Card Shark and Thespis:
exotic tools for hypertext narrative

Mark Bernstein

Eastgate Systems, Inc.
134 Main Street
Watertown MA 02472 USA
bernstein@eastgate.com

ABSTRACT
Card Shark and Thespis are two newly-implemented hypertext systems for creating hypertext narrative. Both systems depart dramatically from the tools currently popular for writing hypertext fiction, and these departures may help distinguish between the intrinsic nature of hypertext and the tendencies of particular software tools and formalisms. The implementation of these systems raises interesting questions about the assumptions underlying recent discussion of immersive, interactive fictions, and suggests new opportunities for hypertext research.

KEYWORDS: hypertext systems, fiction, narrative, Storyspace

EXOTIC NARRATIVE TOOLS
The seeming conflict between the apparent linearity of temporal experience and the nonlinearity of hypertext has provoked much theoretical discussion, especially among those with scant experience of hypertext reading (e.g. [Birkerts 1994], [Miller 1998]). Evidence that the conflict is not inherently irreconcilable is provided by the continued popularity of early hypertext fictions (e.g. [Joyce 1990]), the appearance of new fictions (e.g. [Coverly 2000]), and the flourishing secondary literature on the subject (including Nelson award-winning papers [Walker 1999] and [Tosca 2000] and several recent books [Douglas 2000], [Joyce 1999], [Landow 1997], [Aarseth 1997], [Murray 1997]). The importance of narrative in the craft of hypertext writing has been recognized from the beginning [Bolter and Joyce 1987], for narrative is central not only to works of imagination but also to technical [Bernstein 1991] and scholarly writing [Kolb 1997]. We cannot dispatch the problems of narrative to the province of artists and literary critics, for narrative questions — relating cause and effect, sequence and simultaneity, dependence and explanation — eventually challenge all writers.

I (and others) have long argued that the nature of hypertext is best discerned by studying actual hypertexts. Over the past decade, three systems — HyperCard, Storyspace, and Web browsers — have been the most common choices of hypertext fiction writers.[1] Different writers use the same system in drastically different ways[2], but systems inevitably shape hypertexts. Are the properties of hypertext fiction, such as those observed in [Bernstein 1998] and deplored by [Miller 1998], intrinsic to hypertext, or do they arise from the idiosyncrasies of specific systems? Do the patterns we observe in hypertext fiction arise directly from hypertextuality, or through the mediation of the system?

This paper explores two exotic hypertext systems, tools suitable for hypertext narrative but dramatically unlike the tools currently in use. My motivation for describing these tools is also unusual. The customary reason for building a new system is to build a *better* system; here, I wanted to build a *strange* system, a hypertext environment that might let us step back from Storyspace and the Web in order to gain a better perspective. I do not wish to argue that these systems are better than, say, Storyspace: Storyspace is simpler, more elegant, more flexible, more widely available. I do not suggest that hypertexts written with the new system will be better than those written with other tools.

For our purposes, we need not be better, we need only be different.

In the remainder of this paper, we'll first explore a language or notation, *Card Shark,* that describes *sculptural*

---

[1] The dominance of these particular systems among fiction writers need not be ascribed to any inherent virtue or suitability to the task. Accessibility plays a crucial role, as do the accidents of history. If Guide, Trellis, or NoteCards had survived to develop a literary following, our current impressions of the nature of hypertext narrative might be quite different.

[2] Compare, for example, *afternoon* [Joyce 1990], *Lust* [Arnold 1993], *The In Memoriam Web* [Landow and Lanestedt 1993], and *True North* [Strickland 1998]. All were written with Storyspace, but their use of links varies tremendously.

hypertexts. Shark is small, simple, and appears not to be very expressive, but it can readily describe complex hypertext structures. Next, we embed Shark in a dramatic context; we create a simple theatrical environment that represents characters moving through space, a space through which the reader moves to witness and perhaps to participate in the action. The nature of this participation, though it shares some surface properties of interactive fiction, may escape internal contradictions that confront conventional immersive fictions.. Finally, we conclude with some thoughts on how this approach might be evaluated.

SCULPTURAL HYPERTEXT AND CARD SHARK

Conventional hypertexts take a set of unconnected nodes (or pages, or lexia) and link them together. Card Shark[3] begins with a set of nodes, all of which are connected to each other, and builds structure by removing unwanted connections. We call Card Shark *sculptural* because we create structure by removing unwanted connections, much as a sculptors may create objects by removing unwanted material. Traditional hypertext tools in this sense are *calligraphic;* we create structure by adding lines, one after another, until we have added exactly the necessary degree of connection.

Where this strategy has been employed in the past — most notably in Malloy's *its name was Penelope* and in Malloy and Marshall's *Forward Anywhere*— it has been chosen in part to deemphasise temporal sequence and textual structure [Golovchinsky and Marshal 2000]. Card Shark, as we shall see, foregrounds sequence and emphasizes structure.

