

Immersive Design of DNA Molecules with a Tangible Interface

Steven Schkolne
Caltech
schkolne@caltech.edu

Hiroshi Ishii
MIT Media Lab
ishii@media.mit.edu

Peter Schröder
Caltech
ps@cs.caltech.edu

Abstract

This paper presents an experimental immersive interface for designing DNA components for application in nanotechnology. While much research has been done on immersive visualization, this is one of the first systems to apply advanced interface techniques to a scientific **design** problem. This system uses tangible 3D input devices (tongs, a raygun, and a multipurpose handle tool) to create and edit a purely digital representation of DNA. The tangible controllers are associated with **functions** (not data) while a virtual display is used to render the model. This interface was built in collaboration with a research group investigating the design of DNA tiles. A user study shows that scientists find the immersive interface more satisfying than a 2D interface due to the enhanced understanding gained by directly interacting with molecules in 3D space.

Keywords: tangible user interface, molecular visualization, props, molecular modeling, spatial construction, virtual reality, augmented reality, responsive workbench, DNA design

1 Introduction

Molecular designers encounter a difficult spatial construction task. Their objects of study are so small that they can't touch them. They are so intricate that diagramming them with 2D paper and pencil is insufficient. Prior research has demonstrated the advantages of immersion for spatial design applications such as CAD [25] and artistic design [19, 26], as well as the visualization of molecules [8]. In this paper we extend this work to the domain of scientific design, specifically a DNA engineering task where scientists design DNA tiles that act as computational elements.

Our approach couples a stereoscopic display with tangible controllers. Users wear shutterglasses to view a shape floating above a table surface (the Responsive Workbench [20]). They alternately hold four custom input devices whose physical form is task-specific. Our method builds on the research of the virtual environment, augmented reality, and tangible interface communities with some important distinctions. Traditional tangible interfaces render data in physical form and respond to direct manipulation by changing the state of the physical objects and/or an image projected on the tangible controllers. In our approach we make the **functionality** tangible, not the **data**. The data, displayed virtually, is dynamically controlled by a realtime simulation which uses constraints to loosely approximate the underlying molecular physics. The tan-




Figure 1: Here the user draws part of a DNA cube. The tangible controller we use for this task 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. Directly photographed, one would see a double-image of the molecule lying flat on the table surface.**

gible controllers have distinct physical features which guide their use (affordances). We use each physical form to perform several related operations which we call a **function class**. Each function class is related to the affordances of the device that supports it. For example, the tongs can grab molecules in three different ways to support three different types of move operations.

This project began by studying a group of scientists researching DNA computation. They were unsatisfied with current computer tools which allowed them to view (sometimes stereoscopically) molecules but offered limited ability to change their structure. While they sometimes worked with physical ball-and-stick models similar to Watson and Crick's famous model, they more often found a hybrid solution where ideas were communicated via sketches, special notations, and ad-hoc solutions such as using whiteboard markers to represent DNA helices. These scientists wanted to directly manipulate molecules, and moreover see (via computer simulation) how structures would eventually form in the laboratory. Thus our hybrid solution of virtual display and physical control.

Our primary contribution is to apply prior results in visualization and interface to a molecular design problem. We extend the body of research in tangible interface with our approach of making




Figure 2: 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.

functionality tangible. We also show that scientists were satisfied with this system, preferring it to paper-and-pencil for conceptual design. This final contribution is noteworthy because, historically, computer-based conceptualization tools that rival pencil and paper have been few and far between. This paper shows that immersion has an important role in scientific design.

2 Previous Work

Many advanced interactive techniques have been applied to molecular visualization. In the late 1960's and early 1970's, the GROPE system displayed molecules with force feedback [4]. Since the early 1990's, much molecular visualization has been carried out in the CAVE [11], a room-sized virtual environment. The table-sized Responsive Workbench used in this project is smaller, meaning that it is cheaper and easier to install, at the cost of a smaller virtual volume. For a review of visualization in virtual environments, please see Chen's text [8]. The Nanomanipulator project [30] is a noteworthy prior work because it is one of the few that allow users to change the data they are visualizing. The Nanomanipulator transmits user motions directly to an atomic force microscope. In our prototype users create a molecule from scratch as a digital model, with the intent to realize it in the lab at a later date. To our knowledge, the only work to deal with a molecular construction task is Fjeld's Augmented Chemistry system [13], where users can construct simple molecules (such as water) with a tangible interface. In comparison, our problem has a much higher graphical and syntactic complexity. Our problem is a real one faced by contemporary researchers, while Fjeld's aim is to create an educational tool.

