Practical Modeling of an Architectural Landmark

by John Locke, Rudy Darken, and Michael Zyda
Department of Computer Science
Naval Postgraduate School
833 Dyer Road
Monterey, CA 93943-5118
408/656-2844, 408/656-4083 (fax)

Designing a building, as an architect, is an act of creation. Modeling an existing building, for use in graphical simulation, is an exercise in reproduction and thus presents its own set of problems.

The nature of those problems can vary widely. One model might be as simple as a rectangular form with photo-realistic textures applied to the polygons to give the illusion of complexity. Another model might approach completeness, with the geometry of all architectural features, even finer-grain details like doorknobs and plumbing fixtures, included. Ultimately, the fidelity of the finished model will depend on some balance of the tools and methods of model-building, the availability of time and labor and, most important, the intended use of the model.

In the NPSNET Research Group of the Naval Postgraduate School, we are exploring the value of using real-time graphical simulation to teach navigation through specific real-world environments, both indoor and out. In one experiment we propose to use three groups of subjects. The members of one group will (individually) attempt to learn their way around a complicated building with only floor plans to guide them; a second group, also with floor plans, will explore a virtual re-creation, rather than the actual building; and the third group will have only the plans. None of the subjects will have prior knowledge of the building. In the crux of the experiment, members of all three groups, sans floor plans, will be tested in their ability to determine the best paths between locations in the actual building, i.e. point A to B, B to C, etc. We can then compare the performance of the groups.

Obviously, a virtual model would not likely have been available for any building that met our needs and would therefore need to be custom made. We selected, both for proximity and complexity, Herrmann Hall, the administration building of our campus in Monterey, California. Herrmann Hall is the former Del Monte Hotel, a luxury establishment originally built in 1880, which fell on hard times during the Depression, before being requisitioned by the Navy in WWII. The central part of the building has six major floors, a pair of transverse wings to its main structure, undulating terrain that sets ground level at two different floors, numerous stairwells, including one that rises seven stories, and over a half-mile of zigzagging hallways. Despite its many unique features that serve as landmarks, it is, as new students and professors discover, an easy building to become disoriented in. Of further relevance to our experiment, its complexity allows selection of locations with alternate paths between them.

The implementation tools for the virtual exploration of the building were not at issue. We had extensive experience building Flight format models using MultiGen, a high-end modeling tool for

SGI workstations. Additionally, our homegrown simulation system, NPSNET, uses Flight models for terrain, static models (buildings, vegetation), and dynamic entities (vehicles, human models). For the conduct of the experiment, the subject walks through the virtual building constrained by the dynamics of a human model. The building is viewed on the monitor, with simple keyboard commands controlling the human's speed and direction. Alternatively, NPSNET supports other, more exotic, input devices, e.g. head-mounted displays, the Fakespace PUSHBOOM, an omni-directional treadmill.

Given these parameters, what then should be the fidelity of the model? The standard we seek is "an adequate immersive experience." The "immersive experience" would be a walk-through that allows the subject to sense place, distance, and time, strongly similar to the experience of exploring the actual building. We are satisfied with mere adequacy because our experience with models tells us that the highest level of fidelity we can envision is extravagantly expensive in both model-building time and simulation system resources for a building the complexity of Herrmann Hall, and would not necessarily pay off in improved subject orientation. On the other hand, since we can't predict exactly where "adequacy" lies (indeed, that is an issue the experiment may shed light on), we should err in favor of as much detail as we can afford. Ultimately, "adequacy" will be better defined when the presence of detail can be experimentally shown to improve navigation performance.

(Since the modeling of Herrmann Hall is not likely to be repeated, a pair of secondary issues, outside the scope of the experiment, arose in regard to the fidelity question. First, the building is an active, important facility, so we did not want to foreclose potential, practical uses of the model, e.g. an interactive lobby display that allowed specific offices to be located. Second, since the structure is a local landmark, aesthetic aspects are desirable in case the model is ever used in a historical context.)

We limited the interior modeling to public spaces, since private rooms and offices wouldn't be suitable for the experiment. This included all lobbies, hallways, stairwells, alcoves, and the like. An inherent advantage to Herrmann Hall is its unique and varied interior. Unlike an ordinary office building, which may be complex but composed of nearly identical floors, the different areas of Herrmann Hall are quite distinct in shape and appearance. The ground floor has long, business-like hallways demarcated with arches. One level above, the main lobby is a huge chamber with giant murals on the walls and an enormous pane of glass for a back window. The mezzanine is a pair of short hallways that end on a balcony looking down into the main lobby. Only the second and third floors, containing the hotel rooms, are much similar, and they are distinguished by different colored carpeting. The sum result is that the collection of shapes forming the interior--particularly when textured with accurate floor and wall material--goes a long way toward orienting the subject within the model.

