

Research

ISAAC: a meta-CAD system for virtual environments

Mark R Mine

This paper presents a description of ISAAC; the Immersive Simulation Animation And Construction program designed and built at the University of North Carolina at Chapel Hill (UNC-CH). ISAAC is a scene composition application used for the interactive construction of virtual worlds. In ISAAC you work directly in a virtual environment; you position, orient and scale objects using direct and indirect manipulation techniques. ISAAC stores object configurations in ASCII files that it uses to recreate scenes at a later date for further manipulation and interactive exploration. ISAAC is not a modeling program; you create worlds by manipulating pre-generated three-dimensional models (which can come from sources such as computer-aided design programs or three-dimensional scanning devices). ISAAC was designed to overcome some of the limitations of working in a virtual environment and to take advantage of the natural and intuitive forms of interaction available in a virtual world. © 1997 Elsevier Science Ltd.

Keywords: immersive design, object manipulation, virtual environment, interaction techniques

INTRODUCTION

Despite considerable promise there has been a surprising lack of real-world applications in the virtual world. Why? What is so hard about working in a virtual environment? While many researchers have focused on technological limitations such as low-resolution displays, limited tracking systems, and end-to-end delays, some of the greatest difficulties facing virtual-environment application developers are conceptual rather than technological in nature. We are unfamiliar with this new medium we work in. We do not fully understand that many of the characteristics that give a virtual environment its power are also the source of its problems. Take, for example, two of the key advantages to working in a virtual world: immersion and the direct manipulation of objects.

Immersion within a virtual world means we can distribute controls (and other information) about a

user, gaining advantages in terms of the density and quantity of information presented to the user. It also means, however, that controls can be difficult to find. Even if visible, controls may obscure your view of the virtual environment or you may have to traverse the environment to reach them. This is unlike a through-the-window system where controls typically have a well-defined location on the screen and are always within reach. We lose some of the organization and efficiency forced upon us by two-dimensional computer interfaces. We are stuck with the problem of finding the light switch in the dark.

Similarly, we perform many tasks in a virtual world using direct manipulation because it is such a natural and intuitive form of interaction. Virtual worlds, however, lack many of the constraints that we depend upon to solve problems in the real-world: desktops, to constrain our motion to a plane; rulers to lock us to a line. Without these constraints, the number of real-world problems that can be solved in the virtual world using direct manipulation is limited. Humans lack the fine motor control necessary to precisely manipulate objects. This is compounded by the fact that virtual worlds lack both haptic feedback (which we depend on to constrain our motion in the physical world) and effective techniques for alphanumeric input (which we use for precise manipulation in the computer world). We are stuck with the worst of both worlds.

The virtual world, therefore, though a powerful new medium in which we can use our eyes, ears, and hands much as we do in the real world, is far from the panacea of computer-aided applications we would like to believe. We have many challenges and limitations yet to overcome.

The interaction techniques used in ISAAC and presented in this paper were chosen because they enhance the natural and intuitive forms of interaction possible in a virtual world (often in ways not possible in the real world). They augment user interaction and/or provide the types of constraints necessary for precise interaction with virtual objects. ISAAC is a testbed for virtual world interaction techniques. It is the outgrowth of research by the author into the impact of working in a virtual world. The focus of this research is the development of interaction techniques that compensate for the limitations of working in a virtual world while simultaneously exploiting its power.

Department of Computer Science, University of North Carolina at Chapel Hill, CB# 3175, Sitterson Hall, Chapel Hill, NC 27599-3175, USA

Paper received: 1 September 1995. Revised: 23 October 1996

The rest of the paper consists of: a discussion of related work in the area of interactive creation and manipulation of objects in a virtual world, an overview of the ISAAC system, examples and discussions of some of the more effective virtual-world interaction techniques used in ISAAC, and a discussion of the importance of constraints for precise interaction in a virtual world (with examples given from the ISAAC system). The paper ends with a discussion of future work and a brief conclusion.

PREVIOUS WORK

One of the earliest systems implemented for the creation and manipulation of three-dimensional objects in a virtual world was 3DM. Designed and built at the University of North Carolina⁴, 3DM is an immersive modeling program that allows users to create simple geometric objects (such as cylinders and spheres) interactively in a virtual environment. 3DM includes several grid and snap functions; it lacks, however, many of the other aids and constraints that we since have found to be necessary for accomplishing precise work.

