

Tangible + Virtual = A Flexible 3D Interface for Spatial Construction Applied to DNA

Steven Schkolne
Caltech

Hiroshi Ishii
MIT Media Lab

Peter Schröder
Caltech

ABSTRACT

Ideas from tangible interfaces and VR ease a difficult spatial design task: construct a DNA molecule with desired characteristics. Our hybrid interface has both the physical intimacy of tangible media and the versatility of 3D digital display. Two new physical affordances: a raygun and a grip tool, enable kinesthetic control of the addition and removal of structure. We introduce 3D local menus which select multiple functions for each tool. New interactions for sensed tongs enable the sophisticated multi-object arrangement that the delicate, intricate DNA construction task demands. These flexible tools allow UI designers to create multiple interfaces upon the same physical substrate. In a user study, practicing research scientists expressed a strong preference for Silkworm, our 3D interface, when compared to mouse/monitor UI. We show that 3D tangible interfaces, heretofore only applied to freeform artistic creation, also facilitate intuition in the highly structured task that is our focus.

Keywords

Spatial construction, tangible user interface, virtual reality, 3D interface, tongs, lightsaber, DNA design, molecular biology.

INTRODUCTION

Scientists researching molecular biology encounter a difficult spatial design task. Their objects of study are so small that they can't physically engage them. They are so intricate that diagramming them with 2D paper and pencil is insufficient. Traditional interfaces support viewing, but are not rich enough to aid in creation. These problems are a special case of the *spatial construction* problem encountered by engineers, artists, and scientists alike. In this paper we focus on the task of creating structures out of DNA, currently explored by Seeman [6], Winfree [21], and many others. This specific problem is of increasing importance as we begin to design increasingly complex molecular machines.

submitted to CHI 2003




Figure 1: Tangible creation: Our experimental interface supports drawing, cutting, and bonding DNA in 3D space. Here the user draws part of a DNA cube directly in a head-tracked stereo display. This tangible tool can be repurposed to cut bonds between molecules, or draw single DNA strands.

We use composite images like this one throughout the text as they are more accurate representations of user experience than direct photographs. For footage of the system in use, please see the accompanying video.

The scientific visualization community has used virtual reality (VR) to *view* data [5] but never (to our knowledge) to *construct* it. Perhaps this is because constructing models requires intricate spatial manipulations, and even moving a single object through virtual space is still an area of active research [3, 16]. Our problem of DNA construction involves difficult tasks, such as making three double-helices intersect so that their six strands smoothly merge (pairwise) into three (see user study). This smoothness is crucial to the success of a DNA molecule when realized in the lab [18].

How can we make such manipulation feasible? We draw inspiration from Tangible User Interfaces (TUIs) [10] which use physical objects to add affordances to data. There are many success stories in the literature which show tactile cues enhance spatial understanding [9]. Current work either occurs in 2D space (using projectors to overlay data onto the physical interface [20]), or a 3D




Figure 2: Manipulation with sensed tongs: Holding the molecule with the left hand, the user sweeps in with the right-hand tong to adjust the position of a single DNA base (one of the blue cones).

space that is not a display [2,8]. Neither of these approaches is sufficient because our problem is *three-dimensional* and our data is *more dynamic* than static physical models allow.

Our solution is a hybrid: we use the tangible affordances of TUIs to control a flexible virtual data representation, all of which occurs in an integrated 3D space (the semi-immersive Responsive Workbench [13]). We present new interactions for moving objects and specifying topological changes. More importantly we do this with a lightweight, variable toolset that provides the appropriate affordances for several tasks without cluttering the workspace. In a user study, practicing research scientists compared our interface (which we call Silkworm¹) to a traditional 2D mouse + monitor interface. Our survey shows they all preferred the experimental interface – most enthusiastically so. Several of them suggested that Silkworm is superior to pencil and paper for thinking creatively about the possibilities of DNA.

