The Hohokam Water Management Simulation: A Collaborative Model for Exploring Alternative Pasts

John T. Murphy

Department of Anthropology
University of Arizona
Tucson, Arizona, USA
jtmurphy@email.arizona.edu

Abstract

The consideration of past social systems as complex coupled human and natural systems has led to new approaches to understanding and even discussing the past. The Hohokam Water Management Simulation assembles an array of elements like plant productivity, labor requirements for the construction and maintenance of the canal system, and the dynamics of water flow, and uses these to create a simulation environment in which programmatic agents can enact strategies of cooperation and competition along a developing canal system. The framework allows questions to be asked about the degree of central control, the system's resilience in the face of various stressors, and the overall trajectory of the Hohokam florescence. The approach is explicitly collaborative and reductionist, and focuses on 'what if' questions rather that recreations. This paper summarizes the modeling framework's goals and philosophy, and discusses some implications of these for the construction of archaeological theory and argument.

1 The Hohokam Context

The Hohokam context provides a puzzle that is simple in its outlines (a general overview is given in Reid and Whittlesey 1997). The Hohokam lived in the Phoenix basin around the Salt and Gila rivers, and around the rivers near modern Tucson, for over 1,000 years. The Phoenix basin is not an easy place to live; it is a hot, arid desert. In this environment the Hohokam built an extensive canal system, with canals that could exceed 25 m wide and 8-10 m deep, extended over 20 km, and brought water to the Hohokam fields. The Hohokam trajectory passes through several stages: a period of expansion is followed by one of intensification throughout southern Arizona; on the heels of this is a relatively stable period, followed by a period of reorganization and, eventually collapse. The Hohokam way of life ended shortly before Spanish contact.

This basic outline provides a number of mysteries. Some of these fall under old headings: the rise of an irrigation society and societal "collapse" are two of the most interesting. But there are new ways of approaching these topics, ways that derive from complexity theory and the relatively recent approach to studying complex adaptive systems (i.e., Lansing 2003).

It was with these ideas that a workshop convened in 2003 to study the Hohokam context. Old archaeological questions of cultural chronology and migration were put aside, and the assembled scholars asked whether the Hohokam system could be studied—or even explained—as a complex system. If the assumption that a large canal system implies a state is removed, the range of possibilities opened is quite wide. Archaeological evidence on the central organization of the Hohokam is equivocal at best, and whatever may have been the system in place at one time, there is also clearly change through time. So to attack the problem, the workshop brought together scholars with a wide array of expertise. Some were Hohokam specialists, others were brought to provide comparative views from other areas.

Together, these scholars brought an array of information, knowledge, and insights. There were specialists on the Hohokam canals and on Hohokam agriculture and subsistence, and on climate reconstructions for the Phoenix basin; there were others who knew a bewildering array of strategies for organizing space around floodplain agriculture, i.e., arrangements involving land tenure based on kinship vs. moieties, etc. Perhaps most importantly, there was a push to consider the Hohokam as a system within a framework such as resilience (Holling 2001).

Two major things resulted from this workshop: the first was a push to integrate the existing data—on climate, plants (various varieties, productivity, water needs), landforms, and canal systems—into a single system where they could be considered together. The second was the belief that the proper way to investigate the system was through modeling—some system akin to the Lansing/Kremer model of Balinese water temples (see Lansing 2001), in which individual agents make choices on the landscape. Perhaps the most surprising aspect of this modeling effort was the need to model alternative pasts; only by doing this could the system be tested in ways that would reveal how it might have responded to different stresses, or assess whether a change in the values for a few relevant parameters would have dramatically changed the Hohokam trajectory or if roughly the same outlines would have occurred even under many broad and different conditions.

The Hohokam Water Management Simulation (HWM) is the vessel that has been created to fulfill these goals.

2 A Modeling Philosophy

The specifics of the Hohokam problem combined with the context of the discussion led to a new group of modeling objectives and what might loosely be called a "philosophy." This is in fact the intersection of both theoretical and practical issues that the project and its goals presented. This philosophy is characterized by a number of things that we (meaning me, Ann Kinzig, and Charles Redman, who employed me in this research and discussed it at length with me, but whom I will not blame for any errors in the following) took to be axiomatic.

