The Virtual Light of Other Days

A digital "time machine" for viewing historical events

Ray Gardener, Daylon Graphics Ltd. (rayg@daylongraphics.com)

Last updated: April 26, 2011 (10th edition)

Ideally, educating people about historical events would be done with time machines or remote temporal viewing technology as described in Arthur C. Clarke's book *The Light of Other Days*. To briefly summarize, the book's characters facilitate the creation of tiny wormholes which allow anyone to view events in the past, anywhere on Earth. While this remains science fiction, it is possible with present technology to do the next best thing.

Rationale

Why would we want to do this? One reason is to simplify the current method of research, which is to visit disparate Internet sites and reference texts. Another is to impart understanding by increasing the connection between cause and effect. As we see events unfold and play out in their related spatiotemporal sequences, we can more easily find and understand why events happen, which ones occur, and why certain amounts of time separate events. Combined with annotations of existing historical data (political speeches, biographies, etc.), the ability for history to speak to us "across the ages" is dramatically enhanced. A simple example is to study movements of people, and what looked like a short distance on a normal map is now easily understood to be a slow, difficult trek across mountains.

Historical recreations (both simulated and real) are of course already plentiful, as anyone who has watched the History Channel (*Dogfights*, *Engineering an Empire*, etc.) or attended a Civil War reenactment can testify. The system described here adds the ability to access such recreations non-linearly in one place using a common interface, and also with the convenience of using any Internet-connected computer.

Another advantage to historical researchers is the ability to gain different perspectives of the same event. For example, what a battle looked and sounded like from both sides.

As the gaps in the chain of cause and effect are minimized, the system increases its value to those working in the social sciences. The ability to change the level of detail (explained further ahead) also allows one to appreciate and understand the relationship between actions performed by individual persons and the aggregate, statistical behavior of social groups, and to see firsthand how the latter averages out the former (and also when the former occasionally upsets the latter).

The System

Considerable previous implementation thought is described in Fabian Groffen's 2004 thesis on digital time machines (link available in the References section). He identifies the key data model aspects as time, space, and identity (i.e., keywords), and describes the necessity of optimally representing each. This is of particular value when searching for multiple events.

For the sake of user-friendliness, I propose a low-compexity (intuitive) interface similar to existing Earth viewers such as Google Earth and NASA's World Wind.

Google Earth 6.0, in fact, introduces a basic temporal viewing facility: some places on the planet have historical imagery, and if those places are visited, UI controls appear to let the user choose the era. It's a logical, easy-to-add and easy-to-use extension of the program, and a fine start.

The solution domain would be accurately described as a subset of a generalized system. A temporal element with variable level of detail is added in the form of *event animations*.

Of notable mention is Stellarium, a virtual sky observatory. Stellarium (and similar programs) offer temporal viewing since the position of celestial bodies can be accurately computed and thus the necessary data is available in compact procedural form. For viewing the Earth's surface, however, events cannot be procedurally compressed except for simple situations, due to the inherent unpredictability of people, weather, and geological forces such as earthquakes and volcanoes.

When viewing, the normal geographic facility is always displayed, which acts as a spatial reference for one to indicate where one wishes to view the past. A *current time* is also maintained to indicate what temporal moment is being viewed. Political (and in some cases geographic) boundaries would need to reflect this moment, e.g., as some bodies of water have recently shrunk and the borders of nations change due to conflicts. Cities also become created and abandoned over time. Since the amount of data to represent all these states is prodigious, the system (in its early stages, anyway) will not include all possible states. The broad strokes, however, such as the changing borders of nations, are reasonably inclusive.

As with the normal spatial system, one focuses on a place and then zooms in to see more detail. The temporal viewer then loads those events which intersect the current time and selects the appropriate level of detail in which to render them. World Wind lets one apply animated georeferenced textures to play out events such as hurricanes, so a few aspects of the system are already developed.