Tangible User Interfaces (TUIs) [17] are characterized by multiple physical objects, each with embedded sensors, combined with a digital display. One of the earliest such systems was Hinckley's interface for neurosurgical visualization [15]. In this system either a doll's head or a spherical controller could be used to position the skull of a patient. The distinction between the doll's head and the sphere shows the role of affordances in input device design. The head-shaped input device has orientation cues and a cultural link to the head of the patient that the sphere does not. Norman's [23] notion of an affordance is a visual cue to an object's function. We extend this notion by regarding tactile cues as affordances.




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.

In tangible interfaces, the physical devices themselves can either act as data or as functional elements. In interfaces such as Frazer's architectural experiments of the early 1980's [14], and many others [2, 24, 32], the model representation is tangible. Note that this physical form is often enhanced by a digital projection.

The role of tangible controllers as functional devices is less developed, particularly because in many systems data and function overlap in the same device. In Fitzmaurice's Bricks system [12], physical bricks move digital objects, such as control points of curves, on a two-dimensional display. While each brick is an operator, it is also in a sense a data element because the physical brick is identified with the digital form for small periods of time. In the Shared Space system [5], cards control digital characters and other pieces of information. The cards are functional, dictating changes in digital system state, yet they are also linked to the data that is displayed on each card (in some cases the character printed on a card matches the digital object that is virtually displayed on top of a card).

Our treatment of tangible controllers is closer to Hinckley's seminal work on neurosurgical visualization [15] as tangible objects control the model state. In our system the identification of physical controller and digital object is much weaker. Instead of the affordances of our physical form depicting a relationship to a data element (as Hinckley's head does), our affordances display a **functional** relationship to the data. The simplest example is the tracked tongs that we use to move and stretch objects. The tool's form is related to the grabbing function, not to the data element being grabbed. If the form was to relate to data, then the tongs would instead look like a DNA molecule. One could say that our affordances are weak, while the doll head's affordance is strong. As we will discuss towards the end of the paper, this weakness allows a device to afford many modeling operations.

One of the chronic difficulties of virtual and augmented reality is the lack of tracking precision (see Azuma [3] for a survey of augmented reality). Yet molecular design depends on exactitude. Unlike approaches that use a stylus or sensed glove to pick points in space, our interactions do not require the physical and virtual worlds to precisely overlap, and thus allow quite a bit of tracker noise (we use electromagnetic 6DOF trackers). This noise tolerance stems from the particular flavor of mapping between tangible input device and virtual operation. Moreover, since final molecular positions are derived by constraints, users don't have to accurately specify positions with the trackers.




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.

3. The DNA Construction Task

The experimental interface addresses the task of creating structures out of DNA, currently explored by Seeman [9], Winfree [33], and others. A matching DNA sequence (ACTG...) is the foundation upon which the subtle and difficult spatial design problem rests. For our prototype we used a simplified model of DNA where we ignore sequence and treat each base as an indivisible unit (in reality each base consists of several atoms). Topological relationships between bases are used as constraints by our realtime simulation engine. Scientists can export data from our design tool to a more precise simulator as necessary (accurate molecular simulations currently require hours of computation).

DNA construction ranges from simple positioning to involved tasks, such as making three double-helices intersect so that their six strands smoothly merge (pairwise) into three, as seen in Figure 2. The smoothness of strand intersections is crucial to the success of a DNA molecule when realized in the lab [28].

Figure 3 shows a user completing a hairpin molecule. In this image we can see 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 the molecules envisioned.

The following modeling operations provide for the creation and editing of these structures:

- create double-helices
- create single strands of DNA
- cut bonds between molecules
- create Hydrogen bonds
- create Phosphate bonds
- move molecules
- move bases
- move the whole scene (navigate)
- turn simulation on and off




Figure 5: 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).

4. Local Menus: Repurposing Tangible Controllers

We began with an approach where each physical tool maps to a single task [17]. This quickly led to a lot of clutter in the workspace. Users were frequently confused, spending significant time managing physical devices and their associated wires. We addressed this issue by overloading each tangible input device, so that it performs multiple functions. Local menus (similar to pie [7] and marking [21] menus) allow users to select a function from a set of possibilities. A menu button illuminates a halo of options (see Figure 6). A more prominent action button is used to activate each tool.