First, all models are abstract and reductionist. They are, at best, cartoons of reality; this is true for even the most elaborate simulation. Second, although one important goal of a simulation might be to re-create what happened in the past, a second goal is to explore our concepts of the past. Thus, we may expect to gain in our understanding of how we address the past theoretically even if we do not find out more about what actually happened in the past.

Third, modeling of the type under consideration here is collaborative. We want to bring together an array of experiences and knowledge, and use the combined wisdom of the group to explore an otherwise unmanageable parameter space. Fourth, that modeling requires pushing our thinking into some real representation—be it balsa wood or bits. I use the term "model" *only* to mean creating something other than a collection of ideas. In most cases this is implemented electronically, but the physical nature of the model means that certain rules, especially with respect to completeness, are enforced without exception, to a degree that merely conceptual models often fall short of meeting.

Fifth, the process of modeling is itself instructive. It forces us to think through details and implications that we may otherwise have let slide. The challenges of communicating the model with our peers, and then instantiating it in code, require that all the details of our concepts be laid bare, and this is inherently beneficial.

Sixth, modeling can often be done by thinking about "things" and the dynamic relationship among those "things;" typically this can be mirrored by a split between static data and code (though this is not always a perfect match).

Seventh, because we are not interested in re-creating a specific past, there is an opportunity to explore multiple alternative pasts—to "play the tape again" and see if the same thing happens (that is, to take up Stephen J. Gould's challenge, as suggested by Lansing 2003). This allows interesting and important questions: "what if the Hohokam had had different plants available, what if rainfall had been different, what if the landscape had been different, and on?

Eighth, simpler is better. There should be in any modeling project room for increasing complexity but also room for increasing simplification: we learn as much by determining what details can be omitted as we do from determining which must be included. The implication for a modeling framework is that it must be extensible, and permit expansion and refinement.

Finally, all modeling is part of a broader discussion. Models are used to build arguments and support them, but the link from initial data and supposition to conclusions should be clear.

It should be noted that these elements are more clearly differentiated in a longer publication that is forthcoming (my dissertation); there, they are more appropriately categorized, the importance of each to the more logical versus practical issues they address is made clearer, and they are more firmly linked to broader literature on modeling in philosophy of science.

3 A General Modeling Solution

This philosophy poses a challenge: how can the various goals it sets be achieved in a single modeling framework? The modeling solution proposed here is largely distinguished by having two components: a database and a simulation are joined. The database stores the information centrally and shares it among the collaborative community. The information so stored includes not only the basic data (like plant data that were previously scattered and unintegrated) but also the combinations of data that different researchers have found interesting to use for simulation runs. They can explore the parameter space together and work toward common approaches that combine their various expertise. Provisional information can be marked as such, and alternative forms of dynamics can be substituted as desired by creating and selecting among code options.

The database requires that the input data be structured carefully. The code that implements the dynamics of the simulation is also structured carefully; object-oriented programming leads toward the creation of programmatic objects that mirror the database structure and, more importantly, reflect the vocabulary of the theory underlying the simulation.

The character of the modeling framework is best understood by considering the simulation and the database that support it to be a single engine, whose purpose is to take a collection of suppositions and find their implication: suppose plants A, B, and C had these characteristics, and the Hohokam planted them in these months, and the rainfall were this much, etc., through all the complicated variations that inform on the problem at hand; supposing those things, what is the result? This is in keeping with the most common purpose of simulation, which is to think through, for us, the things we cannot think through ourselves. The difference is that we make the thinking rigorous and a communal rather than individual endeavor. The combination of a database and simulation code in an extensible framework make possible the processes modeling that the framework requires beginning with the construction of appropriate units, moving through the exploration of a wide parameter space and the construction of complete arguments, and finally to the reconsideration of the units for either simpler (if possible) or more complex (if necessary) substitutions.

4 In Practice: The HWM System

The technical details of the HWM System are bare: a SQL Server database stores input data and output data; the website allows review and contributions to source data and output; and the simulation itself is written in Java.