A Card Shark node (or card) contains some text, typically a brief, focused passage. Each card may also specify constraints on the context in which it may appear. For example, **AFTER 10** requires that the node may be visited only after ten other nodes have been seen. A node that appears **BEFORE 25** may only be visited early in the reading; if it is not seen early, it will not be seen at all. A variety of constraints may be applied to a node; as in a conventional hypertext, it is likely that some nodes may never appear in any given reading.

Each card may also specify modifications it makes to the reading context, chiefly by posting assertions on a blackboard. A passage that serves to introduce a new character, for example, could **ASSERT WENDY**. Other cards that **REQUIRE WENDY** can be visited only after this introduction. Later, a passage may remove **WENDY** from the scene and **RETRACT WENDY**.

Given a collection of cards, we read them by following a

Figure 1. Card Shark nodes establish conditions that subsequent nodes must satisfy.,

simple set of rules[4]:

1) The collection of cards is shuffled and the blackboard is wiped clean.
2) The reader receives 7 cards from the deck.
3) The constraints for each of the player's cards are evaluated. Cards whose conditions are not satisfied are disabled; the reader sees at most a brief title and an indication of what conditions need to be satisfied for the card to be seen.
4) The player chooses among the cards whose constraints are met, selecting a node to visit next.
5) The selected node is visited. Its full text appears (or is performed) on the screen. If the node makes assertions or modifies the environment, those actions are performed. The card remains "on the table"; we may look at it again whenever we like, but it will never be "played" again.
6) The player receives a new card, and repeats until the reading is over.

If the granularity of the lexia — the size of the "card" — is large, Card Shark's constraints describe the episodic architecture of the narrative. If the granularity is very small — individual words or phrases — the constraints describe a text generation engine. If the lexia were lines of iambic pentameter, the constraints could describe a rhyme scheme.

We can easily envision other variations. In particular, we might use one deck but maintain two separate blackboards (perhaps called plot and subplot). The two blackboards provide separate contexts in which cards could appear, offering the reader and the writer greater flexibility.

In conventional hypertext tools, connecting nodes in a sequence is easy but connecting nodes in a dense tangle usually requires effort. Card Shark inverts the situation; making a tangle is easy, but crafting a strict sequence is

---

[3] At the first hypertext conference, Halasz drew a distinction between "Card Sharks", the proponents of granular hypertexts (HyperCard, Storyspace, KMS), and "Holy Scrollers." advocates of stretchtext (Guide, Zig-Zag.).

[4] The details of these performance rules are often arbitrary and the reader may easily envision alternatives. Tinkering with rules has been a fruitful source of amusement and time-wasting during development. My current speculation is that these procedural details are not terribly important; they affect pacing and rhythm during performance, but effective writing is far more important to the performance than clever rules.

hard. The tangle, not the link, is a Card Shark primitive. (Some versions of Storyspace include power tools for making lots of links at once, and Shark may benefit from power tools for making lots of assertions at once)

Card Shark is close in spirit to van Parunak's HyperSet [Parunak 1991], but its goals are quite different[5]. If, we arrange the cards in a prescribed sequence before each reading instead of shuffling them, Card Shark appears to be formally equivalent to Trellis [Stotts and Furuta 1989]. Implementing a Shark performance engine is trivial; a prototype was implemented in Macromedia Flash.

LIVENESS, CARD SHARK, AND WRITING EXERCISES

Like its Petri-net ancestor Trellis, Card Shark is naturally concerned with *liveness.* Consider a reader in the midst of reading a Card Shark hypertext. We can readily envision that a time may come when the reader examines her seven options and finds that none of the preconditions are met. The position is dead; the reader is stuck[6]. Sooner or later, this is inevitable; we will run out of cards, the story must eventually end. But the story must have a chance to play out first; writers need to take care to let the story begin before it comes to an end, to avoid stranding the reader at the outset.

Consider, for example, a Card Shark hypertext that describes an twilight encounter in the garden on Tuesday night and its dénouement in the nearby bedroom the following morning[7]. Some actions **require night**; others **require morning**. If it's currently night and our available actions require that it be morning, we need an transitional action that moves from night to morning. Conversely, if we have been reading about events in the bedroom and there is more to learn about last night's encounter in the garden, we require a transition that moves from the morning to the previous night, from the bedroom to the garden. Indeed, if readers are not to constantly encounter dead positions, we need to provide a rich assortment of transitions[8] to facilitate movement, to shift between times, to get characters on and off stage.

These transitional nodes closely resemble a familiar class of writing exercises (see, for example, [Gardner 1983]). It is

---

[5] HypertSet was intended as an aid to taxonomic reasoning, an analytical tool for scholars rather than a performance medium.