Amplifying the basic shapes is a second layer of detail, comprised of smaller-grain features deemed important as visual cues to pedestrians. These include all steps and staircases, pillars and posts, balustrades, etc. Doors and windows are very important, and are handled in different ways. All doors are modeled, with perpetually open doors as openings, and doors to unmodeled interior spaces as door textures. All windows are modeled. Windows onto modeled interior spaces are transparent, allowing the model to be viewed through its own windows, an important aid to orientation. Windows onto unmodeled spaces are left opaque.

Given the size of the building and this scope of detail, the proficiency of the modeler was a critical issue. A one to two year turnaround to complete the model would far exceed the lead-time of the researchers. The scale of the experiment would thus be directly dependent on modeling skill. As background, the modeler had six months of varied modeling experience, both in creating new models and modifying existing ones. This accumulated experience led to a watershed, beyond which modeling Herrmann Hall became practical. MultiGen supplies a rich assortment of tools. When it is understood how these tools can be flexibly used in combination, the software, in Zen-like terms, ceases to be an obstacle course over which the concept must struggle, but rather a pipe through which it flows. The decisions are reduced to what, not how, to model. Still, the model--75% finished at the time of writing, but complete for the purposes of the experiment--has taken approximately 400 hours. It is comprised of over 10,000 user-defined polygons which the graphics subsystem breaks into 21,000-plus triangles.

Because the model is a reproduction, another critical issue was the information available on the building. Many original plans exist but not in the most useful form. Here, history comes into account. The central part of the building burned down--for the second time--in the twenties and was rebuilt along its current Mediterranean design. After WWII, an annex was grafted onto the back of the rebuilt portion. Not surprisingly, the entire building was not completely redrawn for each major change. Instead, there are several sets of plans, in different scales, drawn to varying degrees of accuracy, with incomplete information on how new joins to old. On top of that, it appears the actual building did not always follow the plan. Also, there have been numerous smaller additions and remodels over the years which render the older plans even more obsolete. The interior floor plans are newer, but occasionally inaccurate on the placement of doors and windows. Ultimately, the only true “database” of attributes is the building itself, but the building is impractical as a source of measurement data.

With incomplete and contradictory data, accuracy must necessarily yield to coherence. Since the bulk of the data would come from the less accurate source, the drawings--as opposed to the more accurate, the actual building--the model would have to be built in a way that satisfactorily resolved the discrepancies. That way is a (conceptually) top-down process that treats the building as a set of major shapes which undergo a methodical sequence of subdivisions. The major shapes are the central wing, the two transverse wings, and the tower. The first set of subdivisions are the floors. Both the outer walls and the floors are given thickness with a double layer of outward-facing polygons. Next were added vertical structures that span multiple floors, stairwells, for example. The guiding principle throughout is that the modeler must endeavor to understand the spatial relationships between elements in the building *before* constructing the polygons. To illustrate the potential pitfalls, Herrmann Hall’s main stairwell gives the impression, both from the floor plans and by actually walking it, that it rises in a straight vertical. In fact, closer study of the plans reveals that it meanders horizontally as it rises. On the other hand, an elevator shaft makes a great anchor--for taking relative measurements. We know it must go straight up.

Note that while this phase addresses the largest dimensions of the model, it also requires the greatest attention to accuracy. The significance of accuracy diminishes as the subdivisions become smaller since they have increasingly less impact on the model as a whole. Mistakes made in constructing the fundamental shapes will haunt the model through its lifetime. A potential source of measurement error would be a failure to recognize the importance of thickness, both of floors and walls. Hermann Hall’s thickest structures are on the order of a half-meter which, cumulatively, shrink the interior volume to a noticeable degree. Measurement errors will add up, either squeezing or stretching the interior spaces. Keeping in mind the intent of the project, we can't predict the effect of disproportionate shapes on test subjects.

With all major structures in place, intersections between them can be modeled. This includes cutting holes in the floors where stairwells pass through, removing sections of wall where wings cross. After that, only local features remain to be fitted around the existing structures. These include hallways, doors and windows, steps and stairs, etc.

At this point, the scale of the modeling shrinks and the question of detail--the finishing touches--arises. Here we encounter a classic trade-off. Should detail be modeled or captured in the photo-realistic textures applied to the polygons? Additional geometry results in a more computationally expensive model. But creating custom textures is a labor-intensive process, also adds a resource expense, and, in our experience, is harder to do well. The general procedure for creating a texture is to photograph the area to be modeled with a digital camera, then process the photo with Adobe Photoshop, or other image editing software. Processing includes removing perspective distortion, adjusting color and brightness, retouching defects, trimming and sizing the image, and so forth.