The main difference between a static and a temporal viewer is that placed objects (buildings, people, vehicles, etc.) can be animated and that their very existence depends on time.

In the simplest case, an object simply begins existing at one point and time and ceases at another. Only three datums are needed: the start time, the end time, and the object's geometry. This can be sufficient for fixed objects such as homes, buildings and trees (although trees of interest may need a breakdown into saplings, young trees and mature trees).

As Clarke's book showed, it is vitally important that viewing not be restricted to camera placement. At higher levels of detail, it is necessary to be able to *track* objects of interest as they move, otherwise the user must constantly chase after them.

Historical Accuracy

Like Wikipedia, some system of vetting and editing events for historical accuracy is necessary, lest people rewrite history. This also means that some events, despite their importance, cannot be objectively included (e.g., the life of Christ). Such events must either be omitted or visually tagged with an icon indicating that the events' historical accuracy is not universally accepted as true. Since there exists an entire spectrum of alleged events, it is best to employ a *probability percentage* system of event accuracy, and tag them accordingly. If showing the migration of early humans from Asia into North America, for example, a probability of 75% would be shown (or whatever probability is deemed universally reasonable by the scientific community).

Such probability values, of course, also depend on the level of detail. While the overall migration would have a high probability, zooming in to see individual persons hunting or fishing would have a lower probability, since we have no way of knowing what ancient individuals did in particular.

Over time, historical accuracy rises as more of history belongs to the period in time following the accurate recording of data. For example, we are more certain of the particular actions of people in the Roman Empire than we are of those in ancient Greece, and we're even more certain of people's actions in the twentieth century. Future historians have available to them an ever-increasing amount of accurate events (in fact, they will have more information than they can deal with).

There is also the possibility that an event has no consensus as to accuracy even when the level of detail is reduced. In such cases, suitably described *alternative event histories* would provide event animations for the same event, with each one viewable on its own for users to compare. Where animations differ only slightly (e.g., some critical "who fired first" detail) transparency effects can be used to show two or more versions of history simultaneously.

Data Collection

Like Wikipedia, the reasonable way to gather the large amount of event data is to enlist the general population. This, however, requires modeling tools simple enough for (a) people to use when making submissions, and (b) editors to use when correcting inaccuracies or adding missing data. The use of existing geometry libraries not only aids user-friendliness but helps apply consistency for common objects (and also helps improve the computer's memory use and performance). Some of these facilities have been implemented to a good level in the Second Life network virtual reality system.

The erroneous alteration of submitted data is another negative. Again, Wikipedia shows a solution by (a) persisting prior forms of the data to which can easily be reverted, (b) denying edits once an event has reached a universally acceptable level of completion and accuracy, and (c) maintaining a board of trusted editors. Events can also be visually tagged with icons showing how trusted the event's creators are.

As time goes on, submitters find it more likely that reusable geometry and art assets are already defined, so their work can be given more to event definition rather than representation.

The sheer volume of data involved requires a special enabling technology: **four-dimensional image capture**. Programs like Canoma and Photosynth allow for assisted (or automated) construction of 3D models from photographs; the next step is the construction of animated 3D scenes from video. In the ideal case, a single submitter with a lightweight stereo-lens camcorder records footage, identifies the items of interest in the initial frames, and an automated process determines the relevant models and their events over time.

The massive amount of data involved may require an even more massive collection approach via **smart dust** -- tiny sensors deployed in truly huge numbers (billions or trillions) that can faithfully observe everything going on. Since these sensors need their energy just for sensing, they would report to a parallel system of larger computers. Peter Hartwell of Hewlett-Packard is one person working on such sensors, and progress is slow but steady.

While smart dust or other large sensor deployments raise privacy issues, they may be the only practical solution for monitoring non-people events in remote areas, especially for planet-wide events such as dynamic geology. Tectonic plate shifts are part of a large domain of events "too boring" for people to watch or to care if they are being watched. DARPA, for example, envisions using large sensor deployments in its Transparent Earth project.

