window managers provide various ways to specify and show which window is the listener. The most important options are:

- **click-to-type**, which means that the user must click the mouse button in a window before typing to it. This is used by the Macintosh.
- **move-to-type**, which means that the mouse only has to move over a window to allow typing

[5]There are usually other aspects to the desktop metaphor, however, such as presenting file operations in a way that mimics office operations, as in the Star office workstation [Smith 82].

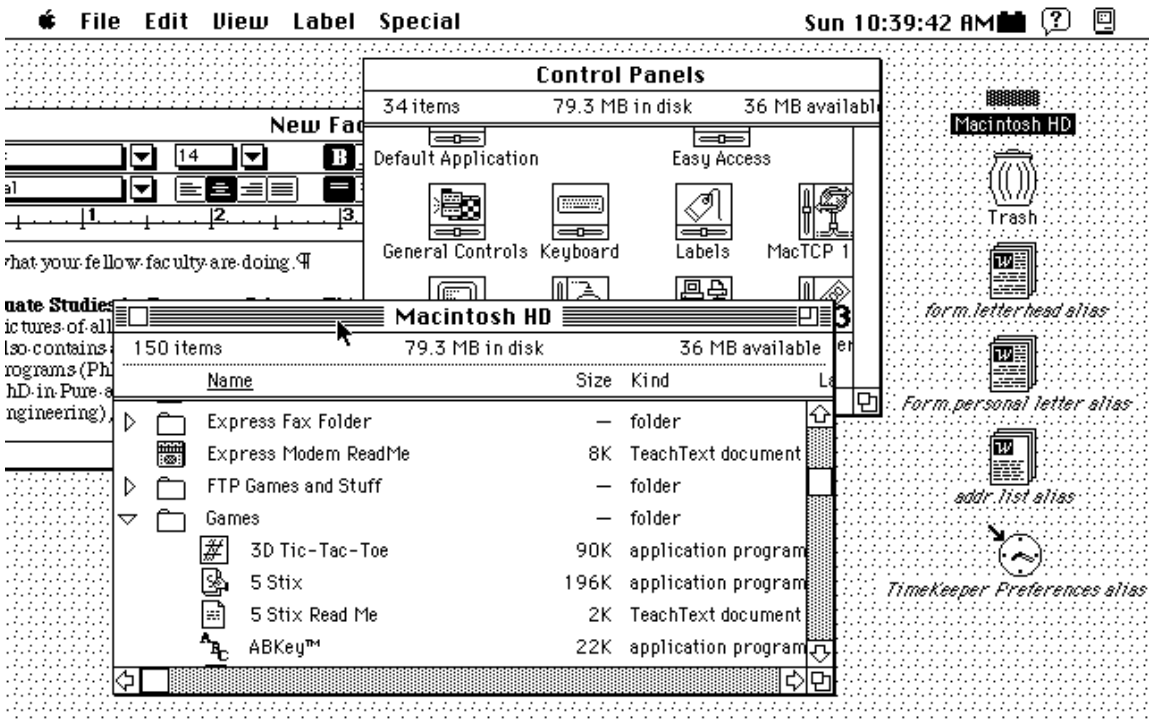

Figure 4: A screen from the Macintosh showing 3 windows covering each other, and some icons along the right margin.

to it. This is usually faster for the user, but may cause input to go to the wrong window if the user accidentally knocks the mouse.

Most X window managers (including the Motif and OpenLook window managers) allow the user to choose which method is desired. However, the choice can have significant impact on the user interface of applications. For example, because the Macintosh requires click-to-type, it can provide a single menu-bar at the top, and the commands can always operate on the focussed window. With move-to-type, the user might have to pass through various windows (thus giving them the focus) on the way to the top of the screen. Therefore, Motif applications must have a menubar in each window so the commands will know which window to operate on.

All covered window systems allow the user to change which window is on top (not covered by other windows), and usually to send a window to the bottom (covered by all other windows). Other commands allow windows to be changed size, moved, created and destroyed.

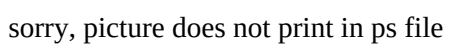

Figure 5: A screen from the Cedar [Swinehart 86] windowing system. Windows are ‘‘tiled’’ into 2 columns. There is a row of icons along the bottom. Each window has a fixed menu of commands below the title line.

6. Toolkits

A **toolkit** is a library of ‘‘widgets’’ that can be called by application programs. A **widget** is a way of using a physical input device to input a certain type of value. Typically, widgets in toolkits include menus, buttons, scroll bars, text type-in fields, etc. Figure 6 shows some examples of widgets. Creating an interface using a toolkit can only be done by programmers, because toolkits only have a procedural interface.

Using a toolkit has the advantage that the final UI will look and act similarly to other UIs created using the same toolkit, and each application does not have to re-write the standard functions, such as menus. A problem with toolkits is that the styles of interaction are limited to those provided. For example, it is difficult to create a single slider that contains two indicators, which might be useful to input the upper and lower bounds of a range. In addition, the toolkits themselves are often expensive to create: ‘‘The primitives never seem complex in principle, but the programs that implement them are surprisingly

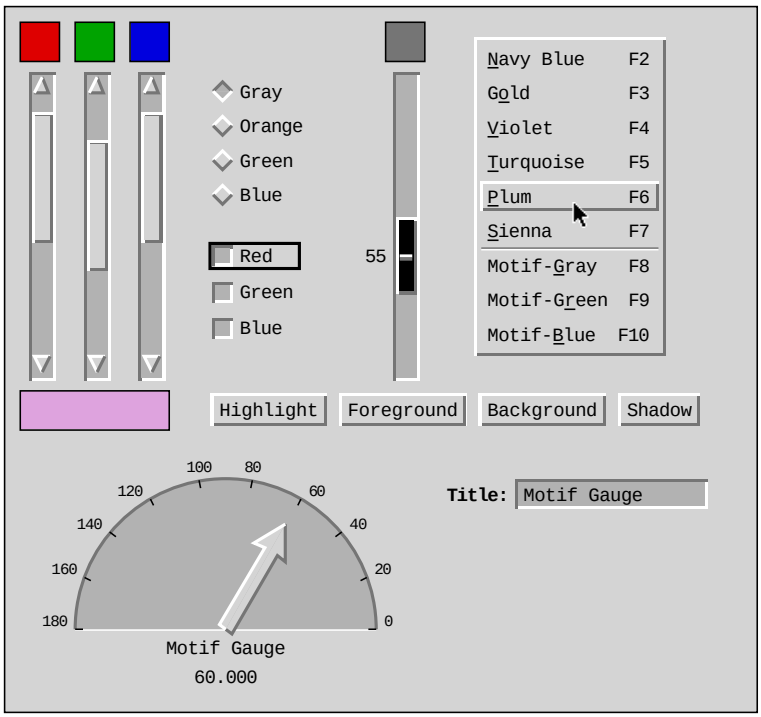

Figure 6: Some of the widgets with a Motif look-and-feel provided by the Garnet toolkit.

intricate'' [Cardelli 85, p.199]. Another problem with toolkits is that they are often difficult to use since they may contain hundreds of procedures, and it is often not clear how to use the procedures to create a desired interface. For example, the documentation for the Macintosh Toolbox now is well over six books, of which about 1/3 is related to user interface programming.

As with the graphics package, the toolkit can be implemented either using or being used by the windowing system (see Figure 3). Early systems provided only minimal widgets (e.g., just a menu), and expected applications to provide others. In the Macintosh, the toolkit is at a low level, and the window manager user interface is built using it. The advantage of this is that the window manager can then use the same sophisticated toolkit routines for its user interface. When the X system was being developed, the developers could not agree on a single toolkit, so they left the toolkit to be on top of the windowing system. In X, programmers can use a variety of toolkits (for example, the Motif, OpenLook, InterViews [Linton 89], Garnet [Myers 90c] or tk [Ousterhout 91] toolkits can be used on top of X), but the window manager must usually implement its user interface from scratch.

Because the designers of X could not agree on a single look-and-feel, they created an **intrinsics** layer on which to build different widget sets, which they called xt [McCormack 88]. This layer provides the common services, such as techniques for object-oriented programming and layout control. The **widget**

set layer is the collection of widgets that is implemented using the intrinsics. Multiple widget sets with different looks and feels can be implemented on top of the same intrinsics layer (Figure 7-a), or else the same look-and-feel can be implemented on top of different intrinsics (Figure 7-b). When Sun announced that it was phasing out OpenLook, the Motif widget set became the standard for X and xt.

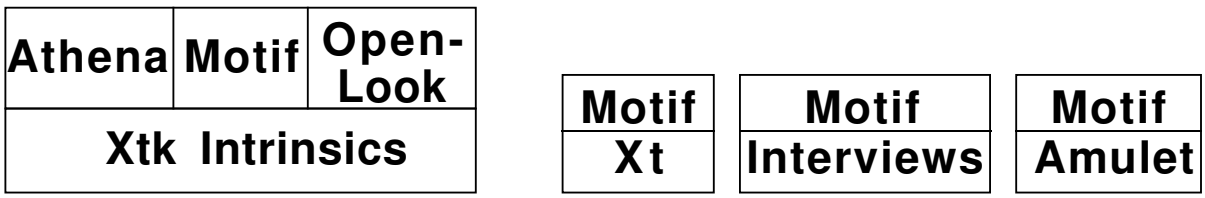

(a) (b)

Figure 7: (a) At least three different widget sets that have different looks and feels have been implemented on top of the xt intrinsics. (b) The Motif look-and-feel has been implemented on many different intrinsics.