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 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).

The following sections detail each physical tool in turn, describing its form, construction, and the functions it supports.

5. The Raygun: Picking Points in Space

Many VR interfaces support object selection by asking users to move a stylus so that it is collocated with an object in space [20]. Since tracking is often imprecise, and more importantly the physical tool blocks the rendered object, picking in this manner breaks the illusion of immersion. We chose to use a pointing metaphor, similar to techniques seen in Mine's work [22], to select DNA bases from within a cluttered molecule.

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 support the metaphor of pointing at objects. A virtual beam emanates from the gun's tip, increasing the precision of this choice. Error is reduced by moving the gun closer to the site, increasing the size of the target in angular space.




Figure 6: Handle tool: The handle 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.

5.1 Specifying Bonds

As illustrated in Figure 3, the raygun draws bonds between bases. This choice 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 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.

Note the raygun is not the only tool that could support this action. This is a weak affordance — the gun’s form supports, but does not require a certain use. Weak affordance allows for **tool generality**. This generality is seen in another application where we use the raygun to spraypaint surfaces. Due to the weak affordance the same raygun controller was successfully used in these two tasks from very different domains [27]. Note how tools with stronger affordances (a bond-shaped stick, a spraycan) could not be shared across applications.

6 The Handle: 3D Location Control

The handle (see Figure 6) enables the accurate positioning of 3D




Figure 7: Doubly sensed tongs. Two foil sensors on each side (top left) 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 with two sets of tongs simultaneously activates combined scaling, translation, and rotation.

objects in space. This device controls double-helix drawing (see Figure 1), single-strand drawing, and the lightsaber cutting tool (see Figure 4). As with the raygun, an action button triggers the operation and a second button displays the local menu.

There is some overlap between the raygun and the handle – if the virtual tools of the raygun are mapped to the handle 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 handle is better at referring to space close to itself. The handle 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 discussion of DNA placement below yields further insight into the differences between the raygun and the handle.

Of our three tools, the handle is the most generic – meaning that if a task is non-specific, the handle is a good choice. In contrast with the pen-shaped VR stylus, the handle is held in the power grip (like a tennis racket), not the precision grip (as one holds a pencil). Power-grip metaphors are better suited to environments with low spatial resolution due to tracker noise.

6.1 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 limit (less than 150 base pairs cannot make a stable ring [28]). In the experimental interface, 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 handle at the beginning and end of this stroke restricts the curvature to physical limits.



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.

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

A lightweight interface idea was to emit DNA from a button on the table’s edge. While lessening the load on our physical devices, this interface has even less control of DNA curvature and stroke placement than the raygun interface. After experimenting with these alternate solutions, we settled on the handle metaphor as it offers users the most control.

6.2 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. 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.

7 The Tongs: Multiscale Manipulation

Schkolne et al. [26] use tongs to manipulate objects which live in a single coordinate system. We extend this metaphor to a more complex setting. For DNA design, three separate movement tasks are required: (1) the whole scene needs to be moved as a single coordinate frame, (2) separate molecules need to move relative to one another and (3) DNA bases within molecules need to be manipulated. We use a combination of proximity information and a differentiation between strong and weak grabs to support these requirements, as detailed in the following two sections.

7.1 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 [29]. In the experimental interface, tong grabs that are not close to any molecule move the whole scene. Due to tracker noise, visual and audio cues are necessary to implement this effectively – we draw lines between a selection cursor (hovering between the tong tips) and any candidate molecules.




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.

7.2 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 (moving them apart increases the size). Although in our sample application every molecule is inherently at the same scale, in other applications we use this interaction to change relative scale.

These tongs cooperate naturally to enable interactions ranging from scaling the scene by grabbing it with both tongs (a similar approach is found in 2D with the metaDESK [31]), to rotating the whole design while changing the placement of a single molecule within it. Figure 5 shows one molecule being moved with the left hand’s tong, while the right hand’s tong moves a base within that molecule. This is accomplished with a weak grab with the left tong and a strong grab with the right. The dominant hand performs fine grain positioning tasks within the coordinate frame established by the nondominant hand. The effectiveness of this style of task division is well established in the literature on bimanual interaction [15].