THE TASK

Figure 3, which shows a user completing a hairpin molecule, illustrates the basic components of the modeling task. The cones represent *bases* – groups of molecules that, for the purposes of this study, can be thought of as atomic. Between the bases are two types of bonds – Phosphate bonds (the thin lines along the curved region in the left of the photo) and Hydrogen bonds (the thick lines seen spanning the helix in the bottom right of the image). This placement of atoms and bonds constitutes a design, whose fulfillment of project-specific goals (such as: can these molecules interlock to tile space?) is dependent on the physical plausibility of the distances and orientations of




Figure 3: Drawing a bond: The user points at a molecule, squeezes the raygun's trigger, and drags to create a bond. Pointing at the second molecule and releasing the trigger completes the operation.

the molecules envisioned. To aid in this process, our software contains a physical simulation which moves the bases into configurations based on the topological structure of the bonds.

Immersive systems have been applied to artistic creation [12, 17]. These systems succeed in VR because of their ability to capture gesture. The DNA design problem raises new issues because gesture and emotional content are *not* important, but rather the demands on spatial properties are more severe (successful design depends on interlocking components). For this reason the 3D intuition gained by merging VR and tangible interface has an important role in nanotechnology.

FLEXIBLE TANGIBLE TOOLS

Silkworm supports several operations. An early implementation had six physical tools, one per task, on the table surface. Users were frequently confused, spending significant time managing physical devices and their associated wires.

Overloading physical tools with local menus (similar to pie [4] and marking [14] menus) solved this difficulty. A menu button illuminates a halo of functionality options (see Figure 6). A more prominent action button is used to activate each tool. It is instructive to think of affordance and action separately: The tangible handle enhances spatial manipulation while the virtual presence supports tool variability. Multipurpose tangible handles present a practical future where a single setup can support many applications with the intimacy and spatial understanding that is the strength of these interaction metaphors.

We considered a number of other options for overloading physical tools. Selecting from a traditional menu (either floating in air or on the table surface) directs attention away from the area of action. A selection device, such as a

¹ because delicate things are being woven




Figure 4: Removing bonds: *The lightsaber (in right hand) cuts a bond between molecules while tongs, held in the left hand, move the structure in space.*

dial or toggle button, forces users to remember which tool is currently active to make quick selections (this information is quickly forgotten when a tool is put down).

RAYGUN: PICKING POINTS IN SPACE

We built a raygun tool by dismantling a toy gun, rewiring the trigger as the action button, inserting a motion tracker, and adding a menu button where the hammer would be (see Figure 8). The physical shape and cultural connotation of this tool supports the metaphor of pointing at objects. A virtual beam emanates from the gun's tip, adding precision to this choice. This collocation of physical and virtual object breaks down the boundary between bodily and digital space.

Historically, many VR interfaces do object selection by asking users to move a stylus so that it is collocated with an object in space [13]. Since tracking is often imprecise, and more importantly the physical tool blocks the rendered object, picking in this manner breaks the illusion of immersion. Research on distant object selection [3] typically focuses on objects beyond the arm's reach. DNA bases are not distant as much as they are small and numerous. By providing the appropriate affordance, the gun allows for precise selection within a dense volume.

Specifying Bonds

In Silkworm, we use the raygun to draw bonds between bases. This decision follows an extensive investigation into the general problem of specifying connections between objects that extended beyond the molecular setting. We think of these links as *glue*. An early interface had a gluepit, an area on the table where portions of objects could be placed to make them sticky. Gluing consisted of dragging an object to the gluepit, then placing it on the target object to form a bond. This approach forces the user to break the spatial layout of their design which is often nontrivial to reconstruct. Another approach dragged the glue itself from




Figure 5: DNA's spatial complexity: *The face of a DNA cube, drawn by a subject in the user study. Note that this structure consists of four double helices (dashed boxes on left) that are precisely rotated so that they form a smooth chain without kinks (highlighted on right). The existence and structural integrity of such chains is crucial to the stability of DNA.*

the gluepit (with tongs) directly onto objects. With this approach it is difficult to place the glue in the right spot – either it accidentally bonds to the wrong location, or the target area is occluded by the tool or the glue. In an application such as DNA construction where the scene is highly cluttered with very small objects, these solutions did not suffice.