[6] The situation is directly analogous to reaching a conventional hypertext node with no outbound links, or to an short, inescapable cycle that signals closure [Bernstein 1998]. Early hypertext systems worried about deadness almost as much as they worried about disorientation [Bernstein 1991]; today, we press the Back button and wonder what the fuss was about.

[7] Beginners often assume that narratives must be related in chronological sequence, that other arrangements are unnatural or "artistic". It often proves desirable to depart from chronology, both for clarity and for pacing. Events in the garden may be interesting only for what follows; unless we know the consequences we may not understand or care about the antecedents.

[8] Multivalence is not a vice [Bernstein, Joyce et al. 1992], and these transitions need not limit themselves to their immediate business. Finding ways to add incident or illuminate character in the course of the transition is an interesting and rewarding chore.


Figure 2 When text and image share the visual frame, the competition for space and for attention creates enormous problems [McCloud, 1993 #529]. Dynamic typography can provide more space for textual expression, although the conflict between textual meaning and letterforms remains problematic. The output of the current Thespis prototype is a script— specification for performance that may be reenacted mentally by the reader or literally by a visualization program. Even if the visualization is complex, however, it need not be literal or cinematic.

important to observe that the text of a node need not merely describe the state change. The text may, for example, proceed from the consequence of the asserted change without describing the transition. Indeed, some transitions may not need to be expressed, either because the reader will understand them or because we want to startle the unwary.

Transitional nodes, when used naturally, tend also to maintain coherence and causality. At the same time, cycles do not naturally occur in Card Shark; to permit recurrence, the same passage must occur on two or more different cards. Incoherence, oscillation, repetition, and cycles are sometimes seen as inherent to hypertext[9] but may inhere more closely to the tools we happen to be using and the inclinations of their creators.

SOCIAL SHARK

Reading is often considered a solitary activity, but we might also enjoy Card Shark hypertexts with company. Extending Card Shark for collaborative reading creates *Social Shark,* and with it some interesting opportunities.

Consider two readers, Mr. Green and Ms. Blue, who meet (perhaps over the net) to read a Social Shark hypertext together. The computer unwraps a fresh deck, shuffles the cards, and deals seven cards apiece to Green and to Blue. The two readers take turns, following the rules of Card Shark; the reading continues until neither reader can continue.

Each card, in addition to its text, its preconditions, and its assertions, is labeled with a green number and a blue

---

[9] Especially by those who don't like the hypertexts they've read [Birkerts 1994] [Miller 1998].

number. Whenever a card is played, Mr. Green receives number of points specified in green, and Blue receives the number of points specified in blue. The reading continues until neither player can continue, and the winner at the end of the reading is the player with the highest score. (Extensions to additional participants are easily envisioned by adding additional score numbers. Alternatively, a third player might seek to maximize the combined blue and green score, a fourth player might seek to minimize it, and so on.)

How might hypertext writers assign these values to cards? One simple approach identifies each player with a goal: Green is rewarded whenever Love grows between two characters, and Blue is rewarded whenever complications or misunderstandings separate them. The goal may extend to plot and subplot: green is rewarded whenever *anyone* falls in love. But the goals of Green and Blue need not be orthogonal: Green might be rewarded when Love advances while Blue is rewarded by anarchy and chaos[10]. In this case, Green's triumph is a romantic comedy — *Pretty Woman* or *The Tempest* — while Blue's triumph is hilarity — *Horsefeathers* or *The Importance of Being Earnest*.

THESPIS

Thespis[11] is a new hypertext system that extends the core idea of Card Shark by allowing many agents to participate in a single hypertext story[12]. Each agent or *actor* receives cards that describe possible actions, and each in turn selects an action to take. One agent represents reader; the reader chooses actions as she wishes. The other agents are computational structures; they choose for themselves.

Each actor has a name and a simple internal state. Each actor also has its unique function **Happiness(state)**.[13] Actors choose actions that are likely to improve their happiness. One character may value money more than another; one may crave excitement while another avoids it. Faced with the same options, different character may choose different actions. These crude behaviors are not meant to model psychology, but merely to provide the appearance of intentionality and individuality. We are not making people; we're making theater [Laurel 1991].

Each actor moves across the bounded, two-dimensional space that represents the *stage*. The reader sees and hears things that happen nearby; more distant actions may be

---

[10] If collaborators aren't conveniently available, we can easily arrange for the computer to play a hand — or even several hands.

[11] Thespis, a poet of sixth century Greece, invented theater by elaborating a traditional musical form for soloist and chorus. Thespis introduced a second soloist, a character who could interact with the main character and with the chorus. Two characters introduce the possibility of dramatic conflict as well as crowded dressing rooms and arguments over top billing; liturgy becomes drama.