6.1 Toolkit Intrinsics

Toolkits come in two basic varieties. The most conventional is simply a collection of procedures that can be called by application programs. Examples of this style include the SunTools toolkit for the SunView windowing system, and the Macintosh Toolbox [Apple 85]. The other variety uses an object-oriented programming style which makes it easier for the designer to customize the interaction techniques. Examples include Smalltalk [Tesler 81], Andrew [Palay 88], Garnet [Myers 90c], InterViews [Linton 89], Xt [McCormack 88], Amulet [Myers 96a], and the Java toolkit AWT [Sun 96].

The advantages of using object-oriented intrinsics are that it is a natural way to think about widgets (the menus and buttons on the screen *seem* like objects), the widget objects can handle some of the chores that otherwise would be left to the programmer (such as refresh), and it is easier to create custom widgets (by sub-classing an existing widget). The advantage of the older, procedural style is that it is easier to implement, no special object-oriented system is needed, and it is easier to interface to multiple programming languages.

To implement the objects, the toolkit might invent its own object system, as was done with Xt, Andrew, Garnet and Amulet, or it might use an existing object system, as was done in InterViews [Linton 89] which uses C++, NeXTStep from NeXT which uses Objective-C, and AWT uses Java [Sun 96].

The usual way that object-oriented toolkits interface with application programs is through the use of **call-back procedures**. These are procedures defined by the application programmer that are called when a widget is operated by the end user. For example, the programmer might supply a procedure to be called

when the user selects a menu item. Experience has shown that real interfaces often contain hundreds of call-backs, which makes the code harder to modify and maintain [Myers 92a]. In addition, different toolkits, even when implemented on the same intrinsics like Motif and OpenLook, have different call-back protocols. This means that code for one toolkit is difficult to port to a different toolkit. Therefore, research is being directed at reducing the number of call-backs in user interface software [Myers 91a].

Some research toolkits have added novel features to the toolkit intrinsics. For example, Garnet [Myers 90c], Rendezvous [Hill 93], Bramble [Gleicher 93] and Amulet [Myers 96a], allow the objects to be connected using **constraints**, which are relationships that are declared once and then maintained automatically by the system. For example, the designer can specify that the color of a rectangle is constrained to be the value of a slider, and then the system will automatically update the color if the user moves the slider.

6.2 Widget Set

Typically, the intrinsics layer is look-and-feel independent, which means that the widgets built on top of it can have any desired appearance and behavior. However, a particular widget set must pick a look-and-feel. The video *All the Widgets* shows many examples of widgets that have been designed over the years [Myers 90b]. For example, it shows 35 different kinds of menus. Like window manager user interfaces, the widgets' look-and-feel can be copyrighted and patented [Samuelson 93].

As was mentioned above, different widget sets (with different looks and feels) can be implemented on top of the same intrinsics. In addition, the same look-and-feel can be implemented on top of different intrinsics. For example, there are Motif look-and-feel widgets on top of the xt, InterViews and Amulet intrinsics (Figure 7-b). Although they all look and operate the same (so would be indistinguishable to the end user), they are implemented quite differently, and have completely different procedural interfaces for the programmer.

6.3 Specialized Toolkits

A number of toolkits have been developed to support specific kinds of applications or specific classes of programmers. For example, the SUIT system [Pausch 92] (which contains a toolkit and an interface builder), is specifically designed to be easy to learn and is aimed at classroom instruction. Amulet [Myers 96a] provides high-level support for graphical, direct manipulation interfaces, and handles input as hierarchical *command objects*, making Undo easier to implement [Myers 96b]. Rendezvous [Hill 93] and Visual Obliq [Bharat 94] are designed to make it easier to create applications that support multiple users on multiple machines operating synchronously. Whereas most toolkits provide only 2-D interaction techniques, the Brown 3-D toolkits [Zeleznik 91, Stevens 94] and Silicon Graphics' Inventor toolkit

[Strauss 92, Wernecke 94] provide preprogrammed 3-D widgets and a framework for creating others. Special support for animations has been added to Artkit [Hudson 93] and Amulet [Myers 96c]. Tk [Ousterhout 91] is a popular toolkit for the X window system because it uses an interpretive language called tcl which makes it possible to dynamically change the user interface. Tcl also supports the Unix style of programming where many small programs are glued together.

7. Virtual Toolkits

Although there are many small differences among the various toolkits, much remains the same. For example, all have some type of menu, button, scroll bar, text input field, etc. Although there are fewer windowing systems and toolkits than there were ten years ago, people are still finding that it to be a lot of work to convert software from Motif to the Macintosh and to Microsoft Windows.

Therefore, a number of systems have been developed that try to hide the differences among the various toolkits, by providing **virtual** widgets which can be mapped into the widgets of each toolkit. Another name for these tools is **cross-platform development systems**. The programmer writes the code once using the virtual toolkit and the code will run without change on different platforms and still look like it was designed with that platform's widgets. For example, the virtual toolkit might provide a single menu routine, which always has the same programmer interface, but connects to a Motif menu, Macintosh menu or a Windows menu depending on which machine the application is run on.

There are two styles of virtual toolkits. In one, the virtual toolkit links to the different *actual* toolkits on the host machine. For example, XVT [XVT 96] provides a C or C++ interface that links to the actual Motif, OpenLook, Macintosh, MS-Windows, and OS/2-PM toolkits (and also character terminals) and hides their differences. The second style of virtual toolkit *re-implements* the widgets in each style. For example, Galaxy [Visix 96], and Open Interface from NeuronData [Open Interface 95], provide libraries of widgets that look like those on the various platforms. The advantage of the first style is that the user interface is more likely to be look-and-feel conformant (since it uses the real widgets). The disadvantages are that the virtual toolkit must still provide an interface to the graphical drawing primitives on the platforms. Furthermore, they tend to only provide functions that appear in all toolkits. Many of the virtual toolkits that take the second approach, for example Galaxy, provide a sophisticated graphics package and complete sets of widgets on all platforms. However, with the second approach, there must always be a large run-time library, since in addition to the built-in widgets that are native to the machine, there is the re-implementation of these same widgets in the virtual toolkit's library.

All of the toolkits that work on multiple platforms can be considered virtual toolkits of the second type. For example, SUIT [Pausch 92] and Garnet [Myers 90c] work on X, Macintosh and Windows. However,

these use the same look-and-feel on all platforms (and therefore do not look the same as the other applications on that platform), so they are not classified as virtual toolkits.

The AWT toolkit that comes with the Java programming language [Sun 96] also can be classified as a Virtual Toolkit, since the programmer can write code once and it will operate on all platforms. Java programs can be run locally in a conventional fashion, or can be downloaded dynamically over the world-wide-web into a browser such as Netscape.

8. Higher Level Tools

Since programming at the toolkit level is quite difficult, there is a tremendous interest in higher level tools that will make the user interface software production process easier. These are discussed next.

8.1 Phases

Many higher-level tools have components that operate at different times. The **design time component** helps the user interface designer design the user interface. For example, this might be a graphical editor which can lay out the interface, or a compiler to process a user interface specification language. The next phase is when the end-user is using the program. Here, the **run-time component** of the tool is used. This usually includes a toolkit, but may also include additional software specifically for the tool. Since the run-time component is ‘‘managing’’ the user interface, the term **User Interface Management System** (UIMS) seems appropriate for tools with a significant run-time component.

There may also be an **after-run-time component** that helps with the evaluation and debugging of the user interface. Unfortunately, very few user interface tools have an after-run-time component. This is partially because tools that have tried, such as MIKE [Olsen 88], discovered that there are very few metrics that can be applied by computers. A new generation of tools are trying to evaluate how people will interact with interfaces by automatically creating cognitive models from high-level descriptions of the user interface. For example, the GLEAN system generates quantitative predictions of performance of a system from a GOMS model [Kieras 95].

8.2 Specification Styles

High-level user interface tools come in a large variety of forms. One important way that they can be classified is by how the designer specifies what the interface should be. As shown in Figure 8, some tools require the programmer to program in a special-purpose language, some provide an application framework to guide the programming, some automatically generate the interface from a high-level model or specification, and others allow the interface to be designed interactively. Each of these types is discussed below. Of course, some tools use different techniques for specifying different parts of the user

interface. These are classified by their predominant or most interesting feature.