The raygun allows precise specification of the beginning and endpoints of glue over a large volume with a small motion. The glue is started by pointing the ray at the starting point and clicking. Dragging the ray to the endpoint and releasing the trigger completes the bond.

Picking, as the ubiquitous mouse pointer attests, is a highly versatile operation. In future work, we hope to explore other applications such as: a flamethrower, which activates a local physical simulation; an ice gun, which freezes a region to turn off simulation; and spraypaint which changes a region's color.

GRIP: 3D LOCATION CONTROL

The grip (see Figure 8) enables the accurate positioning of 3D objects in space. In Silkworm, this prop controls double-helix drawing, single-strand drawing, and the lightsaber (see Figures 1 and 4). As with the raygun prop, an action button triggers the operation and a second button displays the local menu.

There is some overlap between the raygun and the grip – if the virtual tools of the raygun are mapped to the grip and vice versa, what have we lost? The raygun has an explicit connotation of directionality – it refers to space away from the tool, while the grip is better at referring to space close to itself. The grip also affords rotation around its central axis (the lightsaber's axis). We did not take advantage of this in our current application, in part because our implementation has wires which constrain rotation. Our discus-




Figure 6: Grip tool: The grip has an action button (under the thumb), a menu button (under the index finger) and an embedded magnetic motion sensor. Pressing the menu button (bottom left) activates the local menu. Moving the tool towards the single-dot icon (bottom right) activates single-strand drawing. The double-dot icon represents double-helix drawing, and the line represents the lightsaber tool.

sion of DNA placement below yields further insight into the differences between the raygun and the grip.

Of our three tools, the grip is the most generic – meaning that if a task is non-specific, the grip is a good choice. In contrast with the traditional VR stylus (whose selection point is its tip) it implicitly places the virtual part of a tool somewhat away from the hand, preventing the physical affordance from breaking the illusion of immersion.

Creating DNA

Both single strands and double-helices need to be placed in space. Note that these tasks are slightly different in character: single strands can be arbitrarily curved while double-helices have a curvature limits (less than 150 base pairs cannot make a stable ring [18]). In Silkworm, moving the strand tool through space places bases at evenly sampled intervals. The iconographic representation for this task is a sphere (representing a base) sitting in a target circle.

To draw helices, an icon with two bases on opposite sides of the target circle is used. An interpolating curve based on the orientation of the grip at the beginning and end of this stroke restricts the curvature to physical limits.

This interaction allows direct placement of both DNA structures with the appropriate control. We experimented with emitting DNA along the raygun’s ray. This interaction required significant planning to control one endpoint of the DNA, and little control of double-helix curvature. An extra step of moving the DNA was often required.

A lightweight interface idea was to emit DNA from a button on the table’s edge. While lessening the load on our




Figure 7: Doubly sensed tongs: (top right) Two foil sensors detect weak and strong grabs with the tongs. The tongs have three states: open (top right), weakly closed (bottom left) and strongly closed (bottom right). Weakly grabbing a base moves the whole molecule, strongly grabbing a base moves only it. Grabbing empty space with the tongs moves the whole scene, grabbing empty space twice (there are two sets of tongs) activates combine scaling, translation, and rotation.

physical props, this interface had even less control of DNA curvature and stroke placement than the raygun interface.

Severing bonds with the lightsaber

In contrast to the gluing and stroke-drawing tasks, the decision to use the lightsaber as a cutting tool was quite immediate. The use of swords, daggers, and knives to cut is culturally established. Surprisingly, the most functional form of the saber was much shorter than that seen in Star Wars [15]. This is because the user is not engaging a distant enemy, but rather a nearby item amongst a host of other things that shouldn’t be cut. Errors with the lightsaber are also greatly reduced by activating it only when the action button is depressed.

TONGS: MULTISCALE MANIPULATION

Schkolne et al. [17] use tongs to manipulate objects which live in a single coordinate system. Silkworm extends this metaphor to a more complex task in which (1) separate objects need to move relative to one another and (2) objects at different scales need to be manipulated.