JDCAD (by Liang and Green at the University of Alberta), though not an immersive system, tackled many of the issues involved in the interactive creation of three-dimensional objects. By using a 6-DOF input device, users could interact directly in three-dimensions. Liang and Green showed encouraging results when comparing JDCAD with conventional through-the-window modeling systems (using 2D input devices)⁹.

PolyShop, developed at the Institute for Simulation and Training at the University of Central Florida, explored important issues such as two-handed interaction and the use of physical devices (such as a desktop) to provide low-level haptic feedback⁸. Dan Mapes continues this work at MultiGen Inc. with the development of SmartSceneTM.

Conceptual Design Space (CDS), developed at Georgia Tech's Graphics Visualization and Usability Center¹, is an immersive, head-tracked application for the interactive design of architectural spaces. In CDS users create objects by specifying vertices on a ground plane that gets extruded into three-dimensional objects (such as walls). CDS includes a three-dimensional menuing system and several 3D widgets for object manipulation.

The Worlds-In-Miniature (WIM) project at the University of Virginia¹¹ demonstrates how interaction techniques not possible in the real world can be effective tools in the virtual world. In the WIM interface, a user manipulates full sized objects in the immersive world by moving around representations in a miniature copy of the world floating in front of him. The WIM interface extends a user's range of interaction by taking advantage of the ability to interact at multiple levels of scale simultaneously in a virtual environment.

ISAAC OVERVIEW

ISAAC is a scene composition application designed for the interactive construction of virtual worlds. ISAAC is not a modeling program; users create virtual worlds by manipulating pre-generated three-dimensional models. Models can be from any source such as computer-aided design programs or three-dimensional scanning devices.

ISAAC is an immersive environment application; you work directly in the virtual world. The virtual-environment system used for ISAAC consists of a stereo head-mounted display, a tracking device (to measure the position and orientation of the user's head and hand), a five-button input device (the Python) and a host and graphics computer. All images are generated in real time using the Pixel Planes 5 graphics computer⁷. ISAAC was written using the C programming language and was built upon the extensive set of virtual world libraries developed at UNC.

In a typical ISAAC session, you place models into the environment using a two-step process. First you begin by choosing the three-dimensional models to be loaded into the system at program startup. The objects, known as the working set, are the objects that will *potentially* populate your environment. If, for example, you are building a model of a room environment you might specify several types of chairs, tables and other furniture. Note that not all of these objects will necessarily end up in the final environment. During the session, you choose objects from the working set (for example a particular type of chair) and interactively place them in the environment. Note that you may simultaneously have multiple copies of the same object at several locations (several instances of a column to make a colonnade for example). If, at any time during the ISAAC session, you discover you are missing a required model from your working set you can add it using keyboard commands.

Once the user has placed an object in the environment he can position, orient and scale it using direct and indirect manipulation techniques. The user interacts with objects *at-a-distance*; selecting and manipulating them using a laser beam emanating from his hand. To select an object, the user points the laser beam at a desired object and clicks on a button on the input device. The user can select any object visible in the environment for interaction. In ISAAC you manipulate objects using direct gestural interaction (i.e. ISAAC maps movements of your hand directly to movements of the object). Alternately you can manipulate objects indirectly using the laser beam and special manipulation widgets.

Direct object manipulation can be hand-centered or object-centered. This refers to the center of rotation used by the software in mapping rotations of the hand onto rotations of the object. Hand-centered manipulation is most closely analogous to object manipulation in the real world; when you grab an object it pivots about the center of your hand. Object-centered manipulation allows the user to rotate an object about its center while physically removed from the object (standing back to get a better perspective, for example).

ISAAC includes both constrained and unconstrained object manipulation modes. In the unconstrained mode, the software maps all change in position and orientation of the user's hand directly onto the object. In constrained mode the user specifies one- and two-dimensional changes in an object's position (along any principal axis or plane) and one-dimensional rotations about any principal axis using a constrained motion widget.