8 User Satisfaction

Our interest in immersive solutions to the DNA engineering task is founded on a belief that spatial intuition is better supported by 3D interfaces than traditional 2D mouse-based systems. We wish to test this hypothesis, and at the same time we realize that intuition is emotional, fundamentally qualitative. We performed a comparison by exposing scientists both to the immersive 3D system described above and a 2D version that uses traditional mouse/monitor interface (see Figure 9). While no quantitative differences emerged, our survey documents a healthy enthusiasm on the scientists’ part for the experimental system.

We built our own 2D interface, in part to maximize similarity with the 3D tool, and moreover because there are currently no commercial tools available for this task. See the Appendix for a specification of this interface. These two interfaces utilize the same underlying rendering and simulation engines (although note the display in the 2D case is monoscopic).

8.1 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 [28] 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 double 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.

8.2 Results

All of the subjects preferred the 3D interface – many displayed great enthusiasm for working with the experimental interface. For example, one user, being told his time for the 3D task was up, complained “oh, but I’m having so much fun!” The primary strength of the experimental 3D interface seemed to be the natural rotation and placement of objects in space:

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]

Some users found the experimental system 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]

In later research we experimented with a metaphor where tong modes are more explicit [27]. In this system local menus were used to transition between the three different tong modes. This reduced user error.

When asked which interface better supported creative thinking and spatial manipulation, the subjects responded:

3D! 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]

This study focuses on satisfaction rather than effectiveness and efficiency. We did not see a qualitative difference between the molecules designed in the two systems, and would not expect one to emerge without extensive user practice. Moreover, we feel that this test of effectiveness and efficiency is better performed on a system more thoroughly developed than our prototype.

9 Discussion

We sought out to enhance spatial intuition during the design process. Our hybrid approach uses techniques from the tangible UI and virtual reality communities. Tangible devices each control a small set of operations which we call a **function class**. A function class is a set of similar operations that are supported by similar affordances. Tongs grab, the raygun points, and the handle holds tools that are swept through space like a tennis racket. It is instructive to think of function classes as atomic units of interface design. The mouse represents picking and dragging, and this modality is assigned to a set of tasks, a function class that the mouse tangibly affords. The joystick, with identical degrees of freedom, does not support the same function class. With the trackball there is considerable overlap with the mouse.

The raygun supports a function class of pointing and directing. It is not hard to imagine other tools being supported by the raygun. Consider flamethrowers which activate local physical simulation, ice guns which freeze regions to turn off simulation, or spraypaint which changes a region’s DNA base type. It is difficult to imagine these tools being afforded by the tongs. 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.

The immersive tangible interface was enjoyed by a group of demanding, highly knowledgeable users. Spatial construction tasks such as DNA design are inherently three dimensional (unlike the tasks studied by Cockburn [10]), and 3D interfaces support direct

manipulation of 3D space [6, 29]. The interactions are all centered around what 6DOF trackers do best: specifying coordinate frames in space [18]. This hybrid of tangible input and virtual output captures the unique benefits of each approach. 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 hybrid 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.

Acknowledgements

This work was supported in part by NSF (ACI-0219979), the DOE (W-7405-ENG-48/B341492), Alias|Wavefront, Designworks/USA, Pixar, Microsoft, and the Packard Foundation. Special thanks to the Winfree lab at Caltech.