Some details of what you do in the HWM System are appropriate. The HWM System has a central database of plants that the Hohokam might have employed, along with their characteristics. Each plant can have many "varieties," by which is meant different sets of proposed characteristics. A plant's characteristics are simplified down to: how long its stages of life are, what its water need during each stage is, and what happens if it does not get the water it needs.

Users can create a topography and modify it using simple commands. Users can also modify the climate by varying rainfall. A simple system allows even very complicated variations of rainfall regimes to be created, often using the actual streamflow data. For example, one can create a sequence of several years of serious drought, and superimpose them on the actual data derived from dendrochronology at intervals of one's choosing. One can ask, then, what if the Hohokam had faced such a drought early in their trajectory, versus what if such a drought had happened later, when the system had matured and might be more or less able to withstand such provocations.

The collaborative exploration of the model's parameter space is facilitated through this framework. For example, a user whose expertise is in canal construction can review the collections of input data relating to certain plants that a given specialist has assembled and thought worthwhile; he can rely on the expertise of the other researcher by using these configurations, rather than constructing his own.

One key element is that the input data can be associated with references and even comments, which remain permanently linked with all output data derived from them. In this way it is possible to see the complete chain from empirical data (or provisional or hypothetical substitutes) to simulation results, and to collectively assess each step in this chain. Also, tools are available and built into the system for summarizing the results in a way that keeps this chain complete and easily accessible. This allows complete arguments, from beginning to end, to be made transparent and auditable.

5 Some Implications

5.1 Against "Big Real" Models

There is a push in archaeology—and an understandable one—to make use of the latest and most realistic software, and to gobble the most impressive collections of data. This has led to the production of models that make use of huge GIS datasets and software developed in other contexts for other purposes. There is a certain logic to this: if a plant productivity model exists—and even better, if it was developed by people who have dedicated their lives to it—why not use it?

The philosophy given above guides my response: with such complexity, how do you simplify? How do you know what are the essential elements of the system? How can you play the "alternate reality" games that are crucial to understanding how the system would have rebounded to perturbations, and, when you do, know what elements are key? In the modeling framework proposed here, it is possible to work towards the levels of abstraction that are necessary to extract from the specific context being modeled the more general underlying principles of interest.

Note that the argument applies best if the goal pursued is revealing these general principles; if a more specific reconstruction of a particular modeling target is the higher priority, the more detailed and complex approach might be more appropriate.

5.2 The Scope of the Problem is Key

This is not a model of everything; it is not a model that tries to include every detail of everything that was going on. Nor is it a model that tries to do just one thing. As mentioned, we have taken great lengths to ensure that the model can be made more complex or simplified as needed, and one aspect of this is that it can be used to explore a variety of questions. However, this is not an infinite expanse: there are boundaries, and it is these boundaries that make a general approach possible without slipping into the morass of trying to model everything.

5.3 Model as a Repository

There is a conviction in our modeling group that our model should not simply serve one purpose and then be discarded; we hope for it to outlive our involvement. We want it to be a place where scholars who are interested in a problem share data and work together to create knowledge. To be fair, there were a number of other concerns that led us in this direction—a simple one being auditability of what we have done (every aspect of the simulation process should be both transparent and reproducible). But, particularly with respect to one of the main goals to grow from the 2003 workshop, that of integrating the plant data into one place, we are especially happy that the HWM environment is one that stores and shares data effectively.

5.4 Agent-Based Modeling

Agent-Based Modeling is an extremely compelling kind of modeling, and it is appropriate to discuss how (or whether) it relates to the modeling framework presented here. the two are very different, but there is a strong case for finding useful relationships between them.

Agent-Based Modeling means that programmatic agents make decisions based on their own internal states, their own objectives, and their own positions within the larger simulation. There are several kinds: one can have an agent-based model with only one agent on a landscape; one can also have an agent-based model in which there is a population of agents, and, indeed, where the main objectives and challenges for any single agent fall exclusively from its relationship with the other agents. The fundamental principle is that the agent or agents are assigned rule-based behavior, and dynamics of some kind—artifact distribution patterns, population structure, etc.—fall from these lower-level rules, often in surprising ways.