[12] I tried to get a second character onto the hypertext stage in "Hypertext With Characters" [Bernstein 1995]. The earlier work seems to have been completely ignored, perhaps because its implementation was thought to be infeasible. The dearth of new hypertext systems in the late 90's is conspicuous, and may have led us to overestimate the difficulty of implementation.

[13] In the prototype, state variables included Cash, Status, Excitement, and Well-Being,. Happiness was represented as a linear combination of state variables.

unnoticed. This spatial component neatly reifies the hypertextuality of Thespian space; rather than following *this* link and not *that* one, we are standing *here*, not elsewhere. Perhaps we are sitting and talking with Hugh, Wendy, Cathy, and Kai. Across the room, we might see Randy and Stuart arguing with Susana, but if we want to hear them we'll have to walk over there -- and then we'll miss the action now unfolding before us.

We may constrain actions in Thespis by reference to the environment and to the context. Consider a node that reads,

> *"It's getting dark. Winter is coming. I tried to remember winter – the last winter before the war. It seems so long ago."*

Without constraint, any actor might say this. But we can easily impose constraints, choosing who may say this, or to whom it may be said. We could specify when and where it can be spoken. A variety of partial constraints are provided; for example, an action that is **OnlyExcited** can only be performed if the actor is unusually agitated, and a **Private** action can only be performed if the participants don't know they're being observed.

Actors and actions are simple. Rather than create complex actors, we create simple automata that say interesting things about important matters. Actors can move (to a landmark or to another actor), they can use props (eating, for example, if they feel like it), and they can talk. The point of this computational mechanism is merely to keep the actor-automaton from breaking the theatrical illusion. We enforce a naive physics of the stage, decreeing, for example, that actors should never walk through walls. Simple logic can give rise to complex emergent behavior [Resnick 1997], and this aggregate behavior can be convincingly organic.

Perhaps more important, though, is the recruitment of the reader as a dramatic co-conspirator. If simple automata are well written, if they are engaging and convincing, readers will want to attribute agency, intentionality, and emotional depth to them [Reeves and Nass 1996].

A THESPIAN EXAMPLE

What would it be like to read a hypertext written for Thespis? Let's imagine one[14]. We'll call it *The Trojan Kids*, an experimental adaptation of Euripides' *Trojan Women*, in modern dress. It's set in a large, open, metal-roofed shed, a community center for a small village that has recently fallen to the conquering invaders. It's the night of the big school dance; life goes on. It could be France in 1940, it could be a village in Rwanda or Kosovo or Chechnya.

---

[14] I beg the reader's indulgence for this lengthy exposition of the plot of an unimportant hypertext prototype. Demonstrating narrative is a vexing problem; the only way to understand a work is to experience it, and even then the illustration may founder on accidents of taste, interest, or understanding. The intent of this section is to establish the example in sufficient detail to permit the reader to construct a similar hypertext in her own laboratory. The alternative approach — presenting a formalization of Thespis — seems futile; the history of hypertext-as-a-tuple is not bright, and neither Shark nor Thespis lend themselves to elegant formalization. Consider this the experimental section.

The reader is ALICE, uncertain, unsure, unimportant[15].

She is met at the doorway by EMILY, a plain and unpopular student who has done most of the work of arranging the dance, setting out the refreshments, getting permits from the Provisional Government. She greets us warmly,

> *Emily: Come in, come in. I'm so glad you're here. Everybody's here. Come in, let's all be together, together again. Let's celebrate the blessings of peace.*
>
> *Alice: But, Emily, we lost! After ten years, our gates lie in ruins, Greek soldiers patrol the streets, smoke rises from the palace. What blessing is this?*
>
> *Emily: Defeat is bitter, sure. But now we have peace! At last! With honor! And in our time. Now we can have our party. It's our tradition, and the Greeks gave us a permit. Go on in. Try the shrimp -- I hear the dip is really spicy!*

Emily is a Pollyana, an accommodator, a collaborationist in embryo, and the stink of a dark future hangs over her irritating cheeriness. Can she be saved?

Inside, there's quite a crowd. They're kids; most of them have simple motivations. Some hope to get really, really drunk. Some hope for a memorable moment of basketball or Nintendo. Some hope to get lucky -- out back there's an area of secluded dunes. Some of these kids were conceived out there, just five or six years before the War began.

Others have more to say. CASSIE (Cassandra) is dark, sexy, strange. She knows stuff. She's seen Emily, for example, a few years from now, her head shaved, hounded through the street. Cassie knows that some of her friends here tonight will be in the cheering crowd. She knows that others won't make it that far. Cassie isn't popular, nobody pays attention to her stories, but she's hard to ignore; she draws boys like a flame and those boys draw girls. She's rarely alone.