An obvious difficulty for such a system is handling events that divide or combine objects, or when an object is temporarily obscured by another. Increasingly better algorithms may solve these; in the meantime, there remains a wealth of useful data that can be collected by carefully planning to avoid problems, or by manually helping the software along when practical.

Another possibility is that billboarded video (using one or more flat surfaces with animated textures) does not impose impractical memory requirements. However, the usefulness of billboards would be localized spatially, and viewing angles could be constrained. While overviews of hurricanes are normally viewed top-down and are thin enough, exploring the same weather system at closer distances (or from the inside) requires a proper animated 3D model.

Event Definition

What comprises an event? The user may think of point events in singular moments in time, such as the position and state of objects frozen in the moment. An event is also a sequence of such moments evolving from one point in time to another. Like video, much compression can be had by exploiting the fact that events often change little between subsequent moments.

We can also model a singular moment as an event whose start and stop times are equivalent (i.e., an event of zero duration). Such moments, however, need to include a *temporal resolution* so that one need not set the current time precisely to match the event's time, but rather to some reasonably near point. Or we can model the event as unchanging over a temporal span.

Events need to furnish geometry given a time. Some events may find it easier to think of a local normalized time t which is zero at the start of the event and 1.0 at the end. In the most efficient case, the geometry can be interpolated between fixed asset sets. More complicated events need to have more asset sets located at intervals and interpolate between those.

Events can also have level of detail, in which case they break up into subevents, each of which furnishes geometry at smaller sizes and/or greater resolution.

Sound clips can be associated with events, and similar to geometry, can be furnished procedurally (e.g., environmental noises such as chirping birds, rain, etc.) or prerecorded. Sounds would of course only sound correct when the system's current event playback temporal rate was close to realtime. Non-procedural sounds typically add considerable data to an event, and suitable compression schemes are all but mandatory, not just for storage but also to furnish sounds in reasonable time from servers to clients.

Event Spatial Limits

Some events, such as airplane flights, rocket launches, satellite deployments, etc. require events to have geometry located well above the Earth's surface. This, however, is reasonably facilitated by including an altitude component in object positions. Events such as the Apollo moon landing, however, may need to be limited to those actions within Earth orbit. The system could be expanded to include the solar system and thus include interesting space mission events. The navigation interface and other subsystems would need corresponding greater work; e.g., events ocurring on Mars would use that planet as the spatial reference, and viewing events between worlds would not use any planetary navigation interface nor coordinate system (they would likely use the Sun as the origin of a high-resolution solar system space).
While these logistics are complicated, the events themselves can be more easily described, since they behave more procedurally. For example, the motion of space probes in orbit and between worlds is more a matter of simple physics formulae.

Resource Requirements

What kind of resources are necessary to view events at sufficient levels of detail? Let us answer this by modeling an event which has a good combination of features such as wide geographical extent, many levels of detail, and a timespan covering several years: World War II.

A good starting point would be to model the overall troop movements between September 1939 and August 1945. These can be viewed from an altitude of many kilometers and represented using simple objects such as arrows, polygon regions of different colors, and icons representing battles and significant items such as the sinking of the *HMS Hood* and the *Bismarck*. If we model two military events per day on average (one each for the Axis and the Allies), we have 4,380 events. If the events on average require a start/end time, and then for each time a position, size, direction vector, and reference to a identifier and to an item in a geometry asset set, this is about 150 bytes per event, for a total of 642K not including the geometry assets (but those assets would not be very large, given their simplicity, so we could safely assume a total under one megabyte).