Local and Global transformations

When the tongs are closed, they latch on to the closest molecule within range. This enables individual molecules to be moved relative to one another. While this natural interaction seems simple, its implementation raised some difficult issues. How can the whole design be moved? There are many potential solutions in the virtual navigation literature, all of which require extra hardware or a bulkier interface [19]. In Silkworm, tong grabs that are not close to any molecule move the whole scene. Due to tracker noise, visual and audio cues are necessary to im-



Figure 8: The raygun tool: Pressing the menu button (right) displays the two current options for the gun: drawing Phosphate bonds (thin line) and drawing Hydrogen bonds.

plement this effectively – we draw lines between a selection cursor (hovering between the tong tips) and any candidate molecules.

Multiscale manipulation

Observe that the tongs form a second point of contact when they are squeezed tightly (see the shadow in Figure 7). A foil sensor maps this contact to a *strong grab* which moves a single base. This interaction takes advantage of the physical structure of the tongs, avoiding a more abstract form of selection such as the local menus used above.

Successful designs operate on multiple scales – from large-scale placement of helical regions to the angles between individual bases. Grabbing empty space with both tongs and moving them towards one another reduces the size of the model (and vice versa). Although in our sample application every molecule is inherently at the same scale, in other applications we use this interaction to change relative scale.

This pair of tongs combines naturally (see Figure 7) to enable interactions ranging from rotating the whole design while changing the placement of a molecule to moving two individual bases at the same time.

ASSESSING SILKWORM

We evaluated our system by comparing it with a mouse + monitor interface based on established techniques (see appendix for a thorough description of this interface).

Experimental Setup

The subjects are six research scientists (PhD students and postdocs) who study (or studied) DNA, one of whom is female. We asked them to build several molecules with both interfaces (which we referred to as 2D and 3D) and then fill out a questionnaire describing their experience. We alternately started with the 2D and 3D interfaces. Each interaction was described and the subjects demonstrated their understanding of each tool. Following this we ran five timed trials where we asked the subjects to draw a DNA hairpin, then a Holliday junction, followed by a DNA cube. The hairpin is seen in Figure 3. A Holliday junction [18] consists of two aligned helices whose strands cross from one helix to the other. The DNA cube is significantly more complex – each edge of the cube is a dou-

ble helix, each face has one continuous piece of DNA circling it, and at each corner the three intersecting helices swap strands with one another. We gave the subjects five minutes maximum on each molecule.

Hypotheses

We anticipated that each system would have its unique advantages. The strength of each interaction depends on its contribution to the overall experience of the user, the foundation for intuition. Towards this goal, we anticipated the following distinctions between the 2D and 3D systems:

Crisp vs. Mushy: The primary advantage of the 2D system is that it is *crisp*: each action has a definite precision. Conversely VR is *mushy*: tracker noise, coupled with body noise as users hold tools in free space, both of which are accentuated by the head-tracked nature of the camera, make it difficult to act consistently. We predict that crisp 2D, where a mouse pointer can be finely adjusted to a few-pixels of resolution will be stronger for tasks like picking, or careful drawing of regions.

Spatial Mapping: The primary advantage of 3D is that users don't have to constantly make mappings from 2D to 3D and vice versa. This occurs both during visualization, when users make decisions based on the depth of molecules, and more importantly during manipulation, where specific placement and rotation of objects is needed. In these cases, the 2D interface inherently requires multiple steps to specify 6 degrees of freedom, while in 3D all of them can be changed at once.

Other differences which we predict will have less effect on this task: Parallelism is present in the two handed 3D system, but as the users are novices we don't predict that they will be able to take advantage of this. The gestural advantages of 3D systems, seen when they are applied to free-form geometry, have little relevance for this task.

