CHAPTER 10

THE MODELING FRAMEWORK

John H. Christiansen

MASS SIMULATION FRAMEWORK DESIGN PHILOSOPHY AND GOALS

The computer simulation software system that has been constructed to perform socioecological simulations for the MASS project embodies software representations of many diverse aspects of ancient Mesopotamian daily life. These representations are the product of a model-building process in which MASS Group members have contributed their understandings of the structure and dynamics of the ancient Mesopotamian ways of life and have attempted to fit their insights into an evolving conceptual framework. This conceptual framework has been shaped by group consensus and by the inherent requirements of a computer simulation. Principal among these simulation requirements is a need for robustness: any and all contingencies that are not deemed outright impossible must be explicitly accounted for in the simulation's logic pathways. Once initiated, MASS simulations must be able to proceed forward in (simulated) time, minute by minute, for decades to centuries, and must always have logic paths on hand to deal with whatever intermediate states, or situations, that the elements of the simulated settled landscape find themselves in, no matter how unlikely those states may be. Meeting these simulation software requirements has brought a degree of rigor to the development of the MASS project and simulation framework that is unusual in archaeological discourse, illuminating knowledge gaps and data shortfalls that the scientific community had never seen a need to resolve or remedy before, and which indeed may not previously have been identified as deficient. The process of addressing such issues in a structured, focused way, often via numerous *ad hoc* white papers that have in some cases been published in their own right (Widell 2004, 2005 & 2007; Altaweel & Paulette in press; Chapter 12) has perhaps been the most scientifically valuable phase of the project, rivaling the insights derived from actually running the simulation and interpreting its results.

As has been pointed out in Chapter 1, for the most part an holistic, agent-based, 'bottom up' modeling approach has been adopted for the MASS modeling and simulation exercises. In this approach, the simulated historical trajectories of settled landscapes are allowed to emerge as the aggregate outcomes of a welter of finer-scale activities and interactions occurring at the *microscale* level of, e.g., individual persons, households, crop fields, and domesticated animals. An advantage of such an approach is that it bypasses the need to devise and justify global 'top down' control mechanisms that would otherwise be required in order to regulate the *macroscale* (i.e., at the scale of kingdoms, regions, etc.) dynamics of the simulation to arrive at macroscale results. Instead, for the model builder embarked on a 'bottom-up' modeling exercise, the main burden is to identify the important *microscale* processes, articulate and implement appropriate software designs for these processes, and then test and validate them. At the microscale level, the individual microprocess representations must be able to produce results consistent with observation at that scale. For MASS this has been relatively straightforward in that a great many of the relevant dynamics, such as key processes associated with primary food production and with the operation of normative kin-based behaviors, tend to be best understood at the microscale level, informed by archaeological field data, documentation in cuneiform texts, and ethnographic studies (Chapters 5 and 6).

For a bottom-up approach, additional levels of validation come into play. Once the microscale process representations have been individually tested and validated against observations, a second, macroscale level of validation must still be applied: are the aggregate results summed across larger populations and polities and larger swaths of the landscape consistent with macroscale evidence such as, for example, the historically and archaeologically attested rise and fall of cities and kingdoms and major shifts in land cover and land use? Close examination of aggregated simulation results can suggest a need to go back and rethink the design for some of the microscale process representations that, collectively, produced them. For the MASS project, an intriguing example of this cross-scale redesign cycle arose in reconciling modeled total population histories for entire settlements with individuals' modeled childbirth probabilities compiled from diverse ancient textual sources (Coale & Demeny 1966). We will return to this example later in this chapter.

A second, tenet of the MASS modeling philosophy, complementary with the bottom-up approach, is the use of an *agent-based modeling and simulation* (ABMS) approach as the foundation for the MASS simulation framework's representation of social behavior (Epstein & Axtell 1996). In common with several other simulation studies that have addressed societal sustainability in a dynamic natural environment (e.g., Lansing & Kremer 1993; Doran *et al.* 1994; Kohler *et al.* 2000), the simulation's decision making entities, represented by persons, households, and other types of organization, are portrayed as bounded, rational agents (Simon 1982). Thus, each person, household, or other organization in a simulation governs its own behavior based on its own local rules and in response to its own perceptions, preferences, capabilities, and goals.