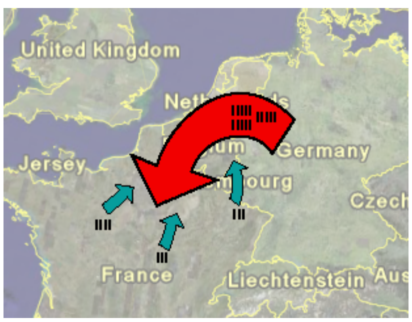

Example military event, low level of detail

While such resource usage is well within the scope of modern desktop computers, the level of detail does not capture the interesting nuances of combat such as, e.g., which particular aircraft and pilots flew in each mission, how they flew during dogfights and ground strafings, where they sustained enemy fire, etc. Without these details, the personal nature of history (which makes history interesting) is unavailable. The event log for a single aircraft could easily run into a hundred kilobytes in one mission, and given that the average squadron has a dozen aircraft, we're over a megabyte of data for a flight mission. The geometry assets also become increasingly detailed. If we assume an event increase factor of 100 and an event size factor increase also of 100, then the entire war at this level of detail would be ten gigabytes. This is storable on modern desktop computers, but not loadable into RAM. However, as in the case of geography, this level of detail is not viewable all at once anyway, and thus only sections (probably several megabytes each) need be loaded. Even if one had such a robust computer, information overload of the user would render the functionality moot -- the virtue of the system is its ability to find and focus on particular sets of events, and smoothly collapse large scale sets into lower levels of detail.

Numerous military events at high level of detail. Resource requirements currently limit such events to short durations, small geographic coverage, or both. As time goes on, data collection and technology improvements allow for representation of not just all of World War II but all of human history.
(Image courtesy of the *Call of Duty 2* videogame).

Search and Storage

Given that our ten gigabyte example not only does not render World War II in full detail -- but that the war itself is just a small part of history -- comprehensive search of the event database becomes crucial. Needless to say, data size is easily in the terabyte or petabyte range, and thus must be stored either on central clusters or distributed. Interestingly, both storage and search have been successfully handled in the non-temporal domain by Google for regular Internet search, and thus the problem can be viewed as a temporal extension of their current infrastructure. An interesting analogue is Google's own caching and the Wayback Machine, a site that searches and displays older versions of Internet sites.

Miguel Helft of the New York Times, in his January 2009 article *At First, Funny Videos. Now, a Reference Tool* noticed that YouTube is being increasingly used as a search engine and could even be the future of Google. There is an indication that the media types involved in search are becoming richer. As such systems become easier to use, a full spatiotemporal system comes closer to being. Timestamping videos and allowing temporal searching, for example, would be a logical next step.

At the same time, Wikipedia has increased its storage by an order of magnitude and also raised the size limit of uploaded files.

Future Uses

Since history can start even a short time before the present, our descendants may well use the system as their main source of information, folding in traditional Web search. E.g., news organizations would not only report the news but archive it in a manner usable by the system, letting viewers understand the events in 3D and in the context of the rest of the system's data.

Indeed, people may wear personal "life recorders" to automatically and continually submit events (mostly about their own movements). Journalists would be the obvious early adopters, but anyone else could do the same. The more witnesses to history, of course, the greater its accuracy. A key technological element in this kind of recording is work being done by Microsoft's Photosynth project, which lets one easily georeference and assemble 3D spatial sets of photographs into a virtual space.

Photosynth in action

A more localized approach is taken by Cambridge University's Qui Pan, with his ProFORMA (Probabilistic Feature-based On-line Rapid Model Acquisition) system, which lets one simply show a physical model to a single camera at different angles, with a textured mesh computed on the fly. While Photosynth is good at acquisition of realworld buildings, ProFORMA can help collect other necessary objects.

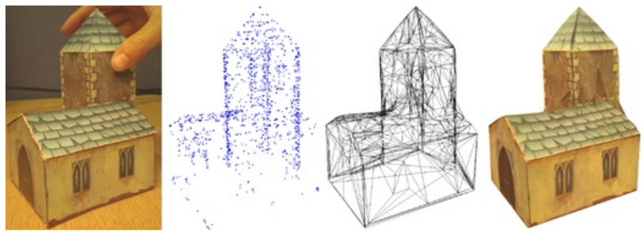

ProFORMA collecting a point cloud from a simple video feed, then computing a mesh and applying surface textures

Such objects are static, since the system does not compute any deformation data. However, David Rosen of Wolfire Games has automated the tedious aspect of "skinning", whereby deformation weights are assigned to a mesh's vertices given a skeleton, using a heat diffusion algorithm.