POLLY (Polixena) is the old king's niece. She shouldn't be here, she doesn't know this crowd, she goes to private school. She was away from the palace when the soldiers came. She's on the run, she's escaped the patrols so far, and she's running out of options. Perhaps, if she can blend in, nobody will notice that there's a member of the royal family still at large. Maybe she can stay free, maybe she can live, a little longer.

FRANK and BILL are drinking from a hip flask and debating the relevance of class struggle to the War. They've

been having this debate since 7th grade, and they're good. Bill has just realized that he is in love with Polly, that her radiant smile makes the bare 60-watt bulbs burn more bright. He thinks he's never seen her before, that she's a new kid; he's seen her on TV a thousand times but not in jeans and a T shirt. Frank has known for years that he's in love with Bill, and he sees this immediately and knows that it cannot come to good.

We always begin at the entrance, with Emily, but after that our experience depends on our choices. Perhaps Cassie and her coterie are hovering over the refreshments; we might join them. Cassie has plenty to say (and she can say it, because she's a prophet and prophets aren't bound by temporal constraints). Or we might first wander over to Polly. Perhaps we stand off a little ways and eavesdrop as Bill tries to chat up Polly while visions of sand dunes dance in his head. Perhaps Emily rushes up with cups of punch, urging Polly to cheer up and have a *great* time and get out and dance!

Each of the characters have things to say. They're kids; they'll tell anyone who'll listen their insights and philosophies. They move in clusters (as kids at a party do), and sometimes individuals or couples will spin off or two groups will coalesce. Topics of conversation are introduced, old topics are exhausted or discreetly abandoned. Questions and conflicts abound -- each individual and unique, but each also connected to the others and to us.

PATTERNS IN THESPIS

Is this a hypertext? *The Trojan Kids* has no links, no blue text, no map view. But these are mere externals. Thespis discloses a chunk of text and then offers the reader a set of choices, and the choice selected determines what is seen next. This sounds like a hypertext. In Thespis, a reader has seven possible choices at any moment, some of which might be unavailable.

Hypertext, Joyce observed, requires rereading [Joyce 1994] [Rau 2000]. Thespis can be enjoyed on first reading, but its game-like qualities encourage rereading as well. In the *Trojan Kids,* Alice isn't assigned a mission — there are no captive princesses to be saved — but there are lots of things she could do differently, and many consequences can be imagined. Nothing Alice does can operate by brute force; she can't save Polly by fighting off the Invading Army with her super-powers, she has no spell to redeem Emily nor elixir to cheer Cassandra. But if she had taken Polly for a midnight walk on the beach, perhaps the police would have missed them in the dark? It's worth a try. Ineptly done, this is mere puzzle play; done well, it's the author's expression of the tension between tragedy, where fate is inexorable, and comedy, where our effort[16] can perhaps be rewarded with triumph.

Coherence, causality, and closure — those suspect qualities whose (perceived) absence bedevils the reputation of

---

[15] The unimportance of the point-of-view character may be essential to Thespian hypertext. If the reader's point of view is a hero protagonist, for example, the reader is led naturally to test the limits of the possible, That's what heroes do. The drama rapidly devolves into a negotiation between the reader and the world model; the reader asks to do unexpected, the system typically responds with incomprehension. Ironic detachment makes things worse, not better; the reader-protagonist still wants to test the rules, and detachment invites the frigidity and sophomoric contempt that so often mar computer entertainments. Because the Thespian protagonist is patently unimportant, unheroic, and constrained, the reader's attention is not constantly drawn to the limits of the possible (and hence the shortcomings of the system).

[16] Or, in romance, our inherent virtue, our intrinsic wonderfulness. [Mamet 1998]

hyperfiction — can be achieved easily in Thespis if we want them. Assertions form a convenient shorthand for episodes, so explicit temporal dependences are less difficult to maintain. (We can also create Thespian dramas that use multiple decks, one after another: a set of characters and actions becomes a scene or an act of a larger drama) Closure, too, can be achieved in all the conventional dramatic ways.

What becomes of the hypertext patterns with which we have become so familiar? Some flourish unchanged; an assertion that opens up a new topic for discussion introduces a Split, and the retraction then becomes the balancing Join. Indeed, because assertions are easy to retract, Split/Join may be more common, larger, and more elaborate in Thespian hypertexts. Basements or mirrorworlds can be constructed in space; one room contains the theme, another room the counterpoint, and access between the two is restricted by a bottleneck. Feints are at once more problematic (we have no maps) and less (characters lie). And tangles, obviously, are Thespis' natural construct.

Other familiar patterns may be less common in Thespian hypertext. For instance, recurrence in Thespis tends to be rare, brief, and deliberate. Writers who like recurrence can easily provide duplicate actions. The more elaborate cycles that are the staple of many Storyspace hyperfictions are harder to reproduce in Thespis. Douglas' cycle, in which repetition signals closure, is difficult to implement, and Joyce's cycle appears infeasible.

