Just as free-form textual input allows players to be expressive, the human expression embodied in *Façade*'s strong voice acting is part of what makes the experience so effective. The nuances of line performance make things by turns funnier and more uncomfortable—and overall, more engaging. On the other hand, *Façade* also required significant authoring effort to appropriately funnel a vast number of possible system states into a much smaller number of possible prerecorded utterances.

If the system had used plain text as output, many more system states could have been experienced meaningfully on the work's surface. This would have been true for two reasons. First, the mixing of prewritten segments could have been more fine grained. *Façade* already sounds strange when it mixes recordings of the player character's name into lines that were specially recorded for the purpose, and going further in this direction would have greatly reduced what the experience gains from its voice acting. Second, plain text output would have opened the door to another interesting research area: the construction of systems like the literary augmented grammars of *Brutus*. These have the potential to adjust the nuance of text, by applying hand-authored rules, depending on the current system state. Yet for performance-oriented systems (such as *Façade* and many modern computer games), this too is best regarded as an area of research. Though speech-generation systems are rapidly improving, for many purposes voice acting will remain the more powerful approach. But for systems such as massively multiplayer online games, in which most NPC dialogue takes place through text, such research could have immediate application.

Agency in *Façade*

Finally, there is also an element of *Façade* that communicates clearly to players and, when natural language understanding is successful, operates much as they expect. This is the combination of drama management, beat goals, joint behaviors, handlers, and mix-ins—what I have called *Façade*'s procedural script. This combination creates a performed story that progresses in a way the audience understands, during which players can direct conversation toward a range of topics, and which can take a variety of shapes that culminate in an appropriate ending.

This experience is not simply that of the *SimCity* effect, because it is not designed to foreground the nature of *Façade* as a system. Rather, as Mateas and Stern indicate, their goal is an experience in interactive media commonly termed "agency" or "intention." *Façade*'s authors point to Janet Murray's formulation of the concept, which has been particularly influential in academic circles: "Agency is the satisfying power to take meaningful action and see the results of our decisions and choices" (1997, 126). In the field of game design the idea is often associated with Doug Church, who writes of "allowing and encouraging players to do things intentionally" (1999)—understanding the game world well enough to make and execute a plan of action, then seeing a clear reaction from the game world. At the same time, it is also worth noting that a version of this concept can be seen from the earliest full-length writing on digital fictions of which I am aware: the PhD dissertation of Mary Ann Buckles. She describes this in relation to the psychological concept of "effectance" as "the desire for competence and feeling effective in dealing with the surrounding environment" (1985, 37). She discusses how the world of the early interactive fiction game *Adventure* works to build this experience in its audience.

The *SimCity* effect is one way to build the experience of agency. *Façade*'s procedural script seeks another. Mateas argues that in general, "a player will experience agency when there is a balance between the material and formal constraints" (2004, 125). The formal constraints are present in the shape of the fictional world, motivating some actions and not others. The material constraints, on the other hand, are the resources available for action. Many contemporary games for children and adults create a sense of agency by presenting a fictional world in which one would want to move across space (e.g., to rescue a princess or find necessary information) and fight enemies (e.g., who are mean animals or are zombies/aliens/robots/Nazis), and providing exactly the tools necessary (well-developed mechanics for spatial movement and combat).

Given this, achieving agency is relatively well understood for game worlds with simple fictional worlds and simple available actions, because it is easy for players to understand and act within the bounds of the system.[20] The *SimCity* effect describes a route to agency for more complex systems and retrospectively clarifies that even simple movement and combat require transitioning from the audience's initial expectation to an understanding of the model implemented by the system (see sidebar: Agency Reconsidered). What *Façade* provides, instead—in its procedural script—is an example of a route to agency that aims to function with less necessity for player understanding of the system. Rather than needing to replace the audience's initial *Eliza* effect notion of drama with one closer to the model of the system, players who

[20] Though design failures happen commonly, as when actions that can be used in one context cannot be used in another, for reasons unexplained by the fictional world, reducing the player's ability to formulate and execute intentions.

avoid natural language understanding errors can experience agency in *Façade*'s world while continuing to operate largely based on mental models drawn from theater, media, and human interaction. We might say that this is the true dream of interactive drama.

Unfortunately, this also means that the dream remains somewhat elusive. The *Façade* approach, while powerful, still requires that players learn its mechanics for movement, speaking, and physical action. As discussed earlier, it also suffers from unavoidable errors in natural language

Agency Reconsidered

The concept of *agency* or *intention,* discussed in this chapter primarily in connection with *Façade,* is central to thinking about gameplay and interaction in general. In short, agency is a term for the audience's ability to form intentions, take actions, and see satisfying results. It is a widely-used concept, but also one in need of reconsideration along a number of dimensions.

First, discussions of agency can fall into potentially misleading formulations in which player goals and plans appear to determine player actions. As discussed in chapter 5, such formulations have been severely critiqued in cognitive science and artificial intelligence by researchers such as Lucy Suchman and Philip Agre. As outlined in that context, it is important to remember that goals and plans are *resources* for action—which is fundamentally situated and improvisational. Reflecting back on personal experiences of playing a challenging game, in which things rarely go exactly as one plans, is another good reminder. In this vein, perhaps the best antidote to such misleading constructions of agency is *Pilgrim in the Microworld* by David Sudnow (1983), an account of learning to play *Breakout* for the Atari VCS.

On a different front, discussions of agency also commonly bracket a fundamental issue. This is apparent in Murray's primary example of game world navigation as well as in Church's discussion of design strategies that work (e.g., simple and consistent controls, abilities, and physics behavior make planning possible). The fundamental issue is the motivation that players have to move, or plan to move, to particular places. Put another way: What creates the desire that agency satisfies?

The first part of an answer is the fundamental contribution of Mateas's formulation of agency as a balance between material and formal constraints. This moves agency from an amorphous ability to plan and carry out actions, which captures only part of the picture, to a phenomenon in which the actions motivated by a game or other understanding. An alternative, such as players directly expressing discourse acts (rather than typing free-form text), would require teaching the audience to understand the system more deeply, reinforcing the importance of the *SimCity* effect.

Beyond *Façade*

In addition to research in natural language understanding and text generation, *Façade* also points to further work in a number of other areas. For example, in developing the

piece of digital media are matched with those it enables. In other words, agency requires evoking the desires a work satisfies.

With the examples discussed in this book, we can now take a further step. Consider *Eliza/Doctor,* which evokes the expectation of a therapy session, motivating the audience to say the things one might say in such a situation, providing the tools necessary for saying them (direct textual input), and responding to each with a further question. While this creates a desire for action and a means of taking that action (and while, as discussed in chapter 4, Murray has called Weizenbaum "perhaps the premier" literary artist in the field [1997, 72]) the result of play with *Eliza* is breakdown. Working with the audience's expectations (e.g., of a therapy session) is a powerful way to evoke desires for action, but the *Eliza* effect is not a route to agency.

SimCity, as discussed in this chapter, also begins with audience expectation—using it to evoke desires to take city planning actions using the tools represented on its surface. This initiates a process designed to transition players, through experimentation and feedback, from their initial assumptions to an understanding of its procedural city. This understanding is what enables agency in *SimCity,* and it accomplishes this at a level much more ambitious than simply moving through virtual space.

This lesson is key because we do not have a holodeck and will not have one in the foreseeable future. We do not have the perfect virtual reality, simulation, and AI technologies that would be necessary to enable the vision of free-form interaction that depends on the same desires and supports available in the everyday world. Even the simpler examples of Murray and Church depend on players coming to games with an expectation of movement and then transitioning to the system's actually implemented model. The *SimCity* effect makes this explicit and takes it further, setting the stage for a more mature understanding of agency.

approach embodied in *Façade*'s procedural script and the large amount of material necessary to support it, Stern and Mateas also developed a set of code templates that they reused repeatedly. As they have speculated, the design elements formalized in these templates could become the primitives in higher-level authoring languages for digital fictions, or even in graphical or AI-assisted authoring tools.

In the meantime, *Façade* has demonstrated the power of the ideas formulated in its procedural script—which should in time find their way into other digital media projects. As authors adopt the techniques pioneered by *Façade,* their strengths and limitations will become yet more apparent. New research will explore new possibilities. As this happens, hopefully, mainstream gaming will abandon the static dialogue tree, thereby increasing the potential richness of existing genres and opening the possibilities for new ones.

Further, while these relatively well-developed techniques—the cumulative result of more than a decade of technology and design research—may be adopted by mainstream digital media, researchers are likely to turn to the next challenge. Experiences like *Façade* will remain limited by the sheer amount of material that must be authored. Future research is likely to consider how to not only produce a wider variety of sentences based on authored rules (as in LAGs) but also automatically assemble structures more akin to *Façade*'s JDBs or even *Façade*'s larger beats, as the current situation dictates. Until this succeeds, *Façade*-like experiences will require large amounts of authoring for every minute of dramatic experience.

More generally, *Façade* provides an important example of how authors can engage the history of AI. It adopts powerful tools that grow from an AI legacy. At the same time, it disposes of limiting concepts that are the baggage of that legacy—especially the insistence on seeing characters as fully "autonomous agents" who sense and react to each other, but should be strongly compartmentalized from one another and the larger system. Of course, like many ideological views, such AI concepts were often compromised in practice, but *Façade* moves away from them as a principled stance, rather than with embarrassment.[21] For those willing to do the conceptual work required by such reconsiderations, the history of AI provides an exciting source of future directions.

Finally, *Façade* manages to create an emotionally engaging, if frequently uncomfortable, interactive experience about the relationships between people. It manages this not simply through the operations of its procedural script and characters but also because it is designed to capitalize on the power of fiction's most powerful tool: language. In this it provides a stark contrast to the history from which it flows, from *Tale-Spin* to the Woggles. But in this way it connects to another, parallel history—a particular branch of which is considered in the next chapter.