Results

All of the subjects preferred the 3D interface – many displayed great enthusiasm for working with Silkworm². For example, one user, being told his time for the 3D task was up, complained “oh, but I'm having so much fun!” This comfort comes despite extensive familiarity with 2D mouse + monitor style interfaces. (only two of our subjects had used semi-immersive systems, and those only very briefly). The primary strength of the experimental 3D interface seemed to be the natural rotation and placement of objects in space:

The 3D interface was less interference between me and the molecule. Working in the 2D interface, I was spending my time figuring out how to position the space so that I could access the relevant parts of the molecule with the 2D tools. That is not the kind of creative thinking I want to be doing!

² We encourage readers to visit <http://wormsilkworm.net/firms.com> to read the full surveys.

Looking at different parts of the molecule by moving my head was very natural. It felt like there was no “interface” at all. Rotating and/or moving the space or molecules with a single pair of tongs was very natural. [quotes from user #1]

In contrast, the users found spatial management quite difficult with the 2D interface:

I had trouble rotating things and understanding what was closer to me and what was farther away. Also, I didn't really know what I was doing with the rotation except when I was rotating about the axis normal to the screen. [#2]

In particular, I couldn't predict the effect of the different types of rotation. [#3]

Our subjects did not find many strengths with the 2D interface. Many subjects said that drawing helices was the best of the 2D interactions. But those same subjects also said that drawing helices was easy in the 3D task. Others cited familiarity and portability as advantages. Cutting was found to be easier in 2D:

Cutting a bond was easier for me in 2D than in 3D, perhaps because I knew I had to position everything so that the bond was clearly visible and distinct from surrounding clutter, so when I was ready to cut it, it was easy to cut.

Cutting bonds with the sword, I had to concentrate to make sure I would be cutting the right thing. I would have to make sure my point of view gave me a good view of the entire blade. [both quotes from #1]

This was our only evidence of the crispness advantage we hypothesized. Another user preferred the 3D interface for cutting:

Cutting bonds precisely with the lightsaber tool is easier because the plane of rotation of the cutting edge can be changed. [#6]

Some users found Silkworm superior to pencil and paper for sketching out ideas, while saying that the 2D system would be best used in addition to pencil and paper:

When using the 2D interface, I wished I had pencil and paper so that I could sit and sketch things, and make a plan of attack. I never thought this with the 3D interface. When using the 3D interface I immediately saw things that would be very difficult to put on paper, and I felt that the interface was a very natural tool for trying things out. [#2]

The 2D tool didn't seem like a big improvement over pencil and paper, even though it was representing a 3D model. It might still be useful, but it was kind of a hassle to use, so I'm not currently inclined to use it. [#1]

The greatest difficulties with 3D were accidentally triggering a strong grab with the tongs when a weak one was attempted

I had trouble with the “weak” vs. “strong” usage of the tongs [#3]

The distinction with the tongs between moving a single atom or an object should be made crisper, the squeeziness of the tongs is a little subtle (but I like it actually). [#6]

We can see this as mushiness – we anticipate fixing this in future iterations of tong design.

Specific aspects of the interface aside, the gestalt of Silkworm aided these scientists in the intuitive manipulation of structure that is crucial for new research insight. When asked which interface better supported creative thinking and spatial manipulation, the subjects responded:




Figure 9: 2D interface: *The beginnings of a DNA cube, drawn with a 2D interface during our user study. The subject drew two planar faces, but had trouble aligning them correctly.*

The 3D interface: it gives me a much more accurate picture of what is really happening. I don't waste time thinking about geometric misunderstandings, and can really think about what I am building. [#2]

Certainly the 3D -- I can more completely see the medium with which I'm working [#3]

3-D! It seems more natural, you don't need to remember which keys are which (though with time, it might not matter), but it's helpful to be able to “grab” something just like you would in reality [#5]

The 3D interface, without question. By just glancing at the image, I have a better understanding of the structure. But the value of the 3D interface is much more than just nice rendering. By being able to intuitively manipulate the structures I could have a manual understanding that augments the visual understanding. I usually think about 3D objects with my hands, and this interface suited me very well. [#2]

While running the study, the difference between agitation using the 2D interface and enthusiasm for Silkworm was apparent. As user #4 simply wrote: “3D is more fun.” The connection between enthusiasm, attention, and insight suggests that interfaces such as Silkworm could have a crucial role in technological innovation.