Specification Format	Examples	Section
Language Based		8.2.1
State Transition Networks	[Newman 68], VAPS [VAPS 95]	8.2.1.1
Context-Free Grammars	YACC, LEX, Syngraph [Olsen 83]	8.2.1.2
Event Languages	Sassafras [Hill 86], EET [Frank 96], HyperTalk	8.2.1.3
Declarative Languages	Cousin [Hayes 85], Open Dialog [Schulert 85], Motif UIL	8.2.1.4
Constraint Languages	Thinglab [Borning 81], C32 [Myers 91b]	8.2.1.5
Screen Scrapers	Easel [Easel 93]	8.2.1.6
Database Interfaces	Oracle [Oracle 96]	8.2.1.7
Visual Programming	LabView [Labview 96] Prograph [Prograph 96]	8.2.1.8
Application Frameworks	MacApp [Wilson 90], Unidraw [Vlissides 90], Amulet [Myers 96a], Andrew [Palay 88], OLE, OpenDoc	8.2.2
Model-Based Generation	MIKE [Olsen 86], ITS [Wiecha 90], UIDE [Sukaviriya 93], Humanoid [Szekely 93], MASTERMIND [Neches 93]	8.2.3
Interactive Graphical Specification		8.2.4
Prototypers	Bricklin's Demo [Lifeboat 95], Director [MacroMedia 96]	8.2.4.1
Cards	Menulay [Buxton 83], HyperCard	8.2.4.2
Interface Builders	DialogEditor [Cardelli 88], NeXT Interface Builder [NeXT 91], Visual Basic, UIM/X [VisualEdge 90]	8.2.4.3
Data Visualization Tools	DataViews [DataViews 95]	8.2.4.4
Graphical Editors	Peridot [Myers 88b], DEMO [Fisher 92], Marquise [Myers 93]	8.2.4.5

Figure 8: Ways to specify the user interface, some tools that use that technique, and the section of this article that discusses the technique.

8.2.1 Language Based

With most of the older user interface tools, the designer specifies the user interface in a special-purpose language. This language can take many forms, including context-free grammars, state transition diagrams, declarative languages, event languages, etc. The language is usually used to specify the syntax of the user interface; i.e., the legal sequences of input and output actions. This is sometimes called the ‘‘dialog.’’ Green [Green 86] provides an extensive comparison of grammars, state transition diagrams, and event languages, and Olsen [Olsen 92] surveys various UIMS techniques.

8.2.1.1 State Transition Networks

Since many parts of user interfaces involve handling a sequence of input events, it is natural to think of using a state transition network to code the interface. A transition network consists of a set of states, with arcs out of each state labeled with the input tokens that will cause a transition to the state at the other end of the arc. In addition to input tokens, calls to application procedures and the output to display can also be put on the arcs in some systems. Newman implemented a simple tool using finite state machines in 1968 [Newman 68] which handled textual input. This was apparently the first user interface tool. Many of the assumptions and techniques used in modern systems were present in Newman’s: different languages for defining the user interface and the semantics (the semantic routines were coded in a normal programming language), a table-driven syntax analyzer, and device independence.

State diagram tools are most useful for creating user interfaces where the user interface has a large number of modes (each state is really a mode). For example, state diagrams are useful for describing the operation of low-level widgets (e.g., how a menu or scroll bar works), or the overall global flow of an application (e.g., this command will pop-up a dialog box, from which you can get to these two dialog boxes, and then to this other window, etc.). However, most highly-interactive systems attempt to be mostly ‘‘mode-free,’’ which means that at each point, the user has a wide variety of choices of what to do. This requires a large number of arcs out of each state, so state diagram tools have not been successful for these interfaces. In addition, state diagrams cannot handle interfaces where the user can operate on multiple objects at the same time. Another problem is that they can be very confusing for large interfaces, since they get to be a ‘‘maze of wires’’ and off-page (or off-screen) arcs can be hard to follow.

VAPS is a commercial system that uses the state transition model, and it eliminates the maze-of-wires problem by providing a spreadsheet-like table in which the states, events, and actions are specified [VAPS 95]. Transition networks have been thoroughly researched, but have not proven particularly successful or useful either as a research or commercial approach.

8.2.1.2 Context-Free Grammars

Many grammar-based systems are based on parser generators used in compiler development. For example, the designer might specify the user interface syntax using some form of BNF. Examples of grammar-based systems are Syngraph [Olsen 83] and parsers built with YACC and LEX in Unix.

Grammar-based tools, like state diagram tools, are not appropriate for specifying highly-interactive interfaces, since they are oriented to batch processing of strings with a complex syntactic structure. These systems are best for textual command languages, and have been mostly abandoned for user interfaces by researchers and commercial developers.

8.2.1.3 Event Languages

With event languages, the input tokens are considered to be ‘‘events’’ that are sent to individual event handlers. Each handler will have a condition clause that determines what types of events it will handle, and when it is active. The body of the handler can cause output events, change the internal state of the system (which might enable other event handlers), or call application routines.

Sassafras [Hill 86] is an event language where the user interface is programmed as a set of small event handlers. The Elements-Events and Transitions (EET) language provides elaborate control over when the various event handlers are fired [Frank 96]. As described below in section 8.2.4.2, the HyperTalk language that is part of HyperCard for the Apple Macintosh, can be considered an event language. Microsoft’s Visual Basic also contains event-language features, since code is generally written as responses to events on objects.

The advantages of event languages are that they can handle multiple input devices active at the same time, and it is straightforward to support non-modal interfaces, where the user can operate on any widget or object. The main disadvantage is that it can be very difficult to create correct code, since the flow of control is not localized and small changes in one part can affect many different pieces of the program. It is also typically difficult for the designer to understand the code once it reaches a non-trivial size. However, the success of HyperTalk, Visual Basic and similar tools shows that this approach is appropriate for small to medium-size programs.

8.2.1.4 Declarative Languages

Another approach is to try to define a language that is declarative (stating what should happen) rather than procedural (how to make it happen). Cousin [Hayes 85] and HP/Apollo’s Open-Dialogue [Schulert 85] both allow the designer to specify user interfaces in this manner. The user interfaces supported are basically forms, where fields can be text which is typed by the user, or options selected using menus or buttons. There are also graphic output areas that the application can use in whatever manner desired. The application program is connected to the user interface through ‘‘variables’’ which can be set and accessed

by both. As researchers have extended this idea to support more sophisticated interactions, the specification has grown into full application ‘‘models,’’ and newer systems are described below in section 8.2.3.

The layout description languages that come with many toolkits are also a type of declarative language. For example, Motif’s User Interface Language (UIL) allows the layout of widgets to be defined. Since the UIL is interpreted when an application starts, users can (in theory) edit the UIL code to customize the interface. UIL is not a complete language, however, in the sense that the designer must still write C code for many parts of the interface, including any areas containing dynamic graphics and any widgets that change.

The advantage of using declarative languages is that the user interface designer does not have to worry about the time sequence of events, and can concentrate on the information that needs to be passed back and forth. The disadvantage is that only certain types of interfaces can be provided this way, and the rest must be programmed by hand in the ‘‘graphic areas’’ provided to application programs. The kinds of interactions available are preprogrammed and fixed. In particular, these systems provide no support for such things as dragging graphical objects, rubber-band lines, drawing new graphical objects, or even dynamically changing the items in a menu based on the application mode or context. However, these languages are now proving successful as intermediate languages describing the layout of widgets (such as UIL) that are generated by interactive tools.

8.2.1.5 Constraint Languages

A number of user interface tools allow the programmer to use constraints to define the user interface [Borning 86a]. Early constraint systems include Sketchpad [Sutherland 63] which pioneered the use of graphical constraints in a drawing editor, and Thinglab [Borning 81] which used constraints for graphical simulation. Subsequently, Thinglab was extended to aid in the generation of user interfaces [Borning 86a].

The discussion of toolkits above mentioned the use of constraints as part of the intrinsics of a toolkit. A number of research toolkits now supply constraints as an integral part of the object system (e.g., Garnet and Amulet). In addition, some systems have provided higher-level interfaces to constraints. Graphical Thinglab [Borning 86b] allows the designer to create constraints by wiring icons together, and NoPump [Wilde 90] and C32 [Myers 91b] allow constraints to be defined using a spreadsheet-like interface.

The advantage of constraints is that they are a natural way to express many kinds of relationships that arise frequently in user interfaces. For example, that lines should stay attached to boxes, that labels should stay centered within boxes, etc. A disadvantage with constraints is that they require a sophisticated run-time system to solve them efficiently. However, a growing number of research systems

are using constraints, and it appears that modern constraint solvers and debugging techniques may solve these problems, so constraints have a great potential to simplify the programming task. As yet, there are no commercial user interface tools using general-purpose constraint solvers.

8.2.1.6 Screen Scrapers

Some commercial tools are specialized to be ‘‘front-enders’’ or ‘‘screen scrapers’’ which provide a graphical user interface to old programs without changing the existing application code. They do this by providing an in-memory buffer that pretends to be the screen of an old character terminal such as might be attached to an IBM mainframe. When the mainframe application outputs to the buffer, a program the designer writes in a special programming language converts this into an update of a graphical widget. Similarly, when the user operates a widget, the script converts this into the appropriate edits of the character buffer. A leading program of this type is Easel, [Easel 93], which also contains an interface builder for laying out the widgets.