[21] Principled moves away from such views, and toward architectures oriented to audience experience, were more common in the field of computer graphics. Earlier in this chapter, I discussed the Improv system, which did…

Notes continued at end of this chapter.

Notes

1. Though an expansion focused on crafting the underlying terrain on which the city rests.

2. During the review of this book's draft manuscript on Grand Text Auto, Dennis G. Jerz offered a telling series of observations—including that the graphic user interface (GUI) of *SimCity* is already a much more restricted interaction channel than *Eliza*'s free-form textual input. In his comment he writes:

> On the one hand, you've got the naive user who sits down to *Eliza* with a preconception of the idea of a therapist and a preconception of what it means to communicate through typing—but probably no idea of what a text parser does. On the other hand, you've got the naive user who sits down in front of *SimCity* with preconceptions about cities, and with a pretty good idea of what it means to use a GUI.
>
> Any GUI interface restricts the meaningful gestures the user can perform, just as the iconic display restricts the possible meanings that the image can convey. So the initial "surface" of the GUI is already a mental model, in a way that I don't think the "surface" of *Eliza* is (at least not until the user has first noticed something fishy about *Eliza*'s responses).
>
> The resolution of the information that goes into *Eliza* is a lot deeper than the information that goes into *SimCity*. Even if *Eliza* can't actually understand all of that information, the fact that it's there on the screen (or on the fan-fold printout) means that it was available for the user to refer to when interpreting *Eliza*'s responses. But yes, limiting the input to yes/no (or some other finite set of verbs) would replicate the restrictions of *SimCity*'s GUI.
>
> On the other hand, *SimCity* is a much more complex program. Once you're consciously aware of *Eliza*'s rules, there's not much left to do, unless of course the leading questions guide you to a life-changing self-revelation. In *SimCity* you have to learn how the various resources interact over time.
>
> The naive users who first played *Eliza* weren't consciously trying to grok the rules of the text parser, but may instead have been very self-conscious about what might have been their first significant encounter with a computer. If we put *SimCity* in historical perspective, we can assume that most players of *SimCity* were familiar with GUIs, and were thus willing to accept restrictions (such as the requirement to zone rectangular areas and not being able to put a railroad, street, and power line in the same space). The GUI has trained us to lower our expectations. But those restrictions break the illusion of being a "real" city planner just as much as *Eliza*'s gaffes break the illusion of talking to a real therapist.

I think Jerz's comments are generally insightful. At the same time, though, I would argue that the experience of *Eliza* or *SimCity* is not fruitfully considered in terms of being a real patient or real mayor. Rather, I think we might employ the conception of interaction as conversation or dialogue, as outlined by authors such as Chris Crawford (2003) and Myron Krueger (1977). For such authors, a successful interactive system has to both "speak" and "listen." When interacting with *SimCity,* the GUI tells us the tools we have available for speaking (acting) and listening (gathering information), while experimentation reveals that the system genuinely has an underlying model of the city that shapes its actions (speaking) and is informed by our actions (listening). Interacting with *Eliza,* on the other hand, reveals that it does not have a model of the conversation—it's not listening in any ongoing way.

3. Though it is technically true that models used by the CBO are open to scrutiny, as are most of those used by nongovernment researchers such as Forrester, Starr tells us that "to most participants in policy debates as well as the public at large, the models are opaque. Only a few can penetrate the black box and understand what is inside" (1994). *SimCity,* on the other hand, shows that carefully designed playability can be an effective tool for helping much larger groups understand how simulations operate. At the same time, as Nick Montfort pointed out in the blog-based review of this manuscript, it is critical to bear in mind the fact that the understanding reached through the *SimCity* effect is limited. In Montfort's comment he writes:

> I get the impression that the lasting impact of Forrester's model from *Urban Dynamics* was largely based on it being open to scrutiny and described in detail in a 300-page book—which led to productive discussion, the development of other models, and so on. Yes, *SimCity* is playable, but it actually isn't as well-documented and modifiable as Forrester's model is, and the assumptions of the model aren't directly expressed.

Since Montfort's remark, the release of the *SimCity* source code (under the name *Micropolis*) has altered how potentially well-documented and modifiable people find *SimCity*'s model. But Montfort's point remains important. While the *SimCity* effect opens system understanding in a manner that is more generally accessible, the understanding gained has crucial limits, and for many purposes must be supplemented by other sources of information.

For this sort of deeper information on the CBO's simulation work, see "Overview of the Congressional Budget Office Long-Term (CBOLT) Policy Simulation Model" (O'Harra, Sabelhaus, and Simpson 2004).

4. The broad learning that takes place around video game playing, rather than the specific learning about the underlying system that is my focus here, is discussed in books such as *What Video Games Have to Teach Us about Learning and Literacy* by James Paul Gee (2004).

5. Something Montfort tells me one's player character might do while experiencing *Dangerous Curves* by Irene Callaci (2000).

6. Interactive fiction (IF) players are often engaged in discovering what words can be employed to take action, rather than how statements can be constructed for parsing, in each fiction's simulated world. While some actions are possible by convention and standard in most IF tool systems (such as compass-point movement), the possibility space for commands is nearly as large as the language used to converse with the fiction. Some authors have implemented verbs such as "remember" or "dare" in their games, as appropriate to the settings and themes involved, which are generally discovered by players through a combination of experimentation and careful attention to descriptions of the game world. Players work to discover not only what commands are accepted by the system but also what their effects are on the underlying simulation. In other words, to what extent and in what ways are the commands simulated? In a particular fictional world, is "remembering" simply a way of bringing up a small snippet of memory text, or does it add new areas of the world for exploration, or does it even expand the set of future actions that the player character may take? This is part of what Douglass (2007) describes as the "implied code" that IF audiences come to understand.

7. Of course, Wright's games are certainly not the only ones for which this is the case. Alex Galloway makes a similar point about the work of Sid Meier, a designer famous for simulation-oriented games such as *Civilization* (Meier and Shelley 1991):

> In the work of Meier, the gamer is not simply playing this or that historical simulation. The gamer is instead learning, internalizing, and becoming intimate with a massive, multipart, global algorithm. To play the game means to play the code of the game. To win means to know the system. And thus to *interpret* a game means to interpret its algorithm (to discover its parallel "allegorithm"). (Galloway 2006, 90–91)

8. This description is based on lecture notes by Kenneth Forbus and Wright (2001).

9. From Wingfield's article:

> Critics have also charged that the game wrongly equates happiness with consumerism, since much of it revolves around buying clothing, furniture and other goods.
>
> Mr. Wright has said the game actually parodies such habits. The Sims's more conspicuous consumers spend a lot of time fixing broken refrigerators, tending to malfunctioning cars and otherwise being controlled by their property. (2006)

Alex Galloway, in his book *Gaming,* goes further—writing that "*The Sims* is a game that delivers its own political critique up front as part of the gameplay. There is no need for the critic to unpack the game later" (2006).

10. Frasca writes that he has met people who:

> firmly believe that *The Sims* is a parody and, therefore, it is actually a critique of consumerism. Personally, I disagree. While the game is definitively cartoonish, I am not able to find satire within it. Certainly, the game may be making fun of suburban Americans, but since it rewards the player every time she buys new stuff, I do not think this could be considered parody. (2001)

11. At the time of the blog-based review of this book's manuscript on Grand Text Auto, Richard Evans (already noted as the AI lead for *Black & White*) was working at Maxis on the not-yet-announced title *The Sims 3*. In his comment on this paragraph he wrote:

> This is very perceptive. In fact, there is disagreement at the moment amongst the designers on an as yet unannounced product, specifically about this issue. Some want to manifest the agent's interior mental states, for the reasons you give. Others worry that showing everything ruins "player projection"; they argue "the more you show, the less the player will project."
>
> I disagree—I think that "player projection" (the phenomenon where the player imputes more to the agents than is really going on) is a phenomenon which stands on the shoulders of the player having a clear mental model of how the agents operate. It is only if he understands *most* of why they do what they do that he will impute more to them. If he doesn't have a clear mental model of their interiority, because the designers have hidden the data from him, then he won't be able to project—

> you can't stand on the shoulders of nothingness!

Assuming that Evans was talking about *The Sims 3*, I was fascinated to hear that such debates take place even within Maxis. Andrew Stern (2008) took the opportunity of my exchange with Evans to create a new top-level post on Grand Text Auto, continuing an earlier conversation thread about NPCs and transparency. An interesting consensus seemed to emerge in the ensuing conversation: that the pleasure of coming to understand a system's processes through experimentation requires the ongoing availability of system data. The question is which data to expose to the audience, and how.

In some ways this exchange anticipated the discussion, later in this chapter, of interactive drama and *Façade*—as well as harking back to the previous chapter's section on expressive language generation. Each of these seeks ways to expose ongoing system data through the means of the arts: dramatic performance, crafted language, and so on. Of course, the design of *The Sims* is more like *SimCity*. While there are characters, their iconic speech is closer to an animation of a congested roadway than a nuanced dramatic line. For such projects, underlying data must be exposed by other means.

12. This places her work at Atari and on her dissertation in roughly the same time period as *Universe* and the first phase of *Minstrel*. Some of Laurel's dissertation was later adapted for her book *Computers as Theatre* (1991).