A third MASS modeling principle has been the enforcement of a balanced representation of both environmental and societal processes, with neither component dominant at the expense of the other. This has required a corresponding balance in the levels of detail and fidelity in the modeling representations of landscape and social themes. Specifically, both the internal structure and the dynamic behaviors of the modeled elements of the ancient Mesopotamian subject domain must be represented in such a way that it is indeed possible for the actions of social 'agents' to influence the structure and subsequent dynamics of portions of the natural environment, and, in turn, for the state and dynamics of an agent's local environment to influence the agent. In practical terms, if, for example, household work crews and individual persons, herds and individual animals are to be empowered in the simulation to have an impact on their natural environment, then the following factors must be taken into account:

1. The salient mechanisms by which agents of environmental change such as work crews, persons, herds and individual fauna affect their natural surroundings must be explicitly represented in the simulation via sub-models of those aspects of the agents' behavior. For example, a plow team working a plot of land should be capable of altering the physical structure and texture of the upper layers of soil in that plot, and a sheep grazing in a pasture plot should be capable of removing standing biomass or converting it to surface litter (by grazing and trampling), and altering the nutrient balance of the plot's soil (by elimination of wastes and natural manuring).
2. The modeled representation of the natural environment must be resolvable at spatial and temporal scales appropriate to the specific modalities by which the work crews, persons, herds and individual fauna produce their impacts. Thus, in terms of the previous examples, the simulated landscape should be capable of being partitioned into relatively small plots to allow appropriate localization of agent impacts, and should include some defining parameters that can be altered by agents' activities.
3. To complete the impact characterization and to support explicit representation of feedback processes from the environment back to the societal component, the modeled representation of the natural environment should incorporate dynamic models of landscape evolution that are sensitive to the types of impact (as in Factor 1, above) that the agents of society can generate. This means that the evolution must depend to some degree upon parameters, such as soil structure and biomass spatial density, that are alterable by agents' actions.

The MASS simulation software framework that has emerged from several years of collaborative development effort by the MASS Group adheres to all of these principles. It embodies a contextually rich, balanced, microscale treatment of both societal and environmental processes, without recourse to macroscale control mechanisms. Social processes are portrayed in accordance with the ABMS paradigm, with agency residing principally at the level of households and individual persons. Direct interactions and process feedbacks between societal and environmental processes are explicitly supported across the entire spectrum of modeled activity.

OVERVIEW OF THE MASS SIMULATION FRAMEWORK DESIGN

We have chosen 'ENKIMDU' as a shorthand name for the MASS Project's simulation software framework, in honor of the ancient Sumerian god of that name. Enkimdu was a god of agriculture and of irrigation, and thus seemed a fitting namesake for our simulation system, with its focus on the agricultural sustainability of ancient Mesopotamian social institutions.

ENKIMDU simulations have addressed natural (weather, crop growth, hydrology, soil evolution, population dynamics, etc.) and societal (farming and herding practices, kinship-driven behaviors, trade, etc.) processes interacting on a daily basis across multi-year to multi-generation time spans. Entities within ENKIMDU's ancient Mesopotamian study domain are resolved and modeled at the level of individual persons and households, and individual cropped/fallowed fields, herds, and flocks. The modeled scenarios can encompass spatial scales and granularities that range from single households and their individual members and resources, to individual hamlets, towns with supporting satellite villages, and regional constellations of such towns arrayed into petty kingdoms. The spectrum of temporal scales addressed is similarly broad, ranging from resource utilization and task prioritization of daily household activities up through multi-year regional climate variations and decades-long evolution of dynastic politics.

Christiansen and Altaweel (2006a) have attempted to situate the MASS simulation framework within the continuum of existing modeling and simulation software tools. ENKIMDU may be classified as (a) object-based, (b) agent-based, (c) a discrete event simulation, and (d) both a simulation and a simulation framework. In brief, the key implications of the inclusion of ENKIMDU in these classification categories are as follows:

a. *ENKIMDU as an object-based application.* The subject domain representation is decomposed into individual software 'objects' to organize and simplify program data flows, an approach pioneered in the SIMULA software language (Dahl & Nygaard 1966). In ENKIMDU, each such software object represents a real-world entity, such as an agricultural field, a goat, a household, a plow, etc. In object-based software terminology, these objects are *instances* of an object *class*. An object class definition carries the blueprint for a particular *type* of entity, specifying what baseline attributes every instance of that class will possess, but without prescribing the values that these attributes will take on for specific instances of that object class. For example, two plots of land may each be represented in a simulation by a Field object: an instance of the Field object class. By virtue of membership in the Field class, each instance will have the same list of attributes, such as area and vegetation cover. But since each Field represents a different, specific real-world, one might be twice as large as the other, and one might be lying fallow while the other supports a crop of barley. Under the object paradigm, each object instance is responsible for maintaining and controlling access to the values of the attributes that define it, and is responsible for expressing any and all of the dynamic behaviors that have been defined as appropriate for objects of its class.

b. *as an agent-based software application*. The ENKIMDU software objects that are used to represent real-world societal entities with decision making power, such as individuals and households, are modeled as bounded rational agents in accordance with the ABMS paradigm; they each make decisions based on imperfect, localized perceptions of their own particular surroundings and situation.