8.2.1.7 Database Interfaces

A very important class of commercial tools support form-based or GUI-based access to databases. Major database vendors such as Oracle [Oracle 96] provide tools which allow designers to define the user interface for accessing and setting data. Often these tools include interactive form editors (which are essentially interface builders) and special database languages. Fourth generation languages (4GLs), that support defining the interactive forms for accessing and entering data, also fall into this category.

8.2.1.8 Visual Programming

‘‘Visual programs’’ use graphics and two (or more) dimensional layout as part of the program specification [Myers 90d]. Many different approaches to using visual programming to specify user interfaces have been investigated. Most systems that support state transition networks (section 8.2.1.1) use a visual representation. Another popular technique is to use dataflow languages. In these, icons represent processing steps, and the data flow along the connecting wires. The user interface is usually constructed directly by laying out pre-built widgets, in the style of interface builders (section 8.2.4.3). Examples of visual programming systems for creating user interfaces include Labview [Labview 96] which is specialized for controlling laboratory instruments, and Prograph [Prograph 96]. Using a visual language seems to make it easier for novice programmers, but large programs still suffer from the familiar ‘‘maze of wires’’ problem. Other papers (e.g., [Myers 90d]) have analyzed the strengths and weaknesses of visual programming in detail.

Another popular language is Visual Basic from Microsoft. Although this is more of a structure editor for Basic combined with an Interface Builder, and therefore does not really count as a visual language, it does make the construction of user interface software easier. Microsoft is pushing Visual Basic as the

extension language that people will use to customize and connect all future Windows-based applications.

8.2.1.9 Summary of Language Approaches

In summary, many different types of languages have been designed for specifying user interfaces. One problem with all of these is that they can only be used by professional programmers. Some programmers have objected to the requirement for learning a new language for programming just the user interface portion [Olsen 87]. This has been confirmed by market research [XBusiness 94, p.29]. Furthermore, it seems more natural to define the graphical part of a user interface using a graphical editor (see section 8.2.4). However, it is clear that for the foreseeable future, much of the user interface will still need to be created by writing programs, so it is appropriate to continue investigations into the best language to use for this. Indeed, an entire book is devoted to investigating the languages for programming user interfaces [Myers 92b].

8.2.2 Application Frameworks

After the Macintosh Toolbox had been available for a little while, Apple discovered that programmers had a difficult time figuring out how to call the various toolkit functions, and how to ensure that the resulting interface met the Apple guidelines. They therefore created a software system that provides an overall **application framework** to guide programmers. This was called MacApp [Wilson 90] and used the object-oriented language Object Pascal. Classes are provided for the important parts of an application, such as the main windows, the commands, etc., and the programmer specializes these classes to provide the application-specific details, such as what is actually drawn in the windows and which commands are provided. MacApp was very successful at simplifying the writing of Macintosh applications. Today, there are multiple frameworks to help build applications for most major platforms, including the Microsoft Foundation Classes for Windows, and the CodeWarrior PowerPlant [Metrowerks 96] for the Macintosh.

Unidraw [Vlissides 90] is a research framework, but it is more specialized for graphical editors. This means that it can provide even more support. Unidraw uses the C++ object-oriented language and is part of the InterViews system [Linton 89]. Unidraw has been used to create various drawing and CAD programs, and a user interface editor [Vlissides 91]. The Amulet framework [Myers 96a] is also aimed at graphical applications, but due to its graphical data model, many of the built-in routines can be used without change (the programmer does not usually need to write methods for subclasses). Even more specialized are various graph programs, such as Edge [Newbery 88] and TGE [Karrer 90]. These provide a framework in which the designer can create programs that display their data as trees or graphs. The programmer typically specializes the node and arc classes, and specifies some of the commands, but the framework handles layout and the overall control.

An emerging popular approach aims to replace today's large, monolythic applications with smaller components that attach together. For example, you might buy a separate text editor, ruler, paragraph formatter, spell checker, and drawing program, and have them all work together seamlessly. This approach was invented by the Andrew environment [Palay 88] which provides an object-oriented document model that supports the embedding of different kinds of data inside other documents. These ''insets'' are unlike data that is cut and pasted in systems like the Macintosh because they bring along the programs that edit them, and therefore can always be edited in place. Furthermore, the container document does not need to know how to display or print the inset data since the original program that created it is always available. The designer creating a new inset writes subclasses that adheres to a standard protocol so the system knows how to pass input events to the appropriate editor. The next generation of operating systems will use this approach extensively: it is the foundation for Microsoft's OLE and Apple's OpenDoc.

All of these frameworks require the designer to write code, typically by creating application-specific sub-classes of the standard classes provided as part of the framework.

8.2.3 Model-Based Automatic Generation

A problem with all of the language-based tools is that the designer must specify a great deal about the placement, format, and design of the user interfaces. To solve this problem, some tools use **automatic generation** so that the tool makes many of these choices from a much higher-level specification. Many of these tools, such as Mickey [Olsen 89], Jade [Vander Zanden 90], and DON [Kim 93] have concentrated on creating menus and dialog boxes. Jade allows the designer to use a graphical editor to edit the generated interface if it is not good enough. DON has the most sophisticated layout mechanisms and takes into account the desired window size, balance, columness, symmetry, grouping, etc. Creating dialog boxes automatically has been very thoroughly researched, but there still are no commercial tools that do this.

UIDE (the User-Interface Design Environment) [Sukaviriya 93] requires that the semantics of the application be defined in a special-purpose language, and therefore might be included with the language-based tools (section 8.2.1). It is placed here instead because the language is used to describe the functions that the application supports and not the desired interface. UIDE is classified as a ''model-based'' approach because the specification serves as a high-level, sophisticated model of the application semantics. In UIDE, the description includes pre- and post-conditions of the operations, and the system uses these to reason about the operations, and to automatically generate an interface. One interesting feature of UIDE is that the pre- and post-conditions are used to automatically generate help [Sukaviriya 90].

Another model-based system is HUMANOID [Szekely 93] which supports the modeling of the presentation, behavior and dialogue of an interface. The HUMANOID modeling language includes abstraction, composition, recursion, iteration and conditional constructs to support sophisticated interfaces. The HUMANOID system, which is built on top of the Garnet toolkit [Myers 90c], provides a number of interactive modeling tools to help the designer specify the model. The developers of HUMANOID and UIDE are collaborating on a new combined model called MASTERMIND that integrates their approaches [Neches 93].

The ITS [Wiecha 90] system also uses rules to generate an interface. ITS was used to create the visitor information system for the EXPO 1992 worlds fair in Seville, Spain. Unlike the other rule-based systems, the designer using ITS is expected to write many of the rules, rather than just writing a specification that the rules work on. In particular, the design philosophy of ITS is that all design decisions should be codified as rules so that they can be used by subsequent designers, which will hopefully mean that interface designs will become easier and better as more rules are entered. As a result, the designer should never use graphical editing to improve the design, since then the system cannot capture the reason that the generated design was not sufficient.

While the idea of having the user interface generated automatically is appealing, this approach is still at the research level, because the user interfaces that are generated are generally not good enough. A further problem is that the specification languages can be quite hard to learn and use. Extensive current research is addressing the problems of expanding the range of what can be created automatically (to go beyond dialog boxes) and to make the model-based approach easier to use.

8.2.4 Direct Graphical Specification

The tools described next all allow the user interface to be defined, at least partially, by placing objects on the screen using a pointing device. This is motivated by the observation that the visual presentation of the user interface is of primary importance in graphical user interfaces, and a graphical tool seems to be the most appropriate way to specify the graphical appearance. Another advantage of this technique is that it is usually much easier for the designer to use. Many of these systems can be used by non-programmers. Therefore, psychologists, graphic designers and user interface specialists can more easily be involved in the user interface design process when these tools are used.

These tools can be distinguished from those that use ‘‘visual programming’’ (section 8.2.1.8) since with direct graphical specification, the actual user interface (or a part of it) is drawn, rather than being generated indirectly from a visual program. Thus, direct graphical specification tools have been called **direct manipulation programming** since the user is directly manipulating the user interface widgets and other elements.

The tools that support graphical specification can be classified into four categories: prototyping tools, those that support a sequence of cards, interface builders, and editors for application-specific graphics.

8.2.4.1 Prototyping Tools

The goal of **prototyping tools** is to allow the designer to quickly mock up some examples of what the screens in the program will look like. Often, these tools cannot be used to create the real user interface of the program; they just show how some aspects will look. This is the chief factor that distinguishes them from other high-level tools. Many parts of the interface may not be operable, and some of the things that look like widgets may just be static pictures. In most prototypers, no real toolkit widgets are used, which means that the designer has to draw simulations that look like the widgets that will appear in the interface. The normal use is that the designer would spend a few days or weeks trying out different designs with the tool, and then completely reimplement the final design in a separate system. Most prototyping tools can be used without programming, so they can, for example, be used by graphic designers.

Note that this use of the term ‘‘prototyping’’ is different from the general phrase ‘‘rapid prototyping,’’ which has become a marketing buzz-word. Advertisements for just about all user interface tools claim that they support ‘‘rapid prototyping,’’ by which they mean that the tool helps create the user interface software more quickly. The term ‘‘prototyping’’ is being used in this paper in a much more specific manner.