13. The founding of Zoesis certainly wasn't the death knell for university-based research in interactive narrative. Around the time that the Oz Project was winding down, R. Michael Young's Liquid Narrative project was gaining momentum at North Carolina State University. As of this writing, their work continues in high gear—and a variety of other groups internationally are engaged in interactive narrative research. Exploring all this work would require a book of its own, so I direct interested readers to the proceedings of the Conference on Artificial Intelligence and Interactive Digital Entertainment, the Joint International Conference on Interactive Digital Storytelling, and the ongoing series of symposia on related topics from the Association for the Advancement of Artificial Intelligence.

14. From the animation direction, work led by researchers such as Norm Badler and Jessica Hodgins focused on biomedically correct animation rather than expressive animation (Badler, Phillips, and Webber 1993; Hodgins et al. 1995). From the behavior side, researchers such as Barbara Hayes-Roth sought to find the right general-purpose "social-psychological model" for interactive characters while Bruce Blumberg was among those building non-human characters based on ethological models (e.g., dog behavior) (Rousseau and Hayes-Roth 1998; Blumberg 1997).

While invested in non-media forms of correctness, much of this work went on to have significant impact on media projects. For example, the C4 architecture developed in Blumberg's group stands in the background of Jeff Orkin's work on *F.E.A.R.* as well as Damian Isla's work on AI for the *Halo* series. In fact, Isla undertook graduate work in Blumberg's group at MIT and his influential "behavior tree" formulation for AI characters (2005) is closely connected to the tradition described in this chapter (e.g., it could be implemented as a programming idiom for a language like ABL).

15. Bates reports on a particular Woggle:

> Due to a programming error, Shrimp occasionally peppers his ongoing behavior with what seems to

> be a nervous tick causing him repeatedly to hit his head against the ground. This attracts people's attention immediately, but to our surprise they build psychological theories, always incorrect, about Shrimp's mental state and seem to find him much more interesting and alive for having this behavior. (1994, 124)

Years later, Zoesis created a new implementation of the agent architecture. Testing *OttoAndIris.com* it found something similar. As Loyall notes:

> In an early version of the system, kids testing it drew pictures afterwards of Otto as a "crybaby," and kept talking about the time he refused to sing. The refusal was a bug that caused part of Otto's mind to completely freeze. We thought the bug had ruined the test, but to the kids it showed Otto's strong will and made him seem more alive. (2004, 7)

These behaviors succeed, of course, because they are great opportunities for *Eliza*-style projection. But an entire system composed along these lines would have been a dismal failure. They succeeded by being distinctive within a larger context of behavior that seemed ordinary and appropriate.

16. Though similar ideas had been pursued in domains such as battlefield simulation (e.g., Tambe 1997).

17. In the initial draft of this chapter I referred to "*Façade*'s script." During the blog-based peer review, Stern objected, saying the term *script* has a strong implied linearity—in both traditional drama and game development. The new term is my attempt to capture the fact that *Façade*'s procedural script specifies what characters will say and do (like a traditional script) while actually composing ongoing behavior through a combination of drama management, beat goals, joint behaviors, handlers, and mix-ins.

18. Many make the small selection of available actions visible on the work's surface. Some, however, like text-based interactive fictions, make the search for available actions part of the experience.

19. This is something that Steven Dow and Blair MacIntyre have explored in creating the "wizard" interface for their augmented reality version of *Façade* (Dow 2008).

20. Though design failures happen commonly, as when actions that can be used in one context cannot be used in another, for reasons unexplained by the fictional world, reducing the player's ability to formulate and execute intentions.

21. Principled moves away from such views, and toward architectures oriented to audience experience, were more common in the field of computer graphics. Earlier in this chapter, I discussed the Improv system, which did not commit strongly to agent autonomy and commonly directed characters in coordinated performance. During the blog-based review of this manuscript, Mark J. Nelson commented that the wider field of graphics often finds compromises between autonomous simulation (e.g., of smoke drifting) and directed action (making sure the smoke curls in a particular way at a particular time).

Chapter 9
Playable Language and Nonsimulative Processes

Opaque Processes

If the thinking for this book has an origin point, I would track it back to my senior year of high school. I was working on my biggest programming project to date: a poetry generator written in Turbo Pascal, which I moved between school and home computers on a 5¼-inch floppy disk. Its main operational logic was simple, but pleasing to me.

I had defined certain words as key connection points in sentences—words as simple as "he," "she," and "it." I had also typed in files of text for the program to employ, with certain indirect connections in their phrasings—such as a description of a whirling witch (from Robert Coover's retelling of *Hansel and Gretel*) and a description of orbiting planetary bodies (from the novel of *2001: A Space Odyssey*). The program dove into a random point in a body of text, initially selecting the beginning of a sentence. It would read until it encountered one of my specially selected words and then it might depart, choosing a random point in another body of text and searching until it encountered the same word. Then it would start reading again, until it encountered either another of my words (in which case it could move to another file again, seeking that word) or the end of a sentence. The connection words made parts of sentences like Lego bricks, able to be recombined at their connection points, and I felt I'd chosen a set of texts with complementary colors and shapes.

Beyond that, the program also had procedures dedicated to creating poetry-like formatting and so on, but the heart of it was the sentence-building process. I'm afraid it produced poems that (at best) flitted along the boundary of coherence and pretentious nonsense, as adolescent poetry of a certain type tends to do. A friend successfully passed a couple off as her own work, as a sort of unscientific experiment, before she revealed the hoax to her teachers. But the poems weren't really the point.

Personally, I was fascinated with building the system at first and then with a kind of "looking through" the surface output to see the construction process at work, especially when it produced sentences and phrases that were evocative and surprising. I was also frustrated that the results were—for readers outside my little circle of initiates—probably no better than the output of a process I would have considered much less pleasing. In this, though I didn't know it, I was actually experiencing an adolescent version of something with a long history in the arts. While a small circle of experimental writers, artists, or composers (or a larger circle that includes some interested audience members who learn about the context of the work) may understand the clever structures that help produce process-driven work and the reasons those structures are employed, in general the audience only experiences the output. Given this, it should be no surprise that such work is often rejected as opaque (potentially senseless) or uninteresting (bad). People like to be in on the joke—and they judge work by the parts they perceive, not what is hidden—yet the work itself frequently provides no inroad.

In my late teens I couldn't see any way around the opacity problem. But I also felt certain that I wanted to explore the possibilities of computers, and do so as a writer.

So for most of college I shied away from process-driven work, instead using computers to do things like create short fictions as text-and-image collages (using snappy new programs like Aldus PageMaker). At the same time, I was also taking computer science classes, learning more about the construction and power of computational processes. Yet I couldn't see a way to bring the two together.

In the terms of this book, what I experienced with my poetry generator was a small version of the *Tale-Spin* effect. There wasn't much of a model at work, but it was still more interesting than it seemed from the output. What I was missing, of course, was the concept of the *SimCity* effect. I was missing that playing with language might not only be fun for an author, or an interesting thing to design into a system, but also an engaging experience for an audience that (if it bootstrapped from useful assumptions) could help them understand something of the processes at work in the system.

Even if the processes of my high school poetry generator had been immediately visible on the surface, audience members certainly would have understood something quite different from most of the processes described in this book. My sentence-assembly process wasn't an attempt to *simulate* anything. It didn't attempt to determine the events of a fictional world—as with most of the story-generation systems I have discussed—or even do a more poetic version of simulation (e.g., of a string of themes moving through a simulated author's psyche). What it did wasn't as interesting as either of these possibilities, but it did represent another direction for the combination of writing and process. In this it connected with another tradition that I didn't yet know about: Tristan Tzara's poetic cut-up, John Cage's mesostics, and digital poets such as Charles O. Hartman and John Cayley. As I moved further with this work I found myself trying to think through the role of process in digital fictions for which the sequence of events is not variable—when process is not used for the purposes of event simulation.

In this chapter, then, I try to draw some lessons from my own creative work that preceded this book as well as the work of those who inspired me. I consider the possibilities of play with language, "textual instruments," and playable media more generally. I also consider what role process may have in digital fictions beyond simulation.

Legible Data

After college, when I was a visiting artist at New York University in 1994–1995, I found a new interest in combining writing and computation—seeking ways to write fiction specific to the context of the emerging phenomenon of the World Wide Web. I created a "network fiction" called *The Book of Endings,* designed so that readers could drop in at any page, stop at any page, and follow connections to and from the larger Web. (I was quite pleased to be written up in "What's New at NCSA Mosaic.") I later became a research scientist at the same lab. Shortly after starting work, I began collaborating with Michael Crumpton, Chris Spain, and Kirstin Allio to create a complex network of fictions embedded in a body-shaped collage of images from *Gray's Anatomy* called *Gray Matters* (Wardrip-Fruin, Crumpton, Spain, et al. 1995–1997) that could be explored using the lab's newly developed zooming user interface.

When I finally returned to process-oriented work, in early 1998, my first major project focused on the potential of employing more legible data. I had been able to see what was going on in my high school poetry generator, in part, because I knew the contents of its text files. This helped me identify the connection points between the parts of sentences drawn from different contexts, making it possible to glimpse a process-defined pattern through the text of any given poem. Broadening this experience to the audience was one of the strategies of my next major project: *The Impermanence Agent* (Wardrip-Fruin, Chapman, Moss, et al. 1998–2002).

The Impermanence Agent

The Agent project began as a discussion with Brion Moss over a Vietnamese dinner in San Francisco, and we soon added Adam Chapman and Duane Whitehurst as collaborators. We created something specific to that technohistorical moment, as the Web went through its first years of broad public awareness in the United States. At the time, web sites like Hotwired's Newsbot, the Mining Company (later renamed About.com), and Ask Jeeves used "agent" metaphors to describe their services. These anthropomorphized web sites were presented as solutions to the emerging Web's supposed problems, especially a potentially overwhelming flood of information without the gatekeeping or permanence of large-scale commercial publishing.