c. *as a discrete event simulation.* Communications among agents and the onset, progression and completion of dynamic processes in a MASS simulation are delineated by time-ordered discrete events that can occur arbitrarily close together in simulation time, rather than being constrained to equal intervals (e.g., days) as in a time-stepped simulation. This facilitates concurrent representation of different processes with widely varying tempos and durations.

d. *as both a simulation and a framework.* Although ENKIMDU's original *raison d'etre,* and its current principal area of application, is analysis of sustainability issues for ancient Mesopotamian settlement systems, the simulation was designed and built in accordance with standard best practices for object-oriented software development, to exploit commonalities in data structures and processes. The result is that its basic software structure and library of fundamental societal and natural landscape entities is applicable to essentially any social and/or ecological subject domain, so that it can be used as an underlying framework to support diverse new simulations. ENKIMDU has been extended to address village-level agroeconomics concerns in modern Southeast Asia, among other milieus.

The ENKIMDU framework has been constructed with the aid of software infrastructure provided by three Argonne-designed modeling, simulation and visualization framework systems:

- *DIAS, the Dynamic Information Architecture System*: An object-based generic framework for building and maintaining complex multidisciplinary simulations. (Christiansen 2000a);
- *FACET, the Framework for Addressing Cooperative Extended Transactions*: An object-based framework for building complex, expressive agent models of social processes. (Christiansen 2000b); and
- *JeoViewer*: An object-based geospatial display and analysis toolkit that interfaces easily with simulation models and accommodates all common geospatial data interchange formats (Lurie *et al.* 2002). JeoViewer provided the means of visualizing the progress and detailed dynamics of ENKIMDU simulations till 2008. At that time, it was replaced by a more generic approach based on the Java2D (Sun Microsystems 2006) open-source visualization software system.

The ENKIMDU simulation system and all three of these enabling frameworks are implemented in the standard Java programming language (Gosling & McGilton 1995).

For multidisciplinary simulations as large and complex as ENKIMDU, use of the DIAS framework as a foundation has brought with it some significant advantages. As has been noted, the object paradigm in computer science calls for software objects to express their own behaviors in response to what information is available to them, rather than relying on some form of overarching central control mechanism to trigger the behaviors[1]. The DIAS framework supports a particularly flexible and powerful extension of the object paradigm, in which objects may invoke simulation model code that may exist anywhere on the network, as long as that model code is 'registered' with the DIAS system as being capable of addressing the specified aspects of those objects' behaviors. Just which model code is chosen to represent a behavior can be determined on the fly during a simulation, as a function of the simulation context. The corresponding model code can have been written especially for the specific modeling problem of interest, or may be part of a proven, existing simulation implemented in essentially any computer language. The DIAS framework supplies the facility to encapsulate such 'foreign' models in Java 'object wrappers' so that they may converse with the other objects representing the subject domain being simulated. This flexible, *ad hoc* form of object-to-model linkage is made possible by requiring that each model be able to express its input needs and output capabilities solely in terms of attributes of types of objects already defined within the DIAS-supported simulation. Each such model converses with the simulation in the language of the relevant subject domain object attributes. Thus models may interact with the domain objects that their owner objects 'know about,' but they never need to interact directly with other models. This approach pays major dividends in scalability when the desired scope of a simulation expands over time, as it almost inevitably will. As more and more process models are added to a DIAS-supported simulation, there is no need to maintain an exponentially growing set of model-to-model linkages and data protocols, because the models have been relieved by DIAS of the responsibility of having to know about each other – they need only know about the subject domain objects that provide their inputs, and that restricted data neighborhood will tend to remain relatively stable.

As required by the DIAS framework, input and output parameter values needed for the simulation models used by ENKIMDU must be drawn from among the attributes of subject domain objects that are defined within the ENKIMDU domain object model. The subsets of domain object attributes that are used as model inputs and outputs may easily overlap from model to model. Such overlaps provide a means of expressing heterogeneous, cross-domain and cross-discipline feedback effects among the dynamic processes represented by the models,

[1] This is tantamount to a more general and inclusive form of the *agent-based* paradigm, in which essentially <u>all</u> software objects in a subject domain can be considered to be agents. For purposes of this discussion, however, we will bow to convention and reserve the term 'agent' for societal entities that have decision making capabilities.

one that arises naturally from the process of selection and defining the objects and models that make up an ENKIMDU simulation.

Interactions and potential feedback loops involving multiple concurrent dynamic processes will therefore ensue implicitly wherever such overlaps occur, and need not be explicitly built into the simulation design. For example, ENKIMDU's present herd foraging model requires as an input the standing plant biomass density of a pasture object (specifically identified as the landscape partition object occupied by the herd