Probably the first prototyping tool was Dan Bricklin’s Demo program [Lifeboat 95]. This is a program for an IBM PC that allows the designer to create sample screens composed of characters and ‘‘character graphics’’ (where the fixed-size character cells can contain a graphic like a horizontal, vertical or diagonal line). The designer can easily create the various screens for the application. It is also relatively easy to specify the actions (mouse or keyboard) that cause transitions from one screen to another. However, it is difficult to define other behaviors. In general, there may be some support for type-in fields and menus in prototyping tools, but there is little ability to process or test the results.

For graphical user interfaces, designers often use tools like Macromedia’s Director [MacroMedia 96] which is actually an animation tool. The designer can draw example screens, and then specify that when the mouse is pressed in a particular place, an animation should start or a different screen should be displayed. Components of the picture can be reused in different screens, but again the ability to show behavior is limited. HyperCard for the Macintosh is also often used as a prototyping tool.

The primary disadvantage of these prototyping tools is that they cannot create the actual code for the user interface. Therefore, the interfaces must be re-coded after prototyping. There is also the risk that the programmers who implement the real user interface will ignore the prototype. Therefore, a new research tool called SILK is trying to provide a quick sketching interface and then convert the sketches into actual

widgets [Landay 95].

8.2.4.2 Cards

Many graphical programs are limited to user interfaces that can be presented as a sequence of mostly static pages, sometimes called ‘‘frames,’’ ‘‘cards,’’ or ‘‘forms.’’ Each page contains a set of widgets, some of which cause transfer to other pages. There is usually a fixed set of widgets to choose from, which have been coded by hand.

An early example of this is Menulay [Buxton 83], which allows the designer to place text, graphical potentiometers, iconic pictures, and light buttons on the screen and see exactly what the end user will see when the application is run. The designer does not need to be a programmer to use Menulay.

Probably, the most famous example of a card-based system is HyperCard from Apple. There are now many similar programs, such as GUIDE [InfoAccess 91], and Tool Book [Asymetrix 95]. In all of these, the designer can easily create cards containing text fields, buttons, etc., along with various graphic decorations. The buttons can transfer to other cards. These programs provide a scripting language to provide more flexibility for buttons. HyperCard’s scripting language is called HyperTalk, and as mentioned above in section 8.2.1.3, is really an event language, since the programmer writes short pieces of code that are executed when input events occur.

8.2.4.3 Interface Builders

An **interface builder** allows the designer to create dialog boxes, menus and windows that are to be part of a larger user interface. These are also called **Interface Development Tools (IDTs)**. Interface builders allow the designer to select from a pre-defined library of widgets, and place them on the screen using a mouse. Other properties of the widgets can be set using property sheets. Usually, there is also some support for sequencing, such as bringing up sub-dialogs when a particular button is hit. The Steamer project at BBN demonstrated many of the ideas later incorporated into interface builders and was probably the first object-oriented graphics system [Stevens 83]. Other examples of research interface builders are DialogEditor [Cardelli 88], vu [Singh 88] and Gilt [Myers 91a]. There are literally hundreds of commercial interface builders. Just two examples are the NeXT Interface Builder [NeXT 91], and UIM/X for X [VisualEdge 90]. Visual Basic is essentially an Interface Builder coupled with an editor for an interpreted language. Many of the tools discussed above, such as the virtual toolkits, visual languages, and application frameworks, also contain interface builders.

Interface builders use the actual widgets from a toolkit, so they can be used to build parts of real applications. Most will generate C code templates that can be compiled along with the application code. Others generate a description of the interface in a language that can be read at run-time. For example, UIM/X generates a UIL description. It is usually important that the programmer not edit the output of the

tools (such as the generated C code) or else the tool can no longer be used for later modifications.

Although interface builders make laying out the dialog boxes and menus easier, this is only part of the user interface design problem. These tools provide little guidance towards creating good user interfaces, since they give designers significant freedom. Another problem is that for any kind of program that has a graphics area (such as drawing programs, CAD, visual language editors, etc.), interface builders do not help with the contents of the graphics pane. Also, they cannot handle widgets that change dynamically. For example if the contents of a menu or the layout of a dialog box changes based on program state, this must be programmed by writing code. To help with this part of the problem, some interface builders, like UIM/X [VisualEdge 90], provide a C code interpreter.

8.2.4.4 Data Visualization Tools

An important commercial category of tools is that of dynamic data visualization systems. These tools, which tend to be quite expensive, emphasize the display of dynamically changing data on a computer, and are used as front ends for simulations, process control, system monitoring, network management, and data analysis. The interface to the designer is usually quite similar to an interface builder, with a palette of gauges, graphers, knobs and switches that can be placed interactively. However, these controls usually are not from a toolkit and are supplied by the tool. Example tools in this category include DataViews [DataViews 95], and SL-GMS [SL-GMS 93].

8.2.4.5 Editors for Application-Specific Graphics

When an application has custom graphics, it would be useful if the designer could draw pictures of what the graphics should look like rather than having to write code for this. The problem is that the graphic objects usually need to change at run time, based on the actual data and the end user's actions. Therefore, the designer can only draw an *example* of the desired display, which will be modified at run-time, and so these tools are called ''demonstrational programming'' [Myers 92c]. This distinguishes these programs from the graphical tools of the previous three sections, where the full picture can be specified at design time. As a result of the *generalization* task of converting the example objects into parameterized prototypes that can change at run-time, most of these systems are still in the research phase.

Peridot [Myers 88b] allows new, custom widgets to be created. The primitives that the designer manipulates with the mouse are rectangles, circles, text, and lines. The system generalizes from the designer's actions to create parameterized, object-oriented procedures like those that might be found in toolkits. Experiments showed that Peridot can be used by non-programmers. Lapidary [Myers 89] extends the ideas of Peridot to allow general application-specific objects to be drawn. For example, the designer can draw the nodes and arcs for a graph program. The DEMO system [Fisher 92] allows

some dynamic, run-time properties of the objects to be demonstrated, such as how objects are created. The Marquise tool [Myers 93] allows the designer to demonstrate *when* various behaviors should happen, and supports palettes which control the behaviors. With Pavlov [Wolber 96], the user can demonstrate how widgets should control a car's movement in a driving game. Research continues on making these ideas practical.

8.3 Specialized Tools

For some application domains, there are customized tools that provide significant high-level support. These tend to be quite expensive, however (i.e., US$20,000 to US$50,000). For example, in the aeronautics and real-time control areas, there are a number of high-level tools, such as AutoCode [AutoCode 91] and InterMAPhics [Gallium 96].

9. Technology Transfer

User interface tools are an area where research has had a tremendous impact on the current practice of software development. Of course, window managers and the resulting ''GUI style'' comes from the seminal research at the Stanford Research Institute, Xerox Palo Alto Research Center, and MIT in the 1970s. Interface builders and ''card'' programs like HyperCard were invented in research labs at BBN, the University of Toronto, Xerox PARC, and others. Now, interface builders are at least a US$100 million per year business and are widely used for commercial software development. Event languages, as widely used in HyperTalk and elsewhere, were first investigated in research labs. The next generation of environments, like OLE and OpenDoc, are based on the component architecture which was developed in the Andrew environment from CMU. Thus, whereas some early UIMS approaches like transition networks and grammars may not have been successful, overall, the user interface tool research has changed the way that software is developed.

10. Research Issues

Although there are many user interface tools, there are plenty of areas in which further research is needed. A report prepared for an NSF study discusses future research ideas for user interface tools at length [Olsen 93]. Here, a few of the important ones are summarized.

10.1 New Programming Languages

The built-in input/output primitives in today's programming languages, such as `printf/scanf` or `cout/cin`, support a textual question-and-answer style of user interface which is modal and well-known to be poor. Most of today's tools use libraries and interactive programs which are separate from

programming languages. However, many of the techniques, such as object-oriented programming, multiple-processing, and constraints, are best provided as *part* of the programming language. Even new languages, such as Java, make much of the user interface harder to program by leaving it in separate libraries. Furthermore, an integrated environment, where the graphical parts of an application can be specified graphically and the rest textually, would make the generation of applications much easier. How programming languages can be improved to better support user interface software is the topic of a book [Myers 92b].

10.2 Increased Depth

Many researchers are trying to create tools that will cover more of the user interface, such as application-specific graphics and behaviors. The challenge here is to allow flexibility to application developers while still providing a high level of support. Tools should also be able to support Help, Undo, and Aborting of operations.

Today's user interface tools mostly help with the *generation* of the code of the interface, and assume that the fundamental user interface *design* is complete. What are also needed are tools to help with the generation, specification, and analysis of the design of the interface [Landay 95]. For example, an important first step in user interface design is task analysis, where the designer identifies the particular tasks that the end user will need to perform. Research should be directed at creating tools to support these methods and techniques. These might eventually be integrated with the code generation tools, so that the information generated during early design can be fed into automatic generation tools, possibly to produce an interface directly from the early analyses. The information might also be used to automatically generate documentation and run-time help.

Another approach is to allow the designer to specify the design in an appropriate notation, and then provide tools to convert that notation into interfaces. For example, the UAN [Hartson 90] is a notation for expressing the end user's actions and the system's responses.