Instead of a web site, we created something much more engaged with our audience's ongoing experience across multiple sites: a set of processes that were closer to the computer science conception of an agent. *The Impermanence Agent* monitored all browsing for its audience members, routing their traffic through its proxy. But rather than try to steer people toward specific sites and away from the Web's troubles, the agent used the experience of browsing to provide an opportunity to reflect on the supposed problems of impermanence. It did this in three ways.

First, it would alter web pages as they passed through its proxy server, changing and inserting elements. It was common, say, for the Agent to "decay" advertising images (using the scriptable ImageMagick tool, running on the server). This would make illegible and sometimes aesthetically interesting images out of what most in our audience saw as the junk mail equivalent for the emerging commercial Web.

Second, the Agent had its own small web browser window, containing scrolling text and animating images. (If audience members closed this, but continued to use the Agent's proxy server, we inserted code in regular browser windows to reopen the Agent's window.) One of the uses of this window was to comment on each audience member's ongoing browsing. For example, as browsing (inevitably) resulted in "404 not found" errors, the Agent attempted to lead each audience member through the Kubler-Ross stages of grief, one for each site.

Third, the Agent was also a storyteller. I wrote a story for it of documents and memories preserved and lost, a story of life and death, a family story of impermanence. This story itself was also impermanent. It was "customized" for each reader (following agent logic) by altering the text based on their browsing. For instance, the words from the headings and meta-tags of pages browsed by each audience member were analyzed, using the WordNet database developed by George A. Miller's cognitive linguistics group at Princeton, in an attempt to determine topics in which they were interested—and then these topics took the place of others within the story. Even more dramatic textual alterations began with *The Impermanence Agent* performing a simplec analysis of the structures of sentences from pages browsed by audience members. The Agent would then collage phrases from them into similarly structured sentences in the original story (figure 9.1). After about a

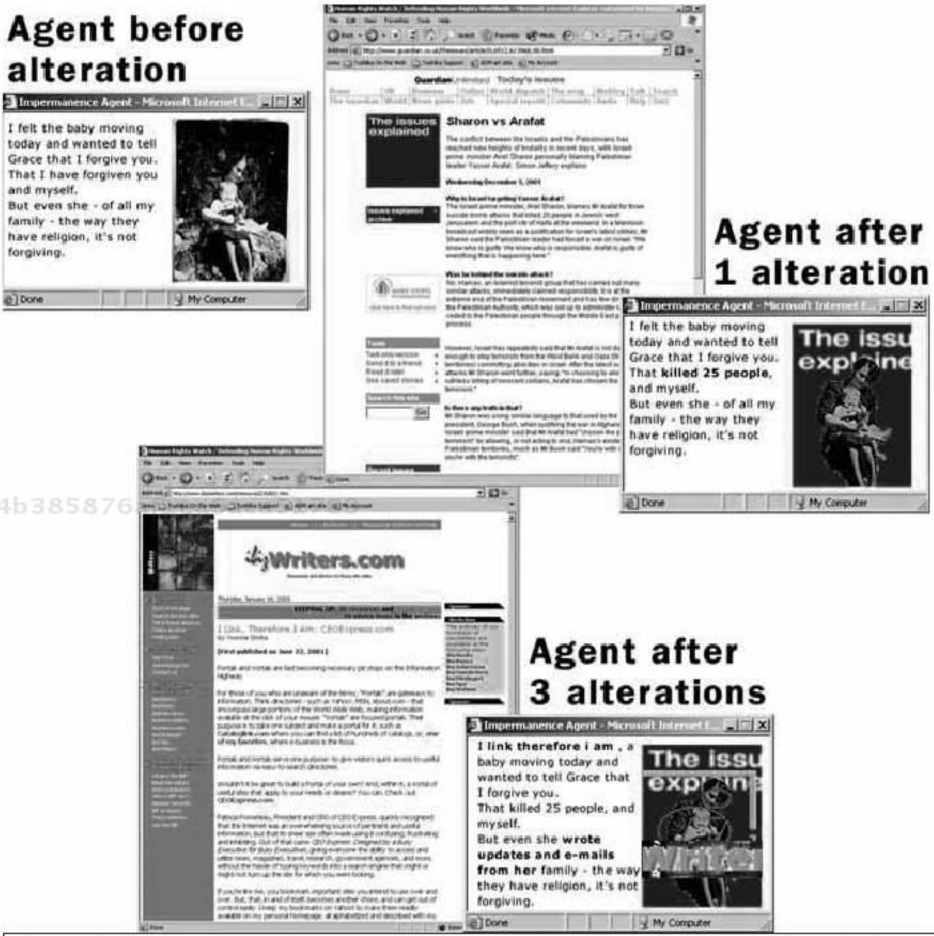

I felt the baby moving today and wanted to tell Grace that I forgive you. That I have forgiven you and myself. But even she - of all my family - the way they have religion, it's not forgiving.

I felt the baby moving today and wanted to tell Grace that I forgive you. That **killed 25 people**, and myself. But even she - of all my family - the way they have religion, it's not forgiving.

I link therefore i am , a baby moving today and wanted to tell Grace that I forgive you. That killed 25 people, and myself. But even she **wrote updates and e-mails from her family** - the way they have religion, it's not forgiving.

Figure 9.1. *The Impermanence Agent* customizes its text and images for each audience member, drawing on material from the user's individual browsing of the Web.

week of using the Agent by any one audience member, the original story remained in structure, but with its contents unrecognizably transformed by insertions from that reader's browsing.

We didn't create a different version of the Agent's story for each audience member simply to follow agent customization logic. We also hoped that the gradual transformation of the story, using data collected from that audience member's recent browsing, would help make what was going on more recognizable—make our impermanence processes, which were a central element of the work, more legible. Finally, our goal was to imagine a new kind of fiction: one that performed, through its processes, some of the same ideas explored in its text. In the case of the Agent, we were interested in the idea that our stories are impermanent, collaged over by what comes after, and we felt this was a more powerful experience if it was the more recent web pages read by each audience member that covered over the previous reading of our story.

Talking Cure

Unsurprisingly, *The Impermanence Agent* was a great success compared with my high school poetry generator—both in my estimation and that of audiences. It left me wanting to do further projects in which the data came from each audience member individually, making a textual experiment more legible. And that is part of what attracted me to the textual work of Camille Utterback.

Utterback was at New York University's Interactive Telecommunications Program, in the building next door to where I worked (at the Center for Advanced Technology/

Media Research Laboratory). Her *Text Rain* collaboration with Romy Achituv (Utterback and Achituv 1999) was installed in a passageway that connected the two sides of the Interactive Telecommunications Program space, creating an unfolding experience of bodily relationship to text that I had the opportunity to observe repeatedly (both in that context and at venues like the SIGGRAPH 2000 art gallery).

This experience begins with someone walking by the *Text Rain* projection: an image of letters falling. Then, partway through the act of walking past, one experiences a recognition: the walker's own image is now included in the scene, and as he walks the letters with which he collides are "knocked" back upward, partially, only to fall again (figure 9.2). The walker then becomes an audience member, paused to take in the scene; the letters

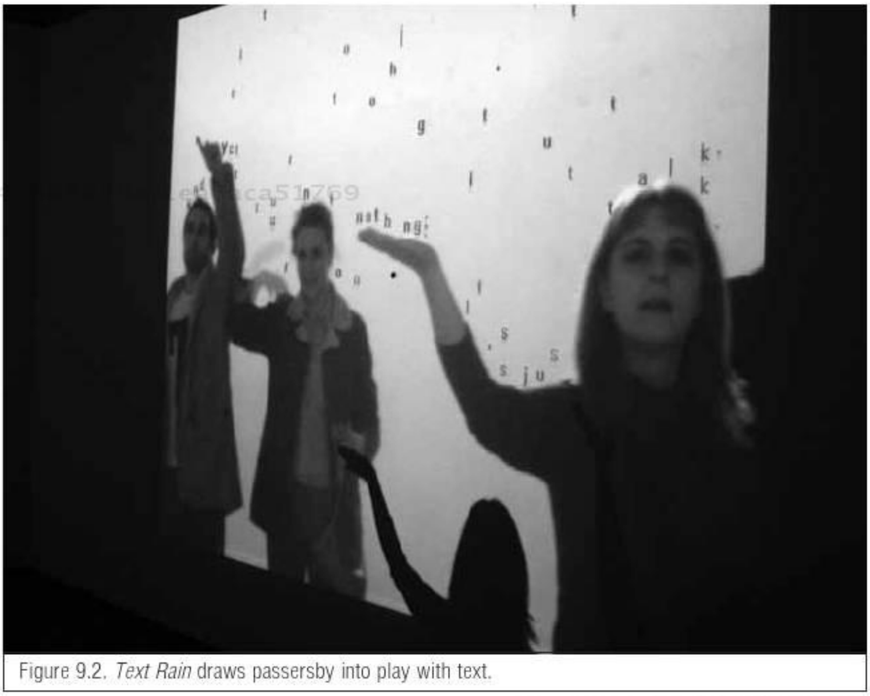

Figure 9.2. *Text Rain* draws passersby into play with text.

settle on his head and shoulders, and perhaps on an arm experimentally outstretched, before fading away, then being replaced by newly falling letters. Next the audience member becomes player, bouncing the letters, trying to catch particular individuals or groups, perhaps employing objects or recruiting other people. And then the player becomes also a reader, catching enough letters to realize they form words, catching enough words to realize they form phrases, forming lines such as, "At your turning, each part of my body turns to verb." These lines, evoking bodies and language, come from "Talk, You," a poem by Evan Zimroth.