In the Hohokam case, there was a need to make the simulation "interactive," so that a user could assess the state of the simulation and make changes to it. To this end a collection of interrogative and imperative commands was developed (and is still growing). The users act through this interface. These employ, of course, the vocabulary of the simulation, which is the same as the vocabulary that has been worked out for discussing the simulation's state and structure. The user can thus make decisions within the milieu of the simulation, using the agreed-upon terms from the simulation glossary.

It is a simple step, then, to making these same terms available to the agents. Not only are the agents and the things with which they interact defined, but the elements of the simulation and the repertoire of decisions they have available are as well. We can use the modeling context as an arena to ask what information would have been available to each agent and what actions they would have had available.

5.5 Comparability and Generalizability

Part of the exercise of Assertion-Based Modeling is simplification; this leads to models that are less closely tied to their original details, and thus can be tested in other contexts. This allows the underlying principles to be abstracted more easily: we learn more general dynamics instead of just specific history.

This also paves the way for tests of comparability, as we can compare approaches developed in one setting with those used in another more easily, because a move toward the abstract domain where the two contexts share particular characteristics is already built-in.

5.6 The Modeling Effort as a Microcosm of Science

Above I discussed the belief that modeling both helps and forces us to clarify our concepts; this is intrinsic in the refinement of vocabulary that is required to construct the database and code that make the simulation work. In our context this is furthered by being explicitly collaborative: groups must find the concepts that are workable, and rigorously define them. Some find this constraining, and I concede that there is always a role for thinking outside the details. But I also believe that modeling is not merely a Procrustean bed, cutting off concepts to fit, but a file that sharpens and a lens that clarifies. Concepts should be well-defined and the implications of them explored in depth; modeling compels this. Philosophers of science would position these views in terms of a *syntactic* versus a *semantic* view of models and their relationship to theory (see Frigg and Hartmann 2006). A further elaboration of this distinction will appear in my dissertation. For now, it is sufficient to put forward the idea that modeling should facilitate the creation of larger arguments, beginning with static data and moving through dynamic relationships implemented in code, and that this should be done in a structure that is congruent with the knowledge we are hoping to create.

6 Conclusion

In summary, the Hohokam presented an opportunity to study a long-time-scale archaeological record from a complex systems perspective. This entailed combining data and perspectives from an array of fields and scholars into a modeling environment in which alternative trajectories of the past could be considered. The modeling environment addresses these logical issues and additional practical ones they elicit by combining a database and simulation code, creating a structure in which a vast parameter space can be explored and cogent arguments created from the results.

The HWM System remains under development and is expected to be "live" beginning in fall of 2006. I will look forward to reporting the results it generates when they are available.

Acknowledgements

This research was funded by a grant from the McDonnell Foundation, and conducted under the aegis of the Global Institute of Sustainability. Many thanks to Charles Redman and Ann P. Kinzig for their discussions of this topic, and to the many other scholars of Hohokam archaeology and anthropology for sharing their insights and knowledge. Peter McCartney, Robin Schroeder, and Wayne Porter of GIOS provided the platform and technical support for the HWM system.

References Cited

Frigg, Roman and Hartmann, Stephan. 2006. Models in Science. *The Stanford Encyclopedia of Philosophy*. Last Updated: 27 February 2006, http://plato.stanford.edu/entries/models-science/ (Accessed 17 August 2006).

Holling, C. S. 2001. Understanding the complexity of economic, ecological, and social systems. *Ecosystems* 4:390-405.

Lansing, J. Stephen. 2001. *Priests and Programmers: Technologies of Power in the Engineered Landscape of Bali*. Princeton: Princeton University Press.

Lansing, J. Stephen. 2003. Complex Adaptive Systems. *Annual Review of Anthropology* 32:183-204.

Reid, J. Jefferson and Whittlesey, Stephanie. 1997. *The Archaeology of Ancient Arizona*. Tucson: University of Arizona Press.