Finally, much work is needed in ways for tools to help evaluate interface designs. Initial attempts, such as in MIKE [Olsen 88], have highlighted the need for better models and metrics against which to evaluate the user interfaces. Research in this area by cognitive psychologists and other user interface researchers (e.g., [Kieras 95]) is continuing.

10.3 Increased Breadth

We can expect the user interfaces of tomorrow to be different from the conventional window-and-mouse interfaces of today, and tools will have to change to support the new styles. For example, most tools today only deal with two-dimensional objects, but there is already a demand to provide 3-D visualizations and animations. New input devices and techniques will probably replace the conventional mouse and menu styles. For example, gesture and handwriting recognition are appearing in mass-market commercial products, such as notepad computers and ‘‘personal digital assistants’’ like Apple’s Newton (gesture recognition has actually been used since the 1970s in commercial CAD tools). ‘‘Virtual reality’’ systems, where the computer creates an artificial world and allows the user to explore it, cannot be handled by any of today’s tools. In these ‘‘non-WIMP’’ applications (**WIMP** stands for Windows, Icons, Menus and Pointing devices), designers will also need better control over the timing of the interface, to support animations and various new media like video [Nielsen 93]. Although a few tools are directed at multiple-user applications, there are no direct graphical specification tools, and the current tools are limited in the styles of applications they support.

10.4 End User Programming and Customization

One of the most successful computer programs of all time is the spreadsheet. The primary reason for its success is that end users can program (by writing formulas and macros). However, **end user programming** is rare in other applications, and where it exists, usually requires learning conventional programming. For example, AutoCAD provides Lisp for customization. More effective mechanisms for users to customize existing applications and create new ones are needed [Myers 92b]. However, these should not be built into individual applications as is done today, since this means that the user must learn a different programming technique for each application. Instead, the facilities should be provided at the system level, and therefore should be part of the underlying toolkit. Naturally, since this is aimed at end users, it will not be like programming in C, but rather at some higher level.

The X Business Group predicts that there will be an increased use of tools by end users, rather than professional software developers, which will present enormous opportunities and challenges to tool creators [XBusiness 94].

10.5 Application and UI Separation

One of the fundamental goals of user interface tools is to allow better modularization and separation of user interface code from application code. However, a survey reported that conventional toolkits actually make this separation more difficult, due to the large number of call-back procedures required [Myers 92a]. Therefore, further research is needed into ways to better modularize the code, and how tools can support this.

10.6 Tools for the Tools

It is very difficult to create the tools described in this paper. Each one takes an enormous effort. Therefore, work is needed in ways to make the tools themselves easier to create. For example, the Garnet toolkit explored mechanisms specifically designed to make high-level graphical tools easier to create [Myers 92d]. The Unidraw framework has also proven useful for creating interface builders [Vlissides 91]. However, more work is needed.

11. Conclusions

The area of user interface tools is expanding rapidly. Ten years ago, you would have been hard-pressed to find any successful commercial higher-level tools, but now there are hundreds of different tools, and tools are turning into a billion dollar a year business. Chances are that today, whatever your project is, there is a tool that will help. Tools that are coming out of research labs are covering increasingly more of the user interface task, are more effective at helping the designer, and are creating better user interfaces. As more companies and researchers are attracted to this area, we can expect the pace of innovation to continue to accelerate. There will be many exciting and useful new tools available in the future.

12. Defining Terms

Application or **application semantics**: the part of the software that is not the user interface.

Application Framework: a software architecture, often object-oriented, that guides the programmer so that implementing user interface software is easier.

Browser: An interactive program for viewing pages on the World-Wide-Web, that might replace Windowing Systems.

Interface Builder: interactive tool that lays out widgets to create dialog boxes, menus and windows that are to be part of a larger user interface. These are also called **Interface Development Tools**.

Intrinsics: the layer of a toolkit on which different widgets are implemented.

Prototyping tool: allow the designer to quickly mock up some examples of what the screens in the program will look like. Often, these tools cannot be used to create the real user interface of the program; they just show how some aspects will look.

Toolkit: a library of widgets that can be called by application programs.

User Interface (UI): the part of the software that handles the output to the display and the input from

the person using the program.

User Interface Development Environments (UIDE): general term for comprehensive user interface tools.

User Interface Management System (UIMS): an older term, not much used now. Sometimes used to cover all ‘‘user interface tools,’’ but usually limited to tools that handle the sequencing of operations (what happens after each event from the user).

User Interface Tool: Any software that helps create user interfaces.

Virtual toolkits: also called **cross-platform development systems**, are programming interfaces to multiple toolkits that allow code to be easily ported to Macintosh, Microsoft Windows and Unix environments.

Widget: a way of using a physical input device to input a certain type of value. Typically, widgets in toolkits include menus, buttons, scroll bars, text type-in fields, etc.

Window manager: the user interface of the windowing system. Also used to mean the entire windowing system.

Windowing system: software that separates different processes into different rectangular regions (**windows**) on the screen.

13. References

[Apple 85] Apple Computer, Inc. *Inside Macintosh.* Addison-Wesley, 1985.

[Asymetrix 95] Asymetrix Corporation. Tool Book. 100 100th Ave., NE, Bellevue, Washington, 98004, (206) 637-1500. 1995.

[AutoCode 91] Integrated Systems. AutoCode. 3260 Jay Street, Santa Clara, CA 94054. (408) 980-1500. 1991.

[Bharat 94] Krishna Bharat and Marc H. Brown. Building Distributed, Multi-User Applications by Direct Manipulation. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 71-81. Proceedings UIST’94, Marina del Rey, CA, November, 1994.

[Bly 86] Sara A. Bly and Jarrett K. Rosenberg. A Comparison of Tiled and Overlapping Windows. In *Human Factors in Computing Systems*, pages 101-106. Proceedings SIGCHI’86, Boston, Mass, April, 1986.

[Booz 92] Booz Allen & Hamilton Inc. NeXTStep vs. Other Development Environments; Comparative Study. 1992. Report available from NeXT Computer, Inc.

[Borning 81] Alan Borning. The Programming Language Aspects of Thinglab; a Constraint-Oriented Simulation Laboratory. *ACM Transactions on Programming Languages and Systems* 3(4):353-387, October, 1981.

[Borning 86a] Alan Borning and Robert Duisberg. Constraint-Based Tools for Building User Interfaces. *ACM Transactions on Graphics* 5(4):345-374, October, 1986.

[Borning 86b] Alan Borning. Defining Constraints Graphically. In *Human Factors in Computing Systems*, pages 137-143. Proceedings SIGCHI'86, Boston, MA, April, 1986.

[Buxton 83] W. Buxton, M.R. Lamb, D. Sherman, and K.C. Smith. Towards a Comprehensive User Interface Management System. In *Computer Graphics*, 17(3), pages 35-42. Proceedings SIGGRAPH'83, Detroit, Mich, July, 1983.

[Cardelli 85] Luca Cardelli and Rob Pike. Squeak: A Language for Communicating with Mice. In *Computer Graphics*, pages 199-204. Proceedings SIGGRAPH'85, San Francisco, CA, July, 1985.

[Cardelli 88] Luca Cardelli. Building User Interfaces by Direct Manipulation. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 152-166. Proceedings UIST'88, Banff, Alberta, Canada, October, 1988.

[DataViews 95] DataViews Corporation. DataViews. 47 Pleasant St., Northampton, MA, 01006, (413) 586-4144. 1995.

[DePalma 93] Donald A. DePalma and Stuart D. Woodring. Client/Server Power Tools Futures. *The Software Strategy Report* 4(1):2-13, April, 1993. Forrester Research, One Brattle Square, Camb, MA 02138.

[Easel 93] Easel. Workbench. 25 Corporate Drive, Burlington, MA 01803. (617) 221-2100. 1993.

[Fisher 92] Gene L. Fisher, Dale E. Busse, and David A. Wolber. Adding Rule-Based Reasoning to a Demonstrational Interface Builder. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 89-97. Proceedings UIST'92, Monterey, CA, November, 1992.

[Frank 96] Martin R. Frank. *Model-Based User Interface Design by Demonstration and by Interview*. PhD thesis, College of Computing, Georgia Institute of Technology, 1996. In progress.

[Gallium 96] Gallium Software Inc. InterMAPhics. 4000-303 Moodie Drive, Nepean, Ontario, CANADA K2H 9R4; (613) 721-0902, http://www.gallium.com/ProductInfo/InterMAPhics. 1996.

[Gaskins 92] Tom Gaskins. *PEXlib Programming Manual.* O'Reilly and Associates, Inc., 103 Morris Street, Suite A, Sebastopol CA, 1992.

[Gleicher 93] Michael Gleicher. A Graphics Toolkit Based on Differential Constraints. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 109-120. Proceedings UIST'93, Atlanta, GA, November, 1993.

[Green 86] Mark Green. A Survey of Three Dialog Models. *ACM Transactions on Graphics* 5(3):244-275, July, 1986.

[Hartson 90] H. Rex Hartson, Antonio C. Siochi, and Deborah Hix. The UAN: A User-Oriented Representation for Direct Manipulation Interface Designs. *ACM Transactions on Information Systems* 8(3):181-203, July, 1990.