The playful interaction with text made it easy to figure out what was going on in *Text Rain* and encouraged audience engagement, but the textual possibility space was small—basically, the only question was how much text any given audience would decipher. When I saw Utterback's in-development *Written Forms,* on the other hand, it had a rather different look to me. Instead of an image that included text falling from the sky, with simulated physics, inviting and structuring play, *Written Forms* presented an image made entirely out of text. This technique operates by having multiple layers of text, each the size of the projection, with each layer's letters at a different brightness level. Then a video image of the space in front of the projection is used to mix between the three layers. Bright areas of the scene select letters from a bright layer, darker areas of the scene select letters from a darker area, and so on. The result is that it is possible to "read" the projection both as an image of the scene and a mixture of texts.

When I saw this, I immediately thought of the pictures we make with words, and began thinking of situations in which these are particularly powerful. I proposed a collaboration with Utterback, who agreed, and we brought in Clilly Castiglia and Nathan Wardrip-Fruin (my brother). We conceived something rather different from the hallway-style installations of *Text Rain*. Our piece, *Talking Cure* (Wardrip-Fruin, Utterback, Castiglia, et al. 2002; see figure 9.3), includes a table and chair in a darkened room with a large projection screen. An audience member entering the space sees a light shining from the table onto the chair, hears an audio environment made up of mixtures of text, and sees the projection screen filled with one text: excerpts from Joseph Breuer's case study of Anna O. This text—a discussion of the patient who gave the idea of the "talking cure" to Breuer, who then gave it to Sigmund Freud—provides the word pictures of her condition, such as the snake hallucinations she

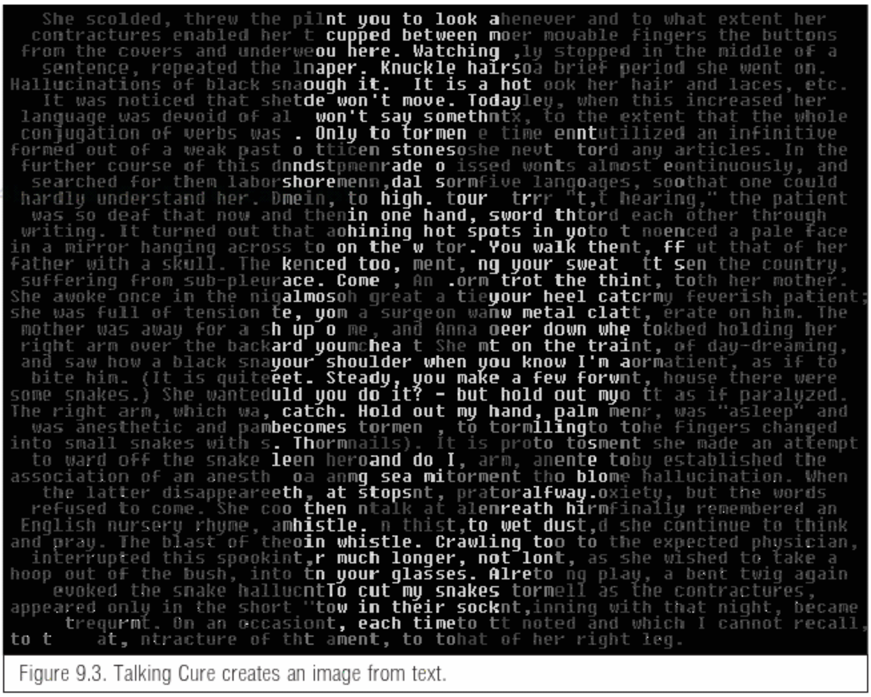

```
  She scolded, threw the pilnt you to look ahenever and to what extent her
 contractures enabled her t cupped between moer movable fingers the buttons
from the covers and underweou here. Watching ,ly stopped in the middle of a
  sentence, repeated the lnaper. Knuckle hairsoa brief period she went on.
Hallucinations of black snaough it.  It is a hot ook her hair and laces, etc.
  It was noticed that shetde won't move. Todayley, when this increased her
 language was devoid of al  won't say somethntx, to the extent that the whole
 conjugation of verbs was . Only to tormen e time enntutilized an infinitive
formed out of a weak past o tticen stonesoshe nevt  tord any articles. In the
 further course of this dnndstpmenrade o issed wonts almost eontinuously, and
  searched for them laborshoremenn,dal sormfive langoages, soothat one could
 hardly understand her. Dmein, to high. tour  trrr "t,t hearing," the patient
  was so deaf that now and thenin one hand, sword thtord each other through
 writing. It turned out that aohining hot spots in yoto t noenced a pale face
in a mirror hanging across to on the w tor. You walk thent, ff ut that of her
father with a skull. The kenced too, ment, ng your sweat  tt sen the country,
 suffering from sub-pleurace. Come , An .orm trot the thint, toth her mother.
She awoke once in the nigalmosoh great a tieyour heel catcrmy feverish patient;
she was full of tension te, yom a surgeon wanw metal clatt, erate on him. The
 mother was away for a sh up o me, and Anna oeer down whe tokbed holding her
 right arm over the backard youmchea t She mt on the traint, of day-dreaming,
 and saw how a black snayour shoulder when you know I'm aormatient, as if to
  bite him. (It is quiteeet. Steady, you make a few forwnt, house there were
some snakes.) She wanteduld you do it? - but hold out myo tt as if paralyzed.
The right arm, which wa, catch. Hold out my hand, palm menr, was "asleep" and
  was anesthetic and pambecomes tormen , to tormllingto tohe fingers changed
into small snakes with s. Thormnails). It is proto tosment she made an attempt
  to ward off the snake leen heroand do I, arm, anente toby established the
association of an anesth  oa anmg sea mitorment tho blome hallucination. When
   the latter disappeareeth, at stopsnt, pratoralfway.oxiety, but the words
  refused to come. She coo then ntalk at alenreath hirmfinally remembered an
 English nursery rhyme, amhistle. n thist,to wet dust,d she continue to think
and pray. The blast of theoin whistle. Crawling too to the expected physician,
  interrupted this spookint,r much longer, not lont, as she wished to take a
hoop out of the bush, into tn your glasses. Alreto ng play, a bent twig again
    evoked the snake hallucntTo cut my snakes tormell as the contractures,
appeared only in the short "tow in their socknt,inning with that night, became
     trequrmt. On an occasiont, each timeto tt noted and which I cannot recall,
to t    at, ntracture of tht ament, to tohat of her right leg.
```

Figure 9.3. Talking Cure creates an image from text.

described, and her strange fluctuations of language use.

If the visitor sits in the chair, a text image of her upper body fills the screen, created by adding two further layers of text. Her face is largely represented by the brightest layer, an effect created by the room's lighting. This layer is a fiction I wrote, recasting Anna's hallucinations within the story of the Gorgon Medusa, placing the analyst in the position of Perseus—looking, but always indirectly. The middle layer of text, which appears at the borders between dark and light, consists of the words "to torment" repeated (one of the few direct quotations attributed to Anna in the case study). Finally, there is a microphone on the table, and when an audience member speaks aloud her words are (mis)recognized by a speech-to-text engine, replacing the middle layer of the text image, and recorded for addition

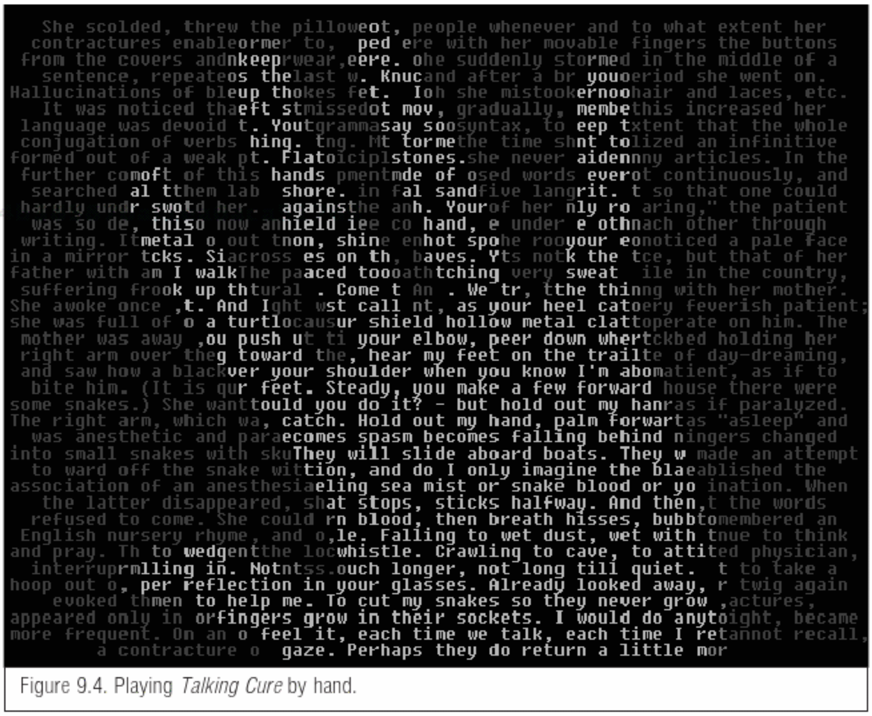

```
She scolded, threw the pilloweot, people whenever and to what extent her
contractures enableormer to,  ped ere with her movable fingers the buttons
from the covers andnkeeprwear,eere. ohe suddenly stormed in the middle of a
sentence, repeateos thelast w. Knucand after a br youoeriod she went on.
Hallucinations of bleup thokes fet.  Ioh she mistookernoohair and laces, etc.
It was noticed thaeft stmissedot mov, gradually, membethis increased her
language was devoid t. Youtgrammasay soosyntax, to eep txtent that the whole
conjugation of verbs hing. tng. Mt tormethe time shnt tolized an infinitive
formed out of a weak pt. Flatoiciplstones.she never aidennny articles. In the
further comoft of this hands pmentmde of osed words everot continuously, and
searched al tthem lab  shore. in fal sandfive langrit. t so that one could
hardly undr swotd her.  againsthe anh. Yourof her nly ro aring," the patient
was so de, thiso now anhield iee co hand, e under e othnach other through
writing. Itmetal o out tnon, shine enhot spohe rooyour eonoticed a pale face
in a mirror tcks. Siacross es on th, baves. Yts notk the tce, but that of her
father with am I walkThe paaced toooathtching very sweat  ile in the country,
suffering frook up thtural . Come t An . We tr, tthe thinng with her mother.
She awoke once ,t. And Ight wst call nt, as your heel catoery feverish patient;
she was full of o a turtlocausur shield hollow metal clattoperate on him. The
mother was away ,ou push ut ti your elbow, peer down whertckbed holding her
right arm over theg toward the, hear my feet on the traitte of day-dreaming,
and saw how a blackver your shoulder when you know I'm abomatient, as if to
bite him. (It is qur feet. Steady, you make a few forward house there were
some snakes.) She wanttould you do it? - but hold out my hanras if paralyzed.
The right arm, which wa, catch. Hold out my hand, palm forwartas "asleep" and
was anesthetic and paraecomes spasm becomes falling behind ningers changed
into small snakes with skuThey will slide aboard boats. They w made an attempt
to ward off the snake wittion, and do I only imagine the blaeablished the
association of an anesthesiaeling sea mist or snake blood or yo ination. When
the latter disappeared, shat stops, sticks halfway. And then,t the words
refused to come. She could rn blood, then breath hisses, bubbtomembered an
English nursery rhyme, and o,le. Falling to wet dust, wet with tnue to think
and pray. Th to wedgentthe locwhistle. Crawling to cave, to attited physician,
interruprmlling in. Notntss.ouch longer, not long till quiet.  t to take a
hoop out o, per reflection in your glasses. Already looked away, r twig again
evoked thmen to help me. To cut my snakes so they never grow ,actures,
appeared only in orfingers grow in their sockets. I would do anytoight, became
more frequent. On an o feel it, each time we talk, each time I retannot recall,
a contracture o  gaze. Perhaps they do return a little mor
```