When using ISAAC, you are free to walk or fly about the environment. ISAAC includes two modes of flying: pointing mode, in which the user points at the desired destination¹⁰ and orbital mode⁵, in which the user orbits a point of interest as controlled by the current orientation of his head (see next section for more details).




Figure 1 Local versus action-at-a-distance

To allow users to make changes encompassing the entire environment, the ISAAC system also includes an implementation of the Worlds-In-Miniature interface. Both grid and constrained motion functions can be used in conjunction with the WIM to allow for more controlled manipulation of objects.

ISAAC includes conventional editing functions such as: cut, copy, paste, duplicate and delete. These are selected using ISAAC's two-dimensional control panels. Also included on the control panels are functions such as grouping controls and functions for the control of ISAAC's grid.

At any time during an ISAAC session the user can save the current object configuration into an ASCII file which he can use at a later date to recreate the configuration for further manipulation and interactive exploration.

ISAAC INTERACTION TECHNIQUES

ISAAC includes several forms of interaction that

illustrate how natural interaction techniques (such as moving your head to set your viewpoint, or moving an object by grabbing it) can be powerfully enhanced in the virtual world.

Action-at-a-distance

An excellent way to augment the direct manipulation of objects (i.e. reaching out and grabbing something) is to extend your normal reach using some form of action-at-a-distance (AAAD). During AAAD interaction, you select objects using a laser beam emanating from your hand (see Figure 1). Once selected you can then remotely manipulate an object by moving your hand (with your hand motion mapped to object motion). Alternately you can use the laser beam to interact with some form of constrained manipulation widget.

A limitation of AAAD is the fact that you must be able to see an object in order to select it. AAAD is therefore most effective when combined with some form of




Figure 2 Orbital mode

viewpoint specification such as orbital mode or WIM navigation discussed below. Also, object manipulation in AAAD mode is imprecise (in particular it can amplify noise in hand-tracking data) and is most effective in directions orthogonal to the direction of the laser beam. More precise object positioning requires additional constraints and control devices.

Orbital mode

Orbital mode takes advantage of one of the most intuitive means of view specification: moving your head, but it does so in a non-obvious way. In orbital mode, objects of interest are constrained to always stay in front of you. The side of the object currently visible depends on which direction you're looking: look up and you see the object's bottom, look down and you see its top (see *Figure 2*). This allows you to quickly and naturally see all sides of an object as it *orbits* your head (or you orbit the object). Chung's experiments showed it to be the method of choice for subjects studying a large, human-scale object in a virtual world⁵. Orbital mode augments viewpoint specification by coupling location in the virtual world with head motion. It has the additional advantage that your hands remain free during viewpoint specification. Note that orbital mode is good for local viewpoint specification and is not a global navigation technique. It does not solve the general problem of navigation through complex virtual worlds.

Worlds-in-Miniature

Worlds-in-Miniature (WIM) is a powerful technique that has applications in two areas: navigation and object manipulation. Worlds-in-Miniature interaction involves the manipulation of objects in a miniature representation of the world floating in front of the user, easily within arms reach (see *Figure 3*). Movement of an object in the WIM results in the corresponding movement of the object in the immersive environment surrounding the user. Movement of a representation of the user in the WIM results in movement of the user through the immersive environment.

Interaction using a WIM extends your reach and

allows you to manipulate objects naturally by reaching out and grabbing them. The WIM brings the world to the user (versus extending the user's reach as in AAAD interaction). Since you are working with a miniature representation of the entire environment you can reach objects anywhere in the world and quickly move then across the entire environment. Working with a WIM has the advantage of giving both global and local contexts at the same time. You can see the entire virtual world in your WIM while immersed in your current location. This brings all objects within reach while simultaneously allowing you to perceive the full-size object. Thus with a WIM it is possible to adjust a painting while standing across the room to see if it is straight. The WIM interface, however, does not solve the problem of direct manipulation and precise object positioning discussed earlier.

Techniques such as WIM and AAAD are particularly effective because they allow people to interact with the entire environment without having to move. We have found conservation of effort to be a dominating principle governing radiologists and molecular structure scientists attempting to do real work in a virtual world^{2,5}.

Menu interaction

ISAAC includes two types of menuing systems: a rotary tool chooser and a set of two-dimensional control panels.