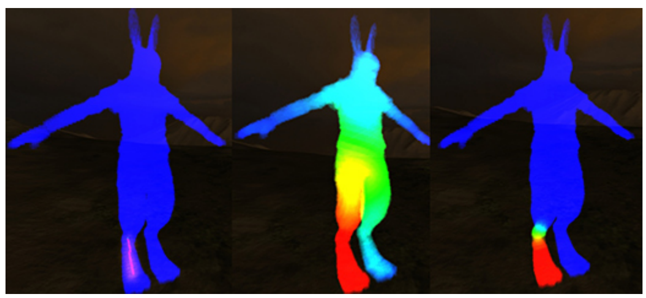

Leg bone "diffuses heat" through voxel model, providing accurate vertex weights for deformation skinning

Privacy and Safety

Privacy and safety issues are unavoidable but legal resolutions would be easier to obtain (indeed, these are main themes in Clarke's book). Where events have been accurately recorded, investigators can readily obtain witness data from any perspective. For reasons explained below, the short-term utility of the system is largely in these rich near-present explorations.

Copyright issues? Certainly -- if, e.g., a sporting event is recorded in sufficient detail, one can view it from the vantage point of anyone who actually attended. While not as detailed as the real thing, the user also enjoys the perks of today's TV watcher, such as being able to pause, rewind, fast forward, etc.

Really Far Ahead

Given ultimate physical limits, the levels of detail attainable are traded off against which events to record. Indiscriminate recording of events is unlikely, and as the system grows more popular, care must be taken not to let dependance on it lead to easy "wholesale rewrites" or "wholesale erasures" of the collective memory.

Farther still, recording is limited by the amount of available matter and energy in the universe. Even assuming that life continues forever, both the speed of light and the (apparently) finite volume of matter require a civilization to increasingly become more selective in what it remembers.

Ultimately, the details of history give way to the summary lessons they can teach, and for this reason, our ability to craft these summaries -- by starting with the right details, even if we must later let go of them -- is what matters. Failing which, we are doomed to repeat history.

References

Amazing Software Turns Cheap Webcam Into Instant 3D Scanner, Wired Magazine, http://www.wired.com/gadgetlab/2009/11/amazing-software-turns-cheap-webcam-into-instant-3d-scanner

Canoma, http://www.canoma.com

With Darpa's 'Transparent Earth,' Underground Doesn't Mean Out of Sight, http://www.wired.com/dangerroom/2010/02/darpas-plan-for-world-domination-map-entire-planets-underground

Google Earth, Google Inc., http://earth.google.com

The History Channel, http://www.history.com

The Light of Other Days, 2000, Arther C. Clarke and Stephen Baxter, ISBN 0-00-224704-6

NASA World Wind, NASA, http://worldwind.arc.nasa.gov

New York Times, Miguel Helft article, http://www.nytimes.com/2009/01/18/business/media/18ping.html

Photosynth, Microsoft Research, http://labs.live.com/photosynth

A preliminary study on the construction of a Digital Time Machine, Fabian Groffen, http://homepages.cwi.nl/~fabian/Teaching/scripties/scriptie.pdf

Second Life, Linden Labs, http://www.secondlife.com

Smart Dust? Not Quite, but We're Getting There, New York Times, http://www.nytimes.com/2010/01/31/business/31unboxed.html

Stellarium, http://www.stellarium.org

Volumetric Heat Diffusion Skinning, Wolfire Games, http://blog.wolfire.com/2009/11/volumetric-heat-diffusion-skinning

Wayback Machine, http://www.archive.org/web/web.php

Wikipedia, http://www.wikipedia.org