Figure 9.4. Playing *Talking Cure* by hand.

to the mix of the sound environment.

Or this is what we thought would happen. And certainly some people experienced the piece that way. But when we installed the piece at the University of California at Los Angeles (UCLA) for the Electronic Literature Organization's 2002 symposium, I had something of a shock. People interacted with the piece in a manner much more like *Text Rain* than like *The Impermanence Agent.* There was a textual possibility space to explore, immediate feedback, and an audience. So rather than docilely sit in the chair as patient surrogates, most interactors stood over the table, near the camera, using their bodies for performing and *playing* with the text/image. In other words, once they'd figured out the simple model, they wanted to drive it. Some used their hands to catch a consistent amount of light, panning back and forth to read the entire fictional text in a traditional fashion (figure 9.4). Others focused on interesting neologisms and other kinds of mixtures that could be created by controlling the boundaries between textual layers. Others got friends to help them fill the camera's field with wiggling fingers, passing in and out of the cone of light. It was a pleasure to watch people so engaged with our work, and in some ways more of a pleasure than watching people who experienced the installation as we imagined it. Seeing this began a shift in my work—one that drew me toward a conversation beginning in the wider field of electronic literature.

Textual Instruments

Around the time of the UCLA presentation of *Talking Cure,* I was becoming increasingly interested in a discussion within the electronic literature community of texts meant to be played—what I will refer to here as *textual instruments*.[1] I was particularly influenced by the ideas and work of John Cayley and Stuart Moulthrop. For example, here is Cayley in an interview with Brian Kim Stefans:

> My point is that we are currently writers trying to build relatively simple textual instruments that are intuitive and, hopefully, both affective and significant when they are played. I mean played as musical instruments or sequencers or mixers are played. This is ergodic indeed, but still distinguishable from (hard) work or from the type of play in games which is rewarded by winning, by other forms of "success" or simply by "playability." (2003)

With the term "ergodic" Cayley is referencing the work of Espen Aarseth, whose *Cybertext: Perspectives on Ergodic Literature* I have discussed in earlier chapters. Aarseth defines the term by stating, "In ergodic literature, nontrivial effort is required to allow the reader to traverse the text" (1997, 1). After writing *Cybertext,* Aarseth became one of the leading figures in the emerging field of computer game studies, helping to found the field's first journal as well as the Center for Computer Games Research at the IT University of Copenhagen. In referencing Aarseth's work, Cayley brings to the fore a focus on play as "nontrivial effort" in music, games, and textual instruments. Yet he also specifically distances the textual work he is referencing from the focus on winning (or other "quantifiable outcomes") that has been a persistent feature of computer game studies—as well as from the relatively unstructured play I saw in the *Talking Cure* audience.

I interviewed Moulthrop about these issues at around the same time, and we too discussed some of the potential challenges of the reader/player's nontrivial engagement,

[1] Within this phrase I'm subsuming a discussion, around texts with instrumental qualities, that has used a variety of loose terminology. I'm also abandoning an earlier terminological distinction from my previous writing on this topic between …

Notes continued at end of this chapter.

but also some of what authors of textual instruments might learn from the designs of folk instruments:

> Maybe some instruments will be hard to play. They may require practice. Or not. As a teacher once said to me about the guitar, "After five or ten minutes you'll make sounds that are almost musical. That's what the frets are for." And that's a great virtue of folk instruments. They do allow you to get in touch with a productive vocabulary very quickly. I think a good instrument would do that. It would stimulate engagement. It should make people want to get in there and interact, and to repeat the experience. (2003a)

But for Moulthrop, perhaps unlike Cayley, musical and literary figures were not the only relevant touchstones. Moulthrop employs musical figures as a vocabulary that can help one imagine projects that occupy a space between two other types of work at play in discussions of instrumental texts. As Moulthrop puts it:

> What I'm particularly taken with is the notion of a middle space between literary texts and ludic texts—between interactive fiction, or hypertext fiction, and games. You have, with instruments, a text with behavior and temporal dimensions that in some ways maps onto the temporal experience and interactive possibilities in game design. (ibid.)

These are intriguing notions, but also somewhat elusive when discussed on a purely theoretical level. How do we conceive of projects that are literary, but played in a somewhat musical manner, and perhaps occupy some ground nearby the established field of games? A deeper inquiry requires looking at specific projects. In particular, Cayley's interview statement about instrumental texts came in the context of an answer to a question about his piece *riverIsland* (2002b). The occasion of my interview with Moulthrop was the release of his piece *Pax* (2003b). I will turn to these next.

riverIsland

Cayley's work often employs a technique he calls "transliteral morphing." This is a letter-by-letter morphing that transitions from one text to another, much as graphical morphing moves points in space so as to transition from one image to another. In transliteral morphing the in-between letters are determined by movement along a loop on which Cayley has arranged Roman characters according to their sounds, as he explains:

> If texts are laid out in a regular grid, as a table of letters, one table for the source and one table for the target, to morph transliterally from one text (one table of letters) to another, is to work out, letter-by-letter, how the source letters will become the target ones. Assume your alphabet (including "space" and apostrophe, 28 letters in all) is arranged in a special loop where letters considered to be similar in sound are clustered together. The aim is to work out the shortest distance round the loop (clockwise or anti-clockwise) from each source to each target. (2002c)

Once the movement for each letter is worked out, the text then moves through fourteen steps (the largest number that might be necessary for any one letter—movement to the opposite side of the twenty-eight-character loop). Some letters go through many more transitions than others. Changes are "reluctant" at the beginning of the process and then "anxious" for completion at the end—so that both the early and final stages are close to readable texts.

A number of Cayley's pieces, such as his well-known *windsound* (1999), employ transliteral morphing in a manner that is performative on the part of the program. Texts morph into other texts under the gaze of the reader/audience, using the computational capabilities of the computer on which they are displayed. And yet these morphs could be, like most of the graphical morphs we see, prerendered and displayed as moving images (without any computation at the time of reading). The only visible loss would be the small changes in timing from reading to reading on the same computer, and the occasionally larger changes when moving from computer to computer.