The rotary tool chooser is provided as a quick and easy way to select one of ISAAC's high level tools (fly, manipulate, scale) or to invoke the more complex two-dimensional control panels. The rotary tool chooser is a one-dimensional menu based upon J. D. Liang's ring menu⁹. With a one-dimensional menu you only have to vary one degree of freedom of your input device (i.e. rotate your hand about an axis or move your hand in a line) to move between menu selections. All other changes in the position and orientation of the input device are ignored¹⁰. The chooser consists of a set of three-dimensional icons (one for each tool) and a selection box, arranged in an arc, floating in front of your face (see *Figure 4*). To select a tool you must simply:

- (1) Invoke the chooser by pressing (and holding) a button on the input device.
- (2) Move the selection box on top of the icon of tool you



Figure 3 Worlds-in-Miniature




Figure 4 Rotary tool chooser

- wish to use by rotating your hand about the pointing axis of the input device (like turning a dial).
- (3) Release the button to select the desired tool.

Since the rotary tool chooser is a one-dimensional menu and the chosen dimension of interaction is rotation about the hand's pointing axis, the current position of the user's hand is immaterial. Thus you do not have to move your hand to coincide with an icon to select it (speeding up menu interaction).

It has been found that you can facilitate interaction with the chooser by locking it in your head space (so that it always remains in front of you). Originally the chooser would show up in an arc about the current location of the user's hand. This meant that it was necessary to look at your hand or to move your hand in view in order to interact with the tool chooser. You avoid these additional actions by locking the chooser in head space; the chooser follows the movement of your head and is always in a consistent orientation and in view. Locking the chooser in head space has the additional benefit that it makes it easier to select a tool by *feel* (based on the sensed orientation of your hand relative to your head); it is no longer necessary to look at the chooser to select a tool (further speeding up menu interaction). Users have found the rotary tool chooser to be an excellent way to choose between a limited number of options (limited by the number of options that can reasonably be displayed at one time).

ISAAC includes more complex functions in its two-dimensional control panels. Two-dimensional control panels are the virtual equivalent of conventional workstation pull down menus and dialog boxes¹⁰. Built upon UNC's Immersive Widgets library, ISAAC's two-dimensional control panels include: single-shot buttons, on/off buttons, radio buttons, sliders, dials, scrollable lists, popup menus and displays (readouts for sliders and dials). Functions included on the two-dimensional control panels include: edit functions (cut, copy, paste, duplicate, delete), grouping functions (group/ungroup), grid control functions (such as translational and rotational grid spacing), functions to add new objects, reset objects and other environmental controls.

To simplify and speed up interaction with the two-dimensional control panels ISAAC uses a point-and-shoot paradigm. The user selects a virtual control panel button or slider by pointing at it using a laser beam emanating from his hand (*Figure 5*). To activate it he



Figure 5 Two-dimensional control panel

must then click on a physical button on his input device. Only the two-dimensional location of the intersection of the laser beam with the control panel plane is used for the selection of control panel widgets; the user does not have to be able to reach a button to invoke it.

ISAAC control panels float in three-dimensional space; they are not locked into screen space (see *Figure 5*). This decision was made based upon a trade off between visibility of the control panel and obscuration of the environment. A control panel locked in screen-space is always visible, but it has the tendency to obscure the user's view of the environment. A control panel floating in three-dimensional space can be moved out of the way (allowing the user to look away from the control panel even as he interacts with it) but has the disadvantage that it can be difficult to locate. To compensate for this ISAAC includes a way to quickly snap the control panel into view.

CONSTRAINED OBJECT MANIPULATION

The absence of constraints when interacting with virtual objects is one of the major limitations in current virtual world applications; without them it is difficult to perform precise object manipulation. Virtual solutions to real-world problems will depend upon the implementation of constraints to augment and control interaction within the virtual world.

The constraints used in virtual world applications can be separated into two categories: virtual and physical. Virtual constraints are those in which the application ignores extraneous degrees of freedom in the user's input data. Thus, for example, even though the user's hand moves in all three directions during object positioning, the software constrains objects to move in a plane (ignoring any out of plane motions of the user's hand). Physical constraints are actual physical objects that you use to restrict the motion of your hand. Examples include a desktop on which you can slide an input device or something like the University of Virginia's tracked tablet¹¹, which you hold in one hand and slide the input device on with the other hand.