OTHER DIRECTIONS

Thespis is a sketch, a prototype. A host of design decisions are arbitrary. Why does the player choose from *seven* alternatives? Should assertions remain on the blackboard forever, or should they fade over time? When the player sees or hears something, the current Thespis prototype adds the description to the end of a long scroll; would it be better to display the text in a large, dynamic collage [Bernstein 2000]? Or in lots of separate windows? Thespis is a chunky hypertext system that generates a smooth (and linear) text; perhaps it should be generating a smooth hypertext, or a chunky one?

It is also interesting to observe that we can add new actions — indeed, entire new characters — to an existing scene. This could give rise to several intriguing possibilities. Not only might we reread a familiar Thespian play, but we might attempt a reading in the presence of a second, supplemental character. We could envision extensible, recombinant fictions, dramas to which readers could add or remove certain characters. Or we could let readers control several characters within the same scene, perhaps through the network; constraints on actions averts the worst faults of the grafitti problem.

MY FRIEND HAMLET: IMMERSION, GAMES, INTERACTION

Thespis shares some characteristics with Interactive Fiction

(IF) — adventure games, MUDs, and MOOs. The resemblance, however, is superficial.

In IF, the reader is the player, the protagonist, the central character. Actions shape the course of events, determining whether how things turn out. In Thespis, the reader is a minor character inhabiting the periphery of the action, a witness to events that unfold. The reader's choices may indeed alter what happens, but the Reader is not the most interesting nor the most active character on stage.

Illusions that place the reader on stage necessarily founder when promised freedom of action is contradicted by the limitations of the simulated environment. IF asks us to find a creative, imaginative, and successful resolution to the dramatic problem. The imaginative reader is bound to think of things the creator never envisioned, and the reader's best thinking inevitably generates the dullest response: "I don't understand." The computational environment can never match our aspirations, and allusions to unlimited computing power of the future (the starship holodeck) can't rectify the fundamental problem: readers will always want to do things nobody (and no computer) could anticipate. That, after all, is why people are interesting, and why we enjoy fiction.

Even if we could experience *Hamlet* on the holodeck, it wouldn't work. Tragedy requires that the characters be blind (as we ourselves, at times, are blind); if you let a sane and sensible reader into the room, everything is bound to collapse. Take Hamlet: it's absolutely obvious that he should go back to school, get roaring drunk, get laid, and await his opportunity. He knows this (I.ii), Horatio knows this, Ophelia knows this. Even Claudius knows — why else send for his college pals? Nobody can bring themselves to say the words — that's the tragedy. But what's to stop the reader? Only brute force and error messages ("You can't do that") that emphasize the arbitrary boundaries of the world. If you make *Hamlet* a game, it has to be rigged.

It's not just Hamlet. Oedipus, for example, needs to get out of town and change his name, to enter the Foreign Legion or the Witness Protection Program. Antigone needs a long talk with her rabbi. Juliet needs to tell her parents exactly what she did last night. She can't, of course, but what's to stop you?

The game is rigged, and constantly calls our attention to deception. Whenever we struggle against the bonds of fate (and the boundaries of the system), we're told, "I don't understand." The more we struggle — the more conviction and intelligence be bring to the action — the greater the likelihood that the system will find no appropriate response. Card Shark avoids this contradiction by foregrounding the familiar convention of reading and drama: we may *want* our favorite characters to prosper, but as readers, spectators, we cannot choose the outcome. Thespis gives us more range of action and might offer us a chance to take a role, but that role is not central and our limitations are evident.

In contrast to the more ambitious interactive fiction projects, Thespis makes no real attempt to model character emotions

Figure 3. Architectural overview of Thespis.

or cognitive state. A trivial mechanism lets the actors choose among possible actions. What really matters is what is said, and everything that can be said in Thespis is written in advance. Authorial control retains its customary (if ambiguous) place[17].

Interactive Fictions tend to be spatial fictions; the implicit narrative of Adventure and Myst is one of travel and discovery[Jenkins and Fully 1994]. Thespis is performed in imaginative space, too, but Thespian spaces tend to differ in scale and design from Adventurous spaces. IF spaces tend to be numerous, small varied, and thinly populated. Thespis, on the other hand, uses space chiefly as a place in which actors move; Thespian spaces tend to be large, bland, and crowded. There may be interesting settings and props in a Thespian world, but these are static and durable. The actors, on the other hand, are moving and speaking; if we don't listen to them now, we may never hear what they say.