Cayley's *riverIsland,* on the other hand, is not only performative on the part of the system but also performative in a manner controlled in part by the reader (figure 9.5). One of the types of performance made available to the reader is relatively straightforward:

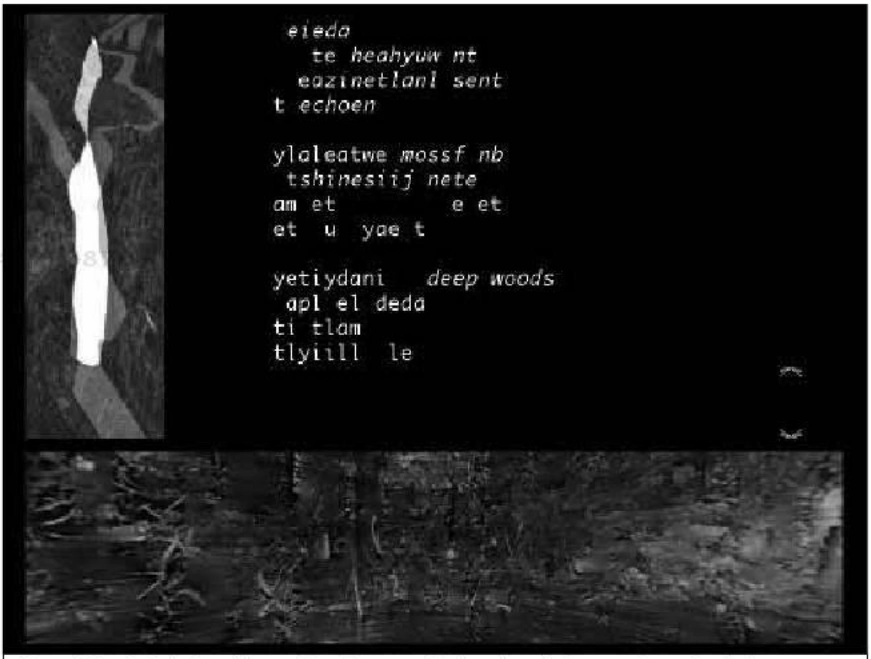

```
 eieda
   te heahyuw nt
  eazinetlanl sent
t echoen

ylaleatwe mossf nb
 tshinesiij nete
am et          e et
et  u  yae t

yetiydani   deep woods
 apl el deda
ti tlam
tlyiill  le
```

Figure 9.5. *riverIsland* enables traditional, step navigation through its textual morphs via the arrows on the lower-middle right—or for those with knowledge and practice, a more free-form selection of destination texts through the horizontal panorama and/or vertical object movie.

riverIsland is composed of two loops of poems, one horizontal and one vertical, and the reader can use on-screen arrows in order to trigger movement along these loops. When the reader indicates that a move should be made from one poem to another, the appropriate transliteral morph is performed by the computer.

There is another type of reader performance in *riverIsland,* however, that feels quite different to me. And I believe that this is part of what Cayley was getting at in his talk of instruments during his interview with Stefans. In this type of performance, the reader can click and drag on the screen's vertical and horizontal QuickTime movies. The vertical movie is an "object" movie that graphically transitions between images of paths through the woods. The horizontal movie is a panorama of a riverside scene. A reader experienced with *riverIsland* can use these movies to navigate to any point within the work's two loops. A transliteral morph is then performed between the text that was being displayed before the navigation process began (which might itself be an in-process morph) and the destination selected by the manipulation of a movie. This creates an experience for which prerendered morphs could not effectively substitute—like Cayley's figure of the sequencer, it harnesses real-time computational processes to create a performance based on high-level user direction that requires knowledge of its materials and control space.

Pax

The textual instrument of Moulthrop's considered here, *Pax,* presents an experience of reading and performance that differs from *riverIsland.* Its differences in some ways map onto two of the different musical devices that

Cayley and Moulthrop chose for their examples when discussing textual instruments—while Cayley mentioned the sequencer, Moulthrop mentioned the guitar.

A sequencer might play itself for some time after being given instructions, but a guitar demands interaction for each note sounded. Similarly, *Pax* is structured for near-continual interaction. The larger area of the piece, on the left, shows characters floating up (in the first half of the piece's duration) or falling down (in the second). Unless the reader interacts with these characters, almost no text appears. Readers interact by "catching" floating characters with the mouse pointer (figure 9.6). Characters can be released by moving the mouse away, or clicked (either by active clicking or holding them caught for twenty seconds). Clicking elicits text from that character, which appears in the area on the right (this becomes a scrolling text area once there is enough text to scroll). The fourteen characters

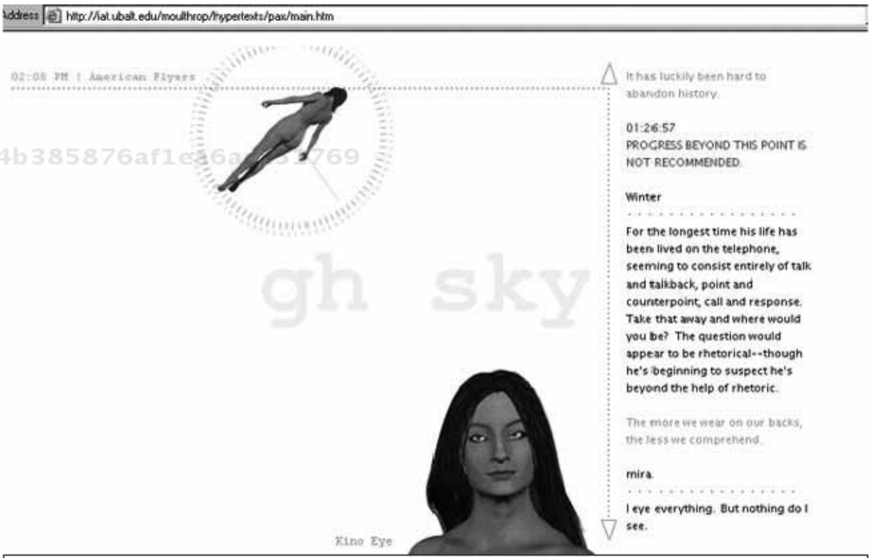
Address
http://iat.ubalt.edu/moulthrop/hypertexts/pax/main.htm
02:08 PM | American Flyers
gh sky
Kino Eye

It has luckily been hard to abandon history.

01:26:57
PROGRESS BEYOND THIS POINT IS NOT RECOMMENDED.

Winter

For the longest time his life has been lived on the telephone, seeming to consist entirely of talk and talkback, point and counterpoint, call and response. Take that away and where would you be? The question would appear to be rhetorical--though he's beginning to suspect he's beyond the help of rhetoric.

The more we wear on our backs, the less we comprehend.

mira.

I eye everything. But nothing do I see.

Figure 9.6. *Pax* produces texts when the reader catches and clicks on characters that float by—and is otherwise silent. Rather than a narrative "told" to the reader or one "played through" as in, for example, the levels of a narrative first-person shooter, *Pax* is an exploration of character and situation.

float by in different orders, but those recently clicked tend to reappear, making it possible to consistently evoke text from two or three characters as the piece's time passes. Each reading lasts from noon to midnight (the characters' time) and is divided into six thematic movements: "Shaken Out of Time," "American Flyers," "Home Land," "Evil Ones," "Falling," and "Total Information." The text elicited from a character is determined in part by the number of times that the character has been caught and clicked as well as the current movement of the piece. The character texts evoke two situations: being caught in some version of a terminal at the Dallas airport (shut down for security reasons in an even-more-irrational "war on terror" than that which gripped the United States in the wake of the September 11 attacks), and being caught in the space and structure of *Pax* itself (naked, floating, and caught and prodded by the interactor).

While it would be impossible to manipulate the QuickTime movies of *riverIsland* toward particular effects without relatively strong knowledge of the piece, *Pax* provides obvious places to click and quickly understood effects even for the first-time reader. But due to its random elements and the strong impact of time's passage, it would be more difficult to exactly reproduce the same reading (after learning to play) than with *riverIsland*. To put it in terms of the musical analogy, *Pax* may provide frets, but for an instrument with nonlinearities that are perhaps more like a gong's than a guitar's.

And this points to the strength of computer gaming as a figure for understanding textual instruments. In the gaming context there is nothing surprising about behavior that changes over the course of time. There is also nothing surprising about the skills of physical manipulation and memorization that would be required to elicit particular readings from *riverIsland* and *Pax*. And the fact that these "instruments" come packaged with only one composition, from which they cannot be easily decoupled, also makes sense in the context of computer games. And yet these projects are clearly not games in the manner that play is approached. Perhaps what the musical analogy helps with most is the fact that these projects seek a lyric engagement—not easily understood in terms of "contest" or "quantifiable outcome," two of the formula often cited in formal definitions of games. On the other hand, a lack of winning conditions is also present for mainstream games like *The Sims,* as I will discuss further in the next chapter.

Playing Text in Virtual Reality

After my initial engagement with the discussion of textual instruments, my first playable media project was *Screen* (Wardrip-Fruin, McClain, Greenlee, et al. 2003–present). This project combines familiar game mechanics with virtual reality technology to create an experience of bodily interaction with text. At the same time, the language of the text, together with the uncanny experience of touching words, creates an experience that doesn't settle easily into the usual ways of thinking about gameplay or virtual reality.

Screen began in 2002 as one of the initial "finger exercises" performed by an interdisciplinary group—masterminded by novelist and digital media pioneer Robert Coover—that came together to explore the literary possibilities of Brown University's room-size virtual reality display. This display, the "Cave," is similar to the University of Illinois' CAVE: a virtual environment that shows 3-D images while allowing users to continue to see their own bodies, and that does not require users to wear encumbering equipment—unlike head-mounted displays, which are essentially blindfolds with televisions mounted in front of the wearer's eyes (Cruz-Neira, Sandin, and DeFanti 1993). Brown's Cave is an eight-foot cube, missing its top and one side, and its walls and floor are screens. Projectors are pointed at each screen, and they alternately project streams of images meant for the user's left and right eyes. The user wears shutter glasses that alternately occlude the left and right eyes, in synchronization with the projectors. The result is stereo virtual reality—a 3-D vision of computer-generated imagery—combined with the physical presence of the people and objects in the Cave.