References

- [1] Alias|Wavefront, Maya 4.0, 2001.
- [2] Anderson, D., Frankel, J.L., Marks, J.W., Agarwala, A., Beardsley, P.A., Hodgins, J.K., Leigh, D.L., Ryall, K., Sullivan, E., Yedidia, J.S., Tangible Interaction + Graphical Interpretation: A New Approach to 3D Modeling, Proceedings of SIGGRAPH 2000, 393-402.
- [3] Azuma, R.T., A Survey of Augmented Reality, Presence: Teleoperators and Virtual Environments 6:4, 355-385, 1997.
- [4] Batter, J.J. and Brooks, F.P., GROPE-I: A Computer Display to the Sense of Feel, Proc. IFIP Congress 1971, 759-763.
- [5] Billingham, M., Shared Space: Collaborative Augmented Reality, SIGGRAPH 1999 Conference Abstracts and Applications, p. 178.
- [6] 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.
- [7] Callahan, J., Hopkins, D., Weiser, M., and Schneiderman, B., An Empirical Comparison of Pie vs. Linear Menus, Proceedings of CHI 1988, 95-100.
- [8] Chen, C., Information Visualization and Virtual Environments, Springer Verlag, 1999.
- [9] Chen, J., and Seeman, N., The Synthesis from DNA of a Molecule with the Connectivity of a Cube, Nature 350:631-633, 1991.
- [10] 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.
- [11] Cruz-Neira, C., Daniel J. Sandin, D.J., DeFanti, T.A., Kenyon, R.V., Hart, J.C., The CAVE: Audio Visual Experience Automatic Virtual Environment, Communications of the ACM 35:6, 64-72, 1992.
- [12] Fitzmaurice, G.W., Ishii, H., Buxton, W., Bricks: Laying the Foundations for Graspable User Interfaces, Proceedings of CHI 1995, 442-449.
- [13] Fjeld, M., B. Voegtli, Augmented Chemistry: An Interactive Educational Workbench, Proceedings of IEEE and ACM International Symposium of Mixed and Augmented Reality 2002, 259-260.
- [14] Frazer, J.H., Frazer, J.M., and Frazer, P.A., Intelligent Physical Three-Dimensional Modelling System, Proceedings of Computer Graphics 80, 359-370.
- [15] Hinckley, K., Pausch, R., Goble, J., Kassell, N., Passive Real-World Interface Props for Neurosurgical Visualization, Proceedings of CHI 1994, 452-458.
- [16] Hinckley, K., Tullio, J., Pausch, R., Proffitt, D., and Kassell, N., Usability analysis of 3D rotation techniques, Proceedings of UIST '97, 1-10.
- [17] Ishii, H. and Ullmer, B., Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms, Proceedings of CHI 1997, 234-241.
- [18] Jacob, R., Sibert, L., The Perceptual Structure of Multidimensional Input Device Selection, Proceedings of CHI 1992, 211-218.
- [19] 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.
- [20] Krüger, W., Fröhlich, B., The Responsive Workbench, IEEE Computer Graphics and Applications, May 1994, 12-15.
- [21] Kurtenbach, G., Buxton, W., User Learning and Performance with Marking Menus, Proceedings of CHI 1994, 258-264.
- [22] Mine, Mark, Exploiting Proprioception in Virtual Environments, PhD thesis, UNC Chapel Hill, 1997.
- [23] Norman, D.A., The Psychology of Everyday Things, Basic Books, 1988.
- [24] Ben Piper, B., Ratti, C., Ishii, H., Illuminating Clay: a 3-D Tangible Interface for Landscape Analysis, Proceedings of CHI 2002, 355-362.
- [25] Sachs, E., Roberts, A. and Stoops, D.: 3-Draw: A Tool for Designing 3D Shapes, IEEE Computer Graphics & Applications 11:6, 18-26, Nov. 1991.
- [26] 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.
- [27] Schkolne, S., 3d Interfaces for Spatial Construction, PhD Thesis, Caltech, 2003.
- [28] Sinden, R., DNA Structure and Function, Academic Press, 1994.
- [29] Tan, D., Robertson, G., Czerwinski, M., Exploring 3D Navigation: Coupling Speed-Coupled Flying with Orbiting, Proceedings of CHI 2001, 418-425.
- [30] Taylor, R.M., Robinett, W., Chi, V.L., Brooks, F.P., Wright, W.V., Williams, R.S., Snyder, E.J., The Nanomanipulator: A Virtual-Reality Interface for a Scanning Tunneling Microscope, Proceedings of SIGGRAPH 1993, 127-134.

[31] Ullmer, B., Ishii, H., The metaDESK: Models and Prototypes for Tangible User Interfaces, Proceedings of UIST 1997, 223-232.

[32] Underkoffler, J., and Ishii, H., Urp: A Luminous Tangible Workbench for Urban Planning and Design, Proceedings of CHI 1999, 386-393.

[33] Winfree, E., Algorithmic Self-Assembly of DNA, PhD thesis, Caltech, 1998.

[34] Zeleznik, R.C., Herndon, K.P., Hughes, J.F., SKETCH: An Interface for Sketching 3D Scenes, Proceedings of SIGGRAPH 1996, 163-170.

Appendix

Implementation

We use 6DOF magnetic trackers (the Ascension MotionStar) 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 custom 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 multithreaded implementation provides continuous interaction during expensive operations. For more information, please contact the authors.

2D Interface

We based our comparison interface, which uses a mouse and a monoscopic monitor, 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. Similar methods are seen in molecular visualization software packages. The icons shown at the top of Figure 9 select 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 [34]

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.