It is important to note that we did *not* see a qualitative difference between the molecules designed in the two systems. While all users managed to build successful hairpin molecules within the allotted time, very few finished the Holliday junction, and none finished the cube. For these more complex designs, the users spent much time building strategies for construction. We feel that differences in design quality would emerge only after more experience.

CONCLUSION

Silkworm was found highly effective by a group of demanding, highly knowledgeable users. Why is this the case, when VR has traditionally had limited success in real applications? Spatial construction tasks such as DNA design are inherently three dimensional (unlike the tasks studied by Cockburn [7]), and Silkworm provides a range of manipulations of 3D space. The interactions are all centered around what 6DOF trackers do best: manipulating coordinate frames in space [11]. This hybrid of tangible

input and virtual output takes each technique for what it does best. We get both a highly flexible data representation and tangible affordances that provide immediacy and control. The interactions span physical and virtual space, providing a direct connection between the user's body and the 3D display space without introducing occlusions that break the illusion of immersion. This intimate connection fosters the comfort and play that lead to insight – a valuable commodity in the emerging field of molecular design.

We are continuing to investigate DNA construction in collaboration with active researchers in the field. This will enable us to place more challenges on the interface as we investigate increasingly complex problems. Placing more functions on each tangible tool will explore the possibility of general-purpose tangible 3D interface. DNA design is just one application of the tangible + virtual approach. One can design micro-electro-mechanical devices (MEMS), investigate novel theorems in geometry, or design wiring schemes for buildings with similar tools. Architecture, sculpture, and the industrial design of 3D shapes are also opportunities for further study.

ACKNOWLEDGMENTS

REFERENCES

1. Alias|Wavefront, Maya 4.0, 2001.
2. Anderson, D., et al, Tangible Interaction + Graphical Interpretation: A New Approach to 3D Modeling, *Proceedings of SIGGRAPH 2000*, 393-402.
3. Bowman, D., and Hodges, L., An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments, *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, 35-38.
4. Callahan, J., Hopkins, D., Weiser, M., and Schneiderman, B., An Empirical Comparison of Pie vs. Linear Menus, *Proceedings of CHI 1988*, 95-100.
5. Chen, C., *Information Visualization and Virtual Environments*, Springer Verlag, 1999.
6. Chen, J., and Seeman, N., The Synthesis from DNA of a molecule with the Connectivity of a Cube, *Nature*, 350:631-633, 1991.
7. Cockburn, A., and McKenzie, B., Evaluating the Effectiveness of Spatial Memory in 2D and 3D Physical and Virtual Environments. *Proceedings of CHI 2002*, 203-210.
8. Frazer, J.H., Frazer, J.M., and Frazer, P.A., Intelligent Physical Three-Dimensional Modelling System, *Proceedings of Computer Graphics 80*, 359-370.
9. Hinckley, K., Tullio, J., Pausch, R., Proffitt, D., and Kassell, N., Usability analysis of 3D rotation techniques, *Proceedings of the 10th Annual ACM Sympo-*

sium on User Interface Software and Technology (UIST '97), 1-10.