IMPLEMENTATION

Thespis was easy to build. It took less than a person-month to implement from a standing start. The tools used were not especially suited for fast prototyping: C++ and its standard libraries. The development environment was not especially powerful. Code reuse was limited to a handful of utility classes (smart rectangles and the like) the author habitually carries from project to project. Roughly five work days elapsed from New Project to a build sufficient to enact the opening of *The Trojan Kids[18]*.

Hypertext systems were once major investments, but this is no longer necessarily the case. Documentation, support, and polish are expensive, of course, but for an experiment they are of little moment.

The script for a Thespian hypertext is composed in a trivial markup language[19]. (It ought to be XML, of course, but it was faster to hand-build a jury-rig parser for these initial experiments.) An excerpt should suffice to give the flavor of the current implementation.

```
//introduce a cast member
Actor: { Name: Emily  X: 5 Y: 3 Gender: f Status:
    30
    Text: "EMILY is a nice girl, a good girl. She is
```

---

[17] The reader chooses where to be and what to do, and so may miss interesting things. A special predicate, MustBeHeard, identifies actions that may only be performed if the reader can witness them, but its ubiquitous use defeats the purpose of Thespis and it is probably best reserved for crucial revelations. The writer may be discouraged at the thought that the best lines might go unheard; of course, in conventional theater they may be left behind in New Haven or lost in the cutting room.

[18] Your results may vary. During the work, I was completing a product release, improvising end-of-year tax planning, packing holiday hypertext orders, and writing early drafts of this paper.

[19] The reviewer with even a passing familiarity with language design and parser construction may by now be saying, "I could do better than that!" It's true; you could. Go ahead. I miss the days when the Hypertext Conference had too many "here is my system" papers. You probably miss them too.

glad the war is over so everyone can be friends. She is hardworking, plain, and not very popular; she was responsible for getting the permits and making the arrangements for tonight's party."}

```
//give Emily something to say
Say: {For: Emily Assert: Dreams After: 25
    Text: "I've been dreaming of peace for
    months. For a time when we could just have
    fun again, when the clouds would lift. Life
    goes on, right!

    These shrimp are really good." }

Monologue: {For: Ilana Requires: Dreams
    Text: "I never dream. Or I don't remember."
    }
```

PERFORMANCE

An action's consequences may be both invisible and visible, and even a conspicuous consequence may be invisible if the reader is not able to observe it. Actions do not generate visible consequences directly; instead, they create an object that describes the consequence to one or more *Views*. The view, in turn, determines how to display the consequence. Conversation is audible to the reader if she is within "hearing" range, while other actions may be visible to the reader if the are performed within a larger "vision" range and if the reader's line of sight is not obstructed by walls or obstacles.

Actions are implemented as a family of action objects that encapsulate:

- methods for selecting partners with which to perform the action

- preconditions for starting the action

- preconditions for ending the action

- a description of the performance of the action

- state changes the action imposes on the actor, and on participants.

Actions do not enact their consequences directly; instead, an object describing the action is passed to the World mediator which passes it, in turn, to various *Views* that render the action for the reader. The current prototype's *Views* follow typographic conventions borrowed from the theater; movement, internal monologue, and description are set as stage directions while conversation is typeset as actor speech. The Reader view omits actions the reader cannot see or hear (and abbreviates or omits those that seem unimportant); writers may find a World view useful as well to inspect all actions that take place without regard to the reader's location.

EVALUATION

Card Shark and Thespis present alternative approaches to hypertext that stand far afield from Storyspace and the Web but share the core values of literary hypertext. The first evaluation question is, "can it be done?" Hard upon its heels, comes, "can it be done for a price?" The answer is unambiguously affirmative[20].

A more difficult question, of course, concerns utility: is Thespis good for anything?

Thespis was designed for works of the imagination. Can it be used for argumentation, pedagogy, or technical documentation? Clearly sequential presentation is invaluable for mathematical proofs and for some kinds of schoolwork. Information retrieval is clearly ideal for answering specific, well-posed questions. If Thespian hypertext has a place outside the world of imagination, that place most probably lies in exploring multifaceted topics for an expert, engaged audience.

One can readily envision, for example, a lively Thespian discussion about areas of professional and scholarly controversy. Was Captain Cook considered a deity, or merely an unwelcome dinner guest? Is manned space flight cost-effective? Should Web site design emphasize familiarity and ease of use, or strive for unique identity and value? Is there a navigation problem? We might also envision Thespian explorations of naturally discursive subjects — the aesthetics of algorithms, the beauty of chemical synthetic pathways, or the experience of life in London in 1680. Here, the weight of the argument lies in the accumulation of detail and in allowing each reader to find the specific details that speak most powerfully to them. If you want to argue that algorithms or sculpture are beautiful, you'd best be prepared with a variety of examples and let the audience tell you what they like. Wandering through Thespian spaces lets the reader see what she wants, yet also indicates unobtrusively that there is more to see, and sketches where she might go next.