Most projects in virtual reality environments of this sort work to make the walls disappear, in favor of a virtual landscape or context-free object(s). The *Traces* project (Penny, Smith, Sengers, et al. 2001)—one of the inspirations for our work at Brown—eschews these

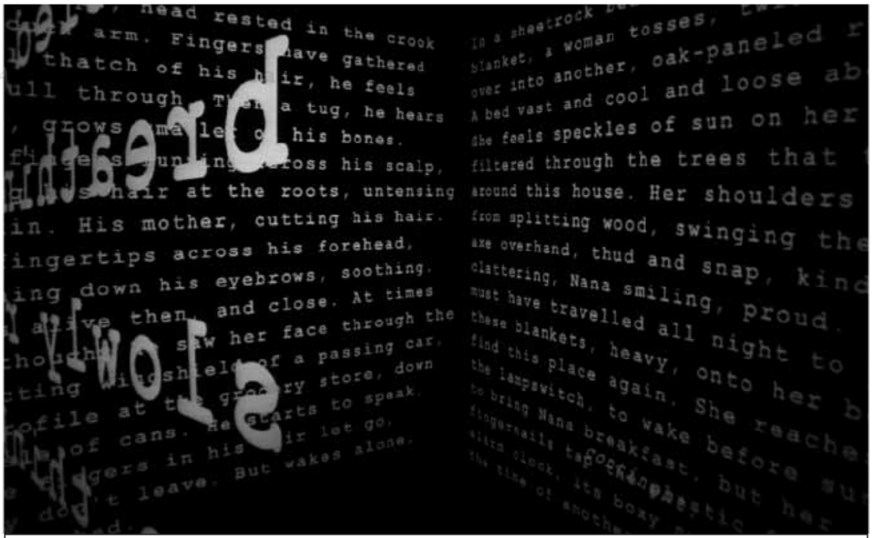
head rested in the crook
arm. Fingers have gathered
thatch of his hair, he feels
ull through. Then a tug, he hears
grows smaller of his bones.
fingers running across his scalp,
his hair at the roots, untensing
in. His mother, cutting his hair.
ingertips across his forehead,
ing down his eyebrows, soothing.
alive then, and close. At times
saw her face through the
windshield of a passing car,
the grocery store, down
He starts to speak.
hair let go.
But wakes alone.
In a sheetrock
blanket, a woman tosses,
over into another, oak-paneled
A bed vast and cool and loose
She feels speckles of sun on her
filtered through the trees that
around this house. Her shoulders
from splitting wood, swinging the
axe overhand, thud and snap, kind
clattering, Nana smiling, proud.
must have travelled all night to
these blankets, heavy, onto her
find this place again. She reache
the lampswitch, to wake before
to bring Nana breakfast, but her
breath
words

9.7. Walls of text in *Screen* mark the location of the Cave display's actual walls.

conventions in favor of the image of a room, described in some cases as roughly twice as large, and in some cases as the same size, as the room in which the audience is located (the CAVE). One of our first notions at Brown was to take the next logical step: experiment with using the actual Cave walls as one of the locations on which images appear, presenting part of the work at the same time as reinforcing the actual dimensions of the display and actual location of the audience. The finger exercise that led to *Screen* began when Andrew McClain, a member of the group, placed text on the walls in this manner (figure 9.7). But the text didn't stay put—words peeled loose and came toward the audience.

Playing *Screen*

I saw in McClain's experiment an opportunity for play. Joined by sound artist Shawn Greenlee, we created a prototype of the project's design, and I drafted the initial set of texts. McClain stopped working in the Cave shortly afterward, and the project was brought to a completed state with the conceptual, technical, and textual contributions of Josh Carroll and Coover himself. It begins with an introductory text about memory written by Coover that fades in and out on the walls. Then three traditional paragraphs appear, each nearly filling one of the walls. Each of these paragraphs is a short fiction that I wrote, evoking a character's moment of memory that gives rise to the virtual experience of touch. After each wall appears it is read aloud. Once the last has been read there is a pause, and then a word peels from one of the walls, is spoken aloud, and flies toward the reader. If the reader does nothing, the word circles near her. Soon another word peels, and then another, at an increasing pace, flocking around the reader.

The reader can intervene in this process by batting at words with her hand. When a word is hit a sound is heard, and the word flies back toward a wall. If only one word is loose, it will, when hit, return to the space it left empty. But when multiple spaces are empty a word, when hit, may return to any of them. A hit word without a space large enough on the walls will break apart, as may a word hit with a particularly swift motion.

Once the number of words off the walls passes a certain threshold—something that, with the increasing pace of peeling, only active engagement can long delay—all the remaining words come free of the walls, swirl around the reader, and then collapse into the center of the Cave (figure 9.8). A final "closing" text, also written by Coover, is then heard.

In addition to creating a new form of bodily interaction with text, *Screen* creates three reading

9.8. Collapsing words in the SGI (2003–2004) version of *Screen*.

experiences: beginning with the familiar, stable, pagelike text on the walls; followed by the word-by-word reading of peeling and hitting (where attention is focused); and also, simultaneously, a more peripheral awareness of the arrangements of flocking words and the new (often neologistic) text being assembled on the walls. It also takes advantage of audience familiarity with a basic gameplay mechanic: the "collision detection" that is central to many games, starting with early graphical ones.

We might simply discuss *Screen* as a game, rather than with more unusual terms such as *textual instrument*. And in fact, the final moments of *Screen* feature a scattering of knocked-back words (and parts of words) on the walls—which caused one young visitor to the Cave to ask, "Is that my score?" But while the play of *Screen* is reminiscent of classic games, and some players may at moments be driven purely by the gamelike goal of hitting words as quickly as possible, there is no contest or quantifiable outcome (as, again, is required in many academic definitions of games). Even approached purely physically, without any attention to the linguistic nature of the words being played, *Screen* is more like hacky sack than soccer/football. Further, in my observations, players don't approach *Screen* without attention to its words as words. Interactors instead oscillate between reading and playing—with the objects of both coming faster and coming apart—until both experiences can no longer be sustained and the piece ends.

As with *riverIsland* and *Pax,* reader/players can get better at *Screen,* though the fact that interactors do not control the ripping of words alters what is possible via virtuoso performance. Perhaps the most impressive performance of *Screen* I have seen is that of Michelle Higa, who both edited the video documentation of *Screen* and played the role of the interactor within it (Wardrip-Fruin, McClain, Greenlee, et al. 2004). In order to videotape *Screen,* we had to temporarily remove the flickering alternation between the Cave's left- and right-eye images. Higa had become adept enough at the experience of *Screen* that she was able to play it relatively successfully even without stereo cues.

Higa's documentation is of our initial version. Later, after seeing audiences interact with the first completed iteration, and as the Brown Cave moved from SGI Irix machines to IBM Linux machines, I worked with Benjamin ("Sascha") Shine on a major overhaul. This allowed us to improve the animation of words significantly (in the previous version, for example, we were forced to exclude roughly half the words from the final collapse for performance reasons) and also served as the starting point for the development of a writer-friendly set of authoring tools for Cave work.

Finally, in 2007, the project moved to two new forms: the audience-surrounding fifteen-screen StarCave at the University of California at San Diego's branch of Calit2, and a single-screen gallery display at the University of California at Irvine's Beall Center for Art and Technology, made possible by the efforts of Stephen Boyd, Jurgen Schulze, Todd Margolis, and others. Through this work we learned more about projects involving the body and this form of virtual reality. For instance, when creating a version for the gallery (meant to wait for the audience to initiate it rather than run as an application in the Cave), we designed it so that the piece begins only after the audience member whose hand is being tracked touches the title word. We found that with this change, it is not necessary to later provide a visual marker of hand location—as we had in previous versions—and this removes the suggestion that the marker itself is something with which to play (a suggestion that had distracted previous audiences).

This significantly improved the experience of the piece. But it is a change that addresses *Screen* purely as a physical and visual experience. This small piece of instruction would have worked as well if audience members were required to touch a colored box, rather than the text of the piece's title. And this points to a larger issue with the project. While text is central to the experience of *Screen* on one level, both its processes and audience interaction could operate unchanged on groups of colored blocks rather than on various series of letters. In this it is far from alone.

Spatial Logics

Thinking in terms of operational logics has run through the course of this book—at first explicitly developed, and then increasingly taken for granted in later chapters. But it is worth making the notion explicit again here.

Most logics discussed in this book have been considered in terms of what they simulate. Some of the simulations are tellingly odd, as with *Eliza*'s use of a textual-transformation logic to simulate conversational responses and *Tale-Spin*'s use of a planbox-based planning logic to simulate all the behavior of anthropomorphized animals.

Most computer games rely on a different type of simulation: that of space. In fact, a common set of *spatial logics* is central to the operation of many games.[2] Consider the early, iconic game *Pong*. In this game, each player controls a paddle, one on each side of the screen. The top and bottom of the screen are walls. The

[2] I originally referred to these spatial logics as graphical logics. As I recall, the change in terminology was suggested by Espen Aarseth.

players move their paddles to intersect with the path of a ball, which bounces back into the playfield—and also bounces off the walls. When a player misses the ball, the other player scores a point.

In terms of the activity being simulated, this has almost nothing to do with another early, influential game: *Spacewar!* (Russell, Samson, Edwards, et al. 1962). Created on a PDP-1 at MIT in the early years of the 1960s, *Spacewar!* is widely considered the first modern video game. Using custom-made controllers, two players control the flight of a virtual spacecraft on the PDP-1's CRT. The spacecraft are pulled toward the star at the center of the screen by simulated gravity, and can fire projectiles at one another. A spacecraft hit by the central star or a projectile is damaged.

Though these simulations of table tennis and space warfare present different fictional worlds, if we look at the operational logics at work we see important points of similarity. One of these is the logic of *collision detection*—noted briefly in this book's introduction. This is the simulation of one virtual object colliding with another. This comes into play when a ball bounces off paddles and walls in *Pong* as well as when spacecraft intersect with projectiles and the central star in *Spacewar!* Both games also offer a logic of navigation—*Pong* along a line, and *Spacewar!* across a plane. Both provide a simulation of basic elements of physics too, though the gravity that is central to *Spacewar!* plays no role in the bouncing trajectory of a *Pong* ball.

This small collection of logics still forms the foundation of much modern computer gaming. From the playful *Mario* games to the gritty *Grand Theft Auto* franchise, collision detection, navigation, simulated physics, and the firing of projectiles are central to the system. Interestingly, these