10. Ishii, H. and Ullmer, B., Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms, *Proceedings of CHI 1997*, 234-241.
11. Jacob, R., Sibert, L., The Perceptual Structure of Multidimensional Input Device Selection, *Proceedings of CHI 1992*, 211-218.
12. Keefe, D., Acevado, D., Moscovish, T., Laidlaw, D., and LaViola, J., CavePainting: A Fully Immersive 3D Artistic Medium and Interactive Experience, *Proceedings of the 2001 Symposium on Interactive 3D Graphics*, 85-93.
13. Krüger, W., Fröhlich, B., The Responsive Workbench, *IEEE Computer Graphics and Applications*, May 1994, 12-15.
14. Kurtenbach, G., Buxton, W., User Learning and Performance with Marking Menus, *Proceedings of CHI 1994*, 258-264.
15. Lucas, G., *Star Wars*, 20th Century Fox/Lucasfilm Ltd., 1977.
16. Pierce, J., and Pausch, R., Comparing Voodoo Dolls and HOMER: Exploring the Importance of Feedback in Virtual Environments, *Proceedings of CHI 2002*, 105-112.
17. Schkolne, S., Pruett, M., and Schröder, P., Surface Drawing: Creating Organic 3D Shapes with the Hand and Tangible Tools, *Proceedings of CHI 2001*, 261-268.
18. Sinden, R., *DNA Structure and Function*, Academic Press, 1994.
19. Tan, D., Robertson, G., Czerwinski, M., Exploring 3D Navigation: Coupling Speed-coupled Flying with Orbiting, *Proceedings of CHI 2001*, 418-425.
20. Underkoffler, J., and Ishii, H., Urp: A Luminous-Tangible Workbench for Urban Planning and Design, *Proceedings of CHI 1999*, 386-393.
21. Winfree, E., *Algorithmic Self-Assembly of DNA*, PhD thesis, Caltech, 1998.
22. Zeleznik, R.C., Herndon, K., Hughes, J., SKETCH: An Interface for Sketching 3D Scenes, *Proceedings of SIGGRAPH 1996*, 163-170.

APPENDIX

Implementation

We use 6DOF magnetic trackers (the Ascension Motion-Star) to track the user's head and the physical tools. The buttons and pressure sensors are controlled by the I-Cube system. Shutterglasses enable stereo viewing on our cus-

tom 1.8 x 1.3m Responsive Workbench. Our computer has two 2GHz x86 processors.

Our software is built from scratch in C++ using OpenGL for graphics. The unifying concept is the Ether, which contains both Things (drawn objects) and Constraints (which continually affect the positions of things). We designed constraints to help DNA maintain its natural structure. All of the sensed Tools affect this Ether, which passes the signals to the underlying Things. The Things all compute basic functions (draw, distance checks, etc) which are combined by the Tools to perform operations. Our multi-threaded implementation provides continuous interaction during expensive operations. For more information, please contact the authors.

2D Interface

We based our comparison interface on Maya[1], a popular commercial 3D modeling package. In particular, we took the camera controls and the method for manipulating individual objects from this system. As shown in Figure 9, clicking icons selects different tools. A modifier key (Alt) temporarily activates the camera: Alt + (left mouse button) rotates the camera, arcball-style. Alt + (middle mouse button) translates the scene in the screen plane. Alt + (left and middle mouse buttons) scales the scene – moving the mouse to the left reduces, moving it to the right enlarges the molecules.

The iconic tools are:

Translate molecule: Clicking on a base illuminates a square around the base in the screen plane, and three coordinate axes in the base's local coordinate system. Clicking

on the center square and moving the mouse translates the molecule in the image plane. This is analogous to a weak grab with the tongs (note that the tongs allow rotation and translation to occur at the same time).

Translate base: Similar to move molecule, but this time only the selected base is moved. This is analogous to a strong grab with the tongs.

Rotate molecule: Clicking on a base, a local arcball tool is drawn which controls rotation of a molecule.

Rotate base: Similar to rotate molecule, this tool only affects a single base.

Hydrogen bond: Clicking on base 1, dragging the mouse, and releasing it on base 2 forms a Hydrogen bond between base 1 and base 2. This is analogous to the raygun.

Phosphate bond: Same as above, creating a Phosphate bond.

Cut tool: Clicking, dragging, and releasing draws a straight line between the endpoints. Any bonds intersecting this line are broken. Analogous to the lightsaber.

Helix draw: Clicking, dragging, and releasing draws a double-helical region between the endpoints in the image plane. Note this is similar to stroke-based screenspace drawing programs such as SKETCH [22]

Strand draw: Clicking and dragging draws a path in the 2d screen plane consisting of linked bases, the orientation of which is specified by the direction of mouse movement.