The current focus of ISAAC system development is the implementation of virtual constraints. The goal is to incorporate constrained manipulation modes and grid functions that will help in the precise manipulation of objects. In designing ISAAC's virtual constraints, every effort was made to keep interaction as direct and intuitive as possible.

For example, direct manipulation technique were enhanced by including controlled manipulation widgets instead of providing sliders and dials on control panels. Several of the virtual constraints currently used by the ISAAC system are discussed.

First, a snap-to-grid function is provided to allow users to place objects on grid intersections. Currently, the snap point (the point on the object aligned to the nearest grid intersection) is defined to be the center of the object's bounding box. Users set the size of the grid increment interactively using a slider on one of ISAAC's control panels. At present the grid is aligned with world coordinates, though efforts are underway to include a user-definable grid that the user can position and orient interactively.

Snap-to-orientation is an important function provided




Figure 6 Constrained motion widgets

in ISAAC to re-orient selected objects to the nearest cardinal orientation (defined to be one of the 24 possible alignments of the object's model coordinates with the eight octants of the world coordinate frame). Snap-to-orientation is a corrective function; i.e. you use it after you move an object (to snap it upright, for example).

ISAAC includes three types of constrained motion widgets for the constrained manipulation of objects in ISAAC: one-dimensional translation, two-dimensional translation, and one-dimensional rotation. The left side of Figure 6 shows the original constrained motion widgets used in ISAAC. The right side shows the same widgets after they had been redesigned to make the affordances more intuitive and to address visibility and occlusion problems suffered by the initial design. A widget is a three-dimensional object with geometry (representing affordances) and encapsulated behavior. To use a constrained motion widget you first select a widget by pointing at it with the laser beam, *grab* it (indicated by pressing a button on the input device), and then (depending on the widget selected) use the widget to move the object along a line, in a plane or rotate it about one of its axes. You can choose to have the constrained motion widgets either aligned with the model's coordinate system or the world's coordinate system.

Careful consideration was given to the interaction

between snap-to-grid and the constrained motion widgets. Initially it was decided that the snap-to-grid function should take precedence. Thus even if an object is being positioned using one of the constrained motion widgets, it would jump to the nearest grid intersection (often not on the constrained path) if snap-to-grid was on (see Figure 7). This resulted in erratic object motion and had the undesirable side effect that it was difficult to move an object back to its original position once moved.

After further evaluation it was realized that it was more useful and intuitive if the constrained motion widgets took precedence over snap-to-grid. Thus an object moved using a constrained motion widget always stays on its constrained path and will not move to the nearest grid point, even if snap to grid is turned on (see Figure 7). To allow for more controlled motion when moving objects with a constrained motion widget an object-centered snap was implemented. With object-centered snap and snap-to-grid turned on, ISAAC constrains objects to move in even increments *relative to their starting position* (their location at the start of the grab—see Figure 7). The size of the motion increment is equal to the size of the current grid spacing. Since you can align the constrained motion widgets with either the world or model coordinate system (see above), you can move an object precisely in world coordinates or relative to the object's current orientation. Using a model-aligned constrained motion widget, for example, you can move an object one unit to its left even if it is not aligned with the world coordinates.

FUTURE WORK

The primary focus of future work on ISAAC will be the implementation of additional forms of constrained interaction. Of particular importance will be the implementation of techniques for the specification of object-to-object alignments and constraints. Both user-specified




Figure 7 Snap-to-grid precedence versus constrained-motion-widget precedence

and automatic techniques (such as the innovative techniques developed by Bukowski³ will be explored. Constrained manipulation techniques such as user definable grids (that the user can move interactively to any position and orientation) and additional constrained motion widgets (such as rotation-only and translation-only object-positioning widgets) will also continue to be explored. Better techniques for alphanumeric input in a virtual world will also need to be developed.