[Hayes 85] Philip J. Hayes, Pedro A. Szekely, and Richard A. Lerner. Design Alternatives for User Interface Management Systems Based on Experience with COUSIN. In *Human Factors in Computing Systems*, pages 169-175. Proceedings SIGCHI'85, San Francisco, CA, April, 1985.

[Hill 86] Ralph D. Hill. Supporting Concurrency, Communication and Synchronization in Human-Computer Interaction — The Sassafras UIMS. *ACM Transactions on Graphics* 5(3):179-210, July, 1986.

[Hill 93] Ralph D. Hill, Tom Brinck, John F. Patterson, Steven L. Rohall, and Wayne T. Wilner. The Rendezvous Language and Architecture. *Communications of the ACM* 36(1):62-67, January, 1993.

[Hudson 93] Scott E. Hudson and John T. Stasko. Animation Support in a User Interface Toolkit: Flexible, Robust, and Reusable Abstractions. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 57-67. Proceedings UIST'93, Atlanta, GA, November, 1993.

[InfoAccess 91] InfoAccess Inc. Guide 2. 2800 156th Avenue SE, Bellevue, WA 98007. (206) 747-3203. 1991.

[Karrer 90] Anthony Karrer and Walt Scacchi. Requirements for an Extensible Object-Oriented Tree/Graph Editor. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 84-91. Proceedings UIST'90, Snowbird, Utah, October, 1990.

[Kieras 95] David E. Kieras, Scott D. Wood, Kasen Abotel, Anthony Hornof. GLEAN: A Computer-Based Tool for Rapid GOMS Model Usability Evaluation of User Interface Designs. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 91-100. Proceedings UIST'95, Pittsburgh, PA, November, 1995.

[Kim 93] Won Chul Kim and James D. Foley. Providing High-level Control and Expert Assistance in the User Interface Presentation Design. In *Human Factors in Computing Systems*, pages 430-437. Proceedings INTERCHI'93, Amsterdam, The Netherlands, April, 1993.

[Labview 96] National Instruments. LabVIEW. 6504 Bridge Point Parkway, Austin, TX, 78730-5039, (512)794-1000. 1996.

[Landay 95] James Landay and Brad A. Myers. Interactive Sketching for the Early Stages of User Interface Design. In *Human Factors in Computing Systems*, pages 43-50. Proceedings SIGCHI'95, Denver, CO, May, 1995.

[Lifeboat 95] Lifeboat Publishing. Dan Bricklin's Demo, 1163 Shrewsbury Ave, Shrewsbury, NJ 07702, 800-336-1166, 908-389-9227. 1995.

[Linton 89] Mark A. Linton, John M. Vlissides and Paul R. Calder. Composing user interfaces with InterViews. *IEEE Computer* 22(2):8-22, February, 1989.

[MacroMedia 96] MacroMedia. Director. 600 Townsend Street, San Francisco, CA 94103, macropr@macromedia.com, http://www.macromedia.com/. 1996.

[McCormack 88] Joel McCormack and Paul Asente. An Overview of the X Toolkit. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 46-55. Proceedings UIST'88, Banff, Alberta, Canada, October, 1988.

[Metrowerks 96] Metrowerks, Inc. PowerPlant for CodeWarrior. 2201 Donley Drive, Suite 310, Austin, TX 78758, (800)377-5416, http://www.metrowerks.com/. 1996.

[Myers 84] Brad A. Myers. The User Interface for Sapphire. *IEEE Computer Graphics and Applications* 4(12):13-23, December, 1984.

[Myers 86] Brad A. Myers. A Complete and Efficient Implementation of Covered Windows. *IEEE Computer* 19(9):57-67, September, 1986.

[Myers 88a] Brad A. Myers. A Taxonomy of User Interfaces for Window Managers. *IEEE Computer Graphics and Applications* 8(5):65-84, Sept, 1988.

[Myers 88b] Brad A. Myers. *Creating User Interfaces by Demonstration.* Academic Press, Boston, 1988.

[Myers 89] Brad A. Myers, Brad Vander Zanden, and Roger B. Dannenberg. Creating Graphical Interactive Application Objects by Demonstration. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 95-104. Proceedings UIST'89, Williamsburg, VA, November, 1989.

[Myers 90a] Brad A. Myers. A New Model for Handling Input. *ACM Transactions on Information Systems* 8(3):289-320, July, 1990.

[Myers 90b] Brad A. Myers. All the Widgets. *SIGGRAPH Video Review* 57, 1990.

[Myers 90c] Brad A. Myers, Dario A. Giuse, Roger B. Dannenberg, Brad Vander Zanden, David S. Kosbie, Edward Pervin, Andrew Mickish, and Philippe Marchal. Garnet: Comprehensive Support for Graphical, Highly-Interactive User Interfaces. *IEEE Computer* 23(11):71-85, November, 1990.

[Myers 90d] Brad A. Myers. Taxonomies of Visual Programming and Program Visualization. *Journal of Visual Languages and Computing* 1(1):97-123, March, 1990.

[Myers 91a] Brad A. Myers. Separating Application Code from Toolkits: Eliminating the Spaghetti of Call-Backs. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 211-220. Proceedings UIST'91, Hilton Head, SC, November, 1991.

[Myers 91b] Brad A. Myers. Graphical Techniques in a Spreadsheet for Specifying User Interfaces. In *Human Factors in Computing Systems*, pages 243-249. Proceedings SIGCHI'91, New Orleans, LA, April, 1991.

[Myers 92a] Brad A. Myers and Mary Beth Rosson. Survey on User Interface Programming. In *Human Factors in Computing Systems*, pages 195-202. Proceedings SIGCHI'92, Monterey, CA, May, 1992.

[Myers 92b] Brad A. Myers (editor). *Languages for Developing User Interfaces.* Jones and Bartlett, Boston, MA, 1992.

[Myers 92c] Brad A. Myers. Demonstrational Interfaces: A Step Beyond Direct Manipulation. *IEEE Computer* 25(8):61-73, August, 1992.

[Myers 92d] Brad A. Myers and Brad Vander Zanden. Environment for Rapid Creation of Interactive Design Tools. *The Visual Computer; International Journal of Computer Graphics* 8(2):94-116, February, 1992.

[Myers 93] Brad A. Myers, Richard G. McDaniel, and David S. Kosbie. Marquise: Creating Complete User Interfaces by Demonstration. In *Human Factors in Computing Systems*, pages 293-300. Proceedings INTERCHI'93, Amsterdam, The Netherlands, April, 1993.

[Myers 94] Brad A. Myers. Challenges of HCI Design and Implementation. *ACM Interactions* 1(1):73-83, January, 1994.

[Myers 96a] Brad A. Myers, Alan Ferrency, Rich McDaniel, Robert C. Miller, Patrick Doane, Andy Mickish, Alex Klimovitski. *The Amulet V2.0 Reference Manual*. Technical Report CMU-CS-95-166-R1, Carnegie Mellon University Computer Science Department, May, 1996. also Human Computer Interaction Institute CMU-HCII-95-102-R1. http://www.cs.cmu.edu/~amulet.

[Myers 96b] Brad A. Myers and David Kosbie. Reusable Hierarchical Command Objects. In *Human Factors in Computing Systems*, pages 260-267. Proceedings SIGCHI'96, Vancouver, BC, Canada, April, 1996.

[Myers 96c] Brad A. Myers, Robert C. Miller, Rich McDaniel, and Alan Ferrency. Easily Adding Animations to Interfaces Using Constraints. 1996.Submitted for Publication.

[Neches 93] Robert Neches, Jim Foley, Pedro Szekely, Piyawadee Sukaviriya, Ping Luo, Srdjan Kovacevic, and Scott Hudson. Knowledgable Development Environments Using Shared Design Models. In *Proceedings of the 1993 International Workshop on Intelligent User Interfaces*, pages 63-70. ACM SIGCHI, Orlando, FL, January, 1993.

[Newbery 88] Frances J. Newbery. An interface description language for graph editors. In *1988 IEEE Workshop on Visual Languages*, pages 144-149. Pittsburgh, PA, October, 1988. IEEE Computer Society Order Number 876.

[Newman 68] William M. Newman. A System for Interactive Graphical Programming. In *AFIPS Spring Joint Computer Conference*, pages 47-54. 1968.

[NeXT 91] NeXT, Inc. NeXTStep and the NeXT Interface Builder. 900 Chesapeake Drive, Redwood City, CA 94063. 1991.

[Nielsen 93] Jakob Nielsen. Noncommand User Interfaces. *Communications of the ACM* 36(4):83-99, April, 1993.

[Olsen 83] Dan R. Olsen, Jr. and Elizabeth P. Dempsey. Syngraph: A Graphical User Interface Generator. In *Computer Graphics*, pages 43-50. Proceedings SIGGRAPH'83, Detroit, MI, July, 1983.

[Olsen 86] Dan R. Olsen, Jr. Mike: The Menu Interaction Kontrol Environment. *ACM Transactions on Graphics* 5(4):318-344, October, 1986.