In a prior model of Herrmann Hall, limited to the exterior skin, we attempted to cover entire walls, with sides measuring in the tens of meters, with single textures. This produced a model that looked interesting from a distance but, on closer observation, exhibited these defects: varying conditions of light and shadow when the photo was taken made every wall a different shade; trees and other foliage impeded a direct view of the walls; taking every photo from a different angle and distance made it impossible to normalize the textures to align horizontal and vertical features on adjacent walls of the model. These defects could be remedied to a limited extent, but only with much effort. With the new Herrmann Hall model, given also that exterior doors and windows had to mate with interior spaces, we concluded that single, detailed textures should be limited to the size of, say, double doors. Or, as a general principle, the fidelity of the texture will be inversely proportional to the area it covers. Therefore, each exterior wall of the new model was subdivided into individual polygons for every door and window. The surrounding area of blank wall was tessellated with a stucco texture applied to a one-meter square area. A similar approach was taken for interior walls. A tall, narrow texture captured the interior wall color. Its bottom end was tile-colored so that when the texture was stretched horizontally across an entire wall, a tile molding abutted the floor, as in the real building. The wall was then subdivided for individual doors, windows, or wall decorations, like framed pictures. Ceilings and other blank surfaces were simply covered with wall color. The lighting function of the display software shades the surfaces, revealing the edges between adjacent polygons.

Another issue is the pixel density of the textures. Herrmann Hall currently has about 80 textures utilizing half the 4Mb of texture memory available on our SGI Onyx RE-2 workstations. Considering that the model has not been lavishly textured, there is an obvious incentive in keeping the textures small for a model on this scale. Most textures were 64x64 pixels or less. Most doors were 64x128. Special images, such as artwork on the walls, were bigger, up to 256x256. Many tessellations, like individual floor tiles, were small, on the order of 32x32. A hidden advantage to smaller textures is the decreased preparation required. At 64x64, or less, it becomes practical to retouch the image on a pixel-by-pixel basis, where needed.

A variety of Photoshop tools were employed in the preparation process. Commonly, for color and contrast, adjustment of Levels and Curves was done first. Perspective was removed with a combination of Perspective, Skew, and Distort. Images were cleaned up with the Despeckle, Sharpen, and Dust & Scratches filters. In some cases, overly clean images, which can lend a synthetic look to the model, were "dirtied" with the Noise filter.

Textures were hand-crafted for windows. For transparent windows, an opaque frame was placed around a 40% transparent glass area. (Transparency is determined by the alpha attribute of each pixel, which can be set with MultiGen's texture editor.) The glass cannot be left clear because then it is simply invisible. Shading the glass pale blue lent a nice effect of a transparent pane. The polygon to which it was applied was set as double-facing to achieve transparency through both sides of the window. The front faces of window polygons were pointed in to the interior spaces. Since NPSNET's collision detection is based on the model center, the human could walk through windows otherwise. To mimic the number of panes in the actual windows, the window texture was tessellated per window polygon accordingly. Windows fronting unmodeled interior spaces were made completely opaque and a slightly darker shade of blue.

As a final note on texturing, though we present it as a phase which follows construction of the model's geometry, the experienced modeler will fold the two processes together to take advantage of the replication of polygons. For example, in modeling a set of stairs, one stair is built, then simply copied, one step to two, two to four, etc. If the initial step is textured, the copied steps will be textured automatically.

As a final note on modeling, it should be understood that regardless of the planning and care that goes into a project of this size, measurement and alignment errors will inevitably creep in, keeping the Holy Grail of perfection ever out of reach. At a point past which rebuilding is a practical option, these errors will make themselves apparent. Indeed, this is not unlike what happens in bricks and mortar construction. Plans are never perfect and neither are workers. Herrmann Hall, for instance, despite its magnificence, suffers in places from a certain inelegance of design, a hint of human nature amidst the dream of luxury. In graphics modeling, therefore, the rule is the same as in the physical world: when it comes down to it, fudge.

The authors belong to the NPSNET Research Group, Department of Computer Science, Naval Postgraduate School, Monterey, California. John Locke is a research modeler. Rudy Darken is an Assistant Professor of Computer Science. Michael Zyda is a Professor of Computer Science and Chair of the Modeling, Virtual Environments and Simulation (MOVES) curriculum.

To contact the authors, visit the NPSNET Research Group Web site, http://www.npsnet.nps.navy.mil/npsnet/.