The relative merits of distributed versus localized control will be investigated. This will involve the investigation of the advantages and disadvantages of moving controls from ISAAC's control panels into the environment (to be registered with the objects they control). Grid spacing controls, for example, might be placed directly on their associated grid lines. Though this reduces the cognitive distance between object and control and is potentially more intuitive, it introduces additional problems of visibility and controllability (due to problems such as those associated with direct manipulation).

Important work will be the development of techniques for the management and control of the increasing number of interaction modes. The goal is to avoid extensive use of menus for mode specification and tool selection. Though a modeless interface is desirable, real-world applications will require numerous modes of interaction (such as the various modes of constrained object manipulation used in ISAAC). The use of voice recognition is one possible avenue of exploration.

CONCLUSIONS

Before we can solve real-world problems in the virtual world we must gain a better understanding of the nature of the virtual environment. Current virtual world applications suffer because we are misusing this new medium we work in; we do not take full advantage of its power and we are unsuccessful in compensating for its limitations. We must learn to better exploit the natural forms of interaction possible in the virtual environment by extending them in ways only possible in the virtual world. We must learn to compensate for the lack of haptic feedback and the absence of effective techniques for alphanumeric input, both essential for precise interaction with virtual objects. The ISAAC system, presented in this paper, gives examples of the kind of effective techniques that must be developed if we are to create real-world applications for the virtual-world; techniques that take advantage of the unique characteristics of this exciting new medium.

ACKNOWLEDGEMENT

This work is funded by DARPA/ETO contract no. DABT63-93-C-0048.

REFERENCES

1. Bowman, D., Conceptual design space. Available at WWW URL <http://www.cc.gatech.edu/gvu/virtual/CDS>, 1995.
2. Brooks, F. P., Jr and Wright, W. V., Twenty-first annual progress report interactive graphics for molecular studies. UNC Chapel Hill Computer Science Technical Report, TR95-010, 1995.
3. Bukowski, R. W. and Sequin, C. H., Object associations: a simple and practical approach to virtual 3D manipulation. *ACM Computer Graphics: Proceedings of 1995 Symposium on Interactive 3D Graphics*, Monterey, CA, 1995, pp. 131-138.
4. Butterworth, J., Davidson, A., Hench, S. and Olano, T. M., 3DM: a three-dimensional modeler using a head-mounted display. *ACM Computer Graphics: Proceedings of 1992 Symposium on Interactive 3D Graphics*, 1992, **25**(2), 135-138.
5. Chung, J., Intuitive navigation in the targeting of radiation therapy treatment beams. Ph.D. dissertation, University of North Carolina, Chapel Hill, NC, 1994.
6. Connor, D., Snibbe, S., Herndon, K., Robbins, D., Zeleznik, R. and van Dam, A., Three dimensional widgets. *ACM Computer Graphics: Proceedings of 1992 Symposium on Interactive 3D Graphics*, 1992, **25**(2), 197-208.
7. Fuchs, H., Poulton, J., Eyles, J., et al., Pixel-planes 5: a heterogeneous multiprocessor graphics system using processor-enhanced memories. *ACM Computer Graphics: Proceedings of SIGGRAPH 89*, 1989, **23**(3), 79-88.
8. IST (Institute for Simulation and Training), PolyShop final report. Available at WWW URL <http://www.vsl.ist.ucf.edu/~polyshop>, 1995.
9. Liang, J. and Green, M., JDCAD: a highly interactive 3D modeling system. *Computers & Graphics: Proceedings of the Conference on Computer Aided Design and Computer Graphics*, 1994, **18**(4), 499-506.
10. Mine, M., Virtual environment interaction techniques. UNC Chapel Hill Computer Science Technical Report, TR95-018, 1995.
11. Stoakley, R., Conway, M. J. and Pausch, R., Virtual reality on a WIM: interactive worlds in miniature. *Proceedings of SIGCHI 1995*.



Mark Mine is a PhD student at the University of North Carolina working under Dr Frederick P. Brooks. His dissertation research involves the exploration of virtual environment interaction techniques and the development of an immersive modeling system for architectural design. He has a BS in aerospace engineering from the University of Michigan and an MS in electrical engineering from the University of Southern California. Before coming to UNC he worked at NASA's Jet Propulsion Laboratory on the Voyager and TOPEX/Poseidon products as a systems engineer.