[Olsen 87] Dan R. Olsen, Jr. Larger Issues in User Interface Management. *Computer Graphics* 21(2):134-137, April, 1987.

[Olsen 88] Dan R. Olsen, Jr. and Bradley W. Halversen. Interface Usage Measurements in a User Interface Management System. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 102-108. Proceedings UIST'88, Banff, Alberta, Canada, October, 1988.

[Olsen 89] Dan R. Olsen, Jr. A Programming Language Basis for User Interface Management. In *Human Factors in Computing Systems*, pages 171-176. Proceedings SIGCHI'89, Austin, TX, April, 1989.

[Olsen 92] Dan R. Olsen, Jr. *User Interface Management Systems: Models and Algorithms.* Morgan Kaufmann, San Mateo, CA, 1992.

[Olsen 93] Dan R. Olsen Jr., James D. Foley, Scott E. Hudson, James Miller, and Brad Myers. Research Directions for User Interface Software Tools. *Behaviour and Information Technology* 12(2):80-97, March-April, 1993.

[Open Interface 95]
NeuronData. Open Interface. 1310 Villa Street, Mountain View, CA 94041, (415) 943-3450, sales@neurondata.com. 1995.

[OpenGL 93] Silicon Graphics, Inc. Open-GL. 2011 N. Shoreline Blvd. Mountain View, CA 94039-7311. (415) 960-1980. 1993.

[Oracle 96] Oracle Tools. Oracle Corporation, 500 Oracle Parkway, Redwood Shores, CA, 94065, (800) 633-0583. 1996.

[Ousterhout 91] John K. Ousterhout. An X11 Toolkit Based on the Tcl Language. In *Winter*, pages 105-115. USENIX, 1991.

[Palay 88] Andrew J. Palay, *et. al*. The Andrew Toolkit - An Overview. In *Proceedings Winter Usenix Technical Conference*, pages 9-21. Dallas, Tex, February, 1988.

[Pausch 92] Randy Pausch, Matthew Conway, and Robert DeLine. Lesson Learned from SUIT, the Simple User Interface Toolkit. *ACM Transactions on Information Systems* 10(4):320-344, October, 1992.

[Prograph 96] Pictorius. Prograph. Pictorius Incorporated, 2745 Dutch Village Road, Suite 200, Halifax, Nova Scotia B3L 4G7, Canada, (800) 927-4847, info@prograph.com. 1996.

[Samuelson 93] Pamela Samuelson. Legally Speaking: The Ups and Downs of Look and Feel. *Communications of the ACM* 36(4):29-35, April, 1993.

[Scheifler 86] Robert W. Scheifler and Jim Gettys. The X Window System. *ACM Transactions on Graphics* 5(2):79-109, April, 1986.

[Schulert 85] Andrew J. Schulert, George T. Rogers, and James A. Hamilton. ADM-A Dialogue Manager. In *Human Factors in Computing Systems*, pages 177-183. Proceedings SIGCHI'85, San Francisco, CA, April, 1985.

[Singh 88] Gurminder Singh and Mark Green. Designing the Interface Designer's Interface. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 109-116. Proceedings UIST'88, Banff, Alberta, Canada, October, 1988.

[SL-GMS 93] SL Corp. Suite 110 Hunt Plaza, 240 Tamal Vista Blvd., Corte Madera, CA, 94925, (415) 927-1724. 1993.

[Smith 82] David Canfield Smith, Charles Irby, Ralph Kimball, Bill Verplank, and Erik Harslem. Designing the Star User Interface. *Byte* 7(4):242-282, April, 1982.

[Stallman 79] Richard M. Stallman. *Emacs: The Extensible, Customizable, Self-Documenting Display Editor*. Technical Report 519, MIT Artificial Intelligence Lab, August, 1979.

[Stevens 83] Albert Stevens, Bruce Roberts, and Larry Stead. The Use of a Sophisticated Graphics Interface in Computer-Assisted Instruction. *IEEE Computer Graphics and Applications* 3(2):25-31, March/April, 1983.

[Stevens 94] Marc P. Stevens, Robert C. Zeleznik, and John F. Hughes. An Architecture for an Extensible 3D Interface Toolkit. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 59-67. Proceedings UIST'94, Marina del Rey, CA, November, 1994.

[Strauss 92] Paul S. Strauss and Rikk Carey. An Object-Oriented 3D Graphics Toolkit. In *Computer Graphics*, pages 341-349. Proceedings SIGGRAPH'92, July, 1992.

[Sukaviriya 90] Piyawadee Sukaviriya and James D. Foley. Coupling A UI Framework with Automatic Generation of Context-Sensitive Animated Help. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 152-166. Proceedings UIST'90, Snowbird, Utah, October, 1990.

[Sukaviriya 93] Piyawadee Sukaviriya, James D. Foley and Todd Griffith. A Second Generation User Interface Design Environment: The Model and The Runtime Architecture. In *Human Factors in Computing Systems*, pages 375-382. Proceedings INTERCHI'93, Amsterdam, The Netherlands, April, 1993.

[Sun 96] Sun Microsystems. Java: Programming for the Internet. 1996. http://java.sun.com/.

[Sutherland 63] Ivan E. Sutherland. SketchPad: A Man-Machine Graphical Communication System. In *AFIPS Spring Joint Computer Conference*, pages 329-346. 1963.

[Swinehart 86] Daniel Swinehart, Polle Zellweger, Richard Beach, and Robert Hagmann. A Structural View of the Cedar Programming Environment. *ACM Transactions on Programming Languages and Systems* 8(4):419-490, October, 1986.

[Szekely 93] Pedro Szekely, Ping Luo, and Robert Neches. Beyond Interface Builders: Model-Based Interface Tools. In *Human Factors in Computing Systems*, pages 383-390. Proceedings INTERCHI'93, Amsterdam, The Netherlands, April, 1993.

[Teitelman 79] Warren Teitelman. A Display Oriented Programmer's Assistant. *International Journal of Man-Machine Studies* 11:157-187, 1979. Also Xerox PARC Technical Report CSL-77-3, Palo Alto, CA, March 8, 1977.

[Tesler 81] Larry Tesler. The Smalltalk Environment. *Byte Magazine* 6(8):90-147, August, 1981.

[Vander Zanden 90] Brad Vander Zanden and Brad A. Myers. Automatic, Look-and-Feel Independent Dialog Creation for Graphical User Interfaces. In *Human Factors in Computing Systems*, pages 27-34. Proceedings SIGCHI'90, Seattle, WA, April, 1990.

[VAPS 95] Virtual Prototypes Inc. VAPS. 4700 de la Savane, Suite 300, Montreal, Quebec CAN, H4P 1T7, (514) 341-3874. 1995.

[Visix 96] Visix Software Inc. Galaxy Application Environment. 11440 Commerce Park Drive, Reston VA 22091. (800) 832-8668, http://www.visix.com. 1996.

[VisualEdge 90] Visual Edge Software Ltd. UIM/X. 3950 Cote Vertu, Montreal, Quebec H4R 1V4. Phone (514) 332-6430. 1990.

[Vlissides 90] John M. Vlissides and Mark A. Linton. Unidraw: A Framework for Building Domain-Specific Graphical Editors. *ACM Transactions on Information Systems* 8(3):204-236, July, 1990.

[Vlissides 91] John M. Vlissides and Steven Tang. A Unidraw-Based User Interface Builder. In *ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 201-210. Proceedings UIST'91, Hilton Head, SC, November, 1991.

[Wernecke 94] Josie Wernecke. *The Inventor Mentor*. Addison-Wesley Publishing Company, Reading, MA, 1994.

[Wiecha 90] Charles Wiecha, William Bennett, Stephen Boies, John Gould, and Sharon Greene. ITS: A Tool for Rapidly Developing Interactive Applications. *ACM Transactions on Information Systems* 8(3):204-236, July, 1990.

[Wilde 90] Nicholas Wilde and Clayton Lewis. Spreadsheet-based Interactive Graphics: from Prototype to Tool. In *Human Factors in Computing Systems*, pages 153-159. Proceedings SIGCHI'90, Seattle, WA, April, 1990.

[Wilson 90] David Wilson. *Programming with MacApp.* Addison-Wesley Publishing Company, Reading, MA, 1990.

[Wolber 96] David Wolber. Pavlov: Programming by Stimulus-Response Demonstration. In *Human Factors in Computing Systems*, pages 252-259. Proceedings SIGCHI'96, Vancouver, BC, Canada, April, 1996.

[XBusiness 94] X Business Group, Inc. *Interface Development Technology*. 3155 Kearney Street, Suite 160, Fremont, CA 94538. (510) 226-1075, , 1994.

[XVT 96] XVT Software, Inc. XVT. 4900 Pearl East Circle, Boulder, CO, 80301, USA, 1-800-678-7988 or (303) 443-4223, info@xvt.com, http://www.xvt.com/xvt. 1996.

[Zeleznik 91] Robert C. Zeleznik, et.al. An Object-Oriented Framework for the Integration of Interactive Animation Techniques. In *Computer Graphics*, pages 105-112. Proceedings SIGGRAPH'91, July, 1991.