REFERENCES

[Aarseth, E. J. 1997] Cybertext: Perspectives on Ergodic Literature. Baltimore, Johns Hopkins University Press.

[Arnold, M.-K. 1993] "Lust." Eastgate Quarterly Review of Hypertext **1**(2).

[20] I did not know, a month ago, what the answer would be. I delayed starting for about ten months, suspecting the cost would exceed my meager budget.

[Bernstein, M. 1991] "Deeply Intertwingled Hypertext: The Navigation Problem Reconsidered." Technical Communication **1991**: 41-47.

[Bernstein, M. 1991]. The Navigation Problem Reconsidered. Hypertext/Hypermedia Handbook. E. B. a. J. Devlin. New York, McGraw-Hill**:** 285-297.

[Bernstein, M. 1995]. Conversations With Friends: Hypertexts With Characters. Hypermedia Design. S. Fraïse, F. Garzotto, T. Isakowitz, J. Nanard and M. Nanard. London, Springer**:** 207-15.

[Bernstein, M. 1998]. Patterns of Hypertext. Hypertext '98, Pittsburgh, PA, ACM.

[Bernstein, M. 2000]. More Than Legible: On Links that Readers Don't Want to Follow. Hypertext 2000, San Antonio, Texas, ACM.

[Bernstein, M., M. Joyce, et al. 1992]. Contours of Constructive Hypertext. European Conference on Hypernedia Technology, Milano, Asssociation for Computing Machinery.

[Birkerts, S. 1994] The Gutenberg Elegies. Boston, Faber and Faber.

[Bolter, J. D. and M. Joyce 1987]. Hypertext and Creative Writing. Hypertext '87, Chapel Hill, ACM.

[Coverly, M. D. 2000] Califia. Watertown, MA, Eastgate Systems, Inc.

[Douglas, J. Y. 2000] The End of Books -- or Books without End? Ann Arbor, University of Michigan Press.

[Gardner, J. 1983] The Art of Fiction: Notes On Craft for Young Writers. New York, Vintage Books.

[Golovchinsky, G. and C. C. Marshal 2000]. Hypertext Interaction Revisited. Hypertext 2000, San Antonio, Texas, ACM.

[Jenkins, H. and M. Fully 1994]. Nintendo and New World Narrative. Communications in Cyberspace. S. Jones. Los Angeles, Sage.

[Joyce, M. 1990] afternoon, a story. Watertown, MA, Eastgate Systems, Inc.

[Joyce, M. 1994] Of Two Minds:Hypertext Pedagogy and Poetics. Ann Arbor, University of Michigan Press.

[Joyce, M. 1999] Othermindedness: the emergence of network culture. Ann Arbor, University of Michigan Press.

[Kolb, D. 1997]. Scholarly Hypertext: Self-Represented Complexity. Hypertext 97, Southampton, U.K., ACM.

[Landow, G. P. 1997] Hypertext 2.0: The Convergence of Contemporary Critical Theory and Technology, 2nd edn. Baltimore, Johns Hopkins Press.

[Landow, G. P. and J. Lanestedt 1993] The In Memoriam Web. Watertown MA, Eastgate Systems, Inc.

[Laurel, B. 1991] Computers as Theatre, Addison-Wesley.

[Mamet, D. 1998] Three Uses Of The Knife: on the nature and purpose of drama. New York, Columbia University Press.

[Miller, L. 1998] www.claptrap.com. New York Times Book Review. New York**:** 43.

[Murray, J. 1997] Hamlet On The Holodeck: The Future of Narrative in Cyberspace. New York, The Free Press.

[Parunak, H. V. 1991] Don't Link Me In: Set Based Hypermedia for Taxonomic Reasoning. Hypertext'91. San Antonio, ACM Press**:** 233-242.

[Rau, A. 2000] "Wreader's Digest: How To Appreciate Hyperfiction." Journal of Digital Information **1**(7).

[Reeves, B. and C. Nass 1996] The Media Equation - How People Treat Computers, Television , and New Media Like Real People and Places. Cambridge, MA, Cambridge University Press.

[Resnick, M. 1997] Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds. Cambridge, MIT Press.

[Stotts, P. D. and R. Furuta 1989] "Petri-net based hypertext: Document structure with browsing semantics." ACM Transactions on Office Information Systems **7**(1): 3-29.

[Strickland, S. 1998] True North. Watertown, MA, Eastgate Systems, Inc.

[Tosca, S. P. 2000]. A Pragmatics of Links. Hypertext 2000 Proceedings, San Antonio, Texas, ACM.

[Walker, J. 1999]. Piecing together and tearing apart: finding the story in *afternoon*. Hypertext '99, Darmstadt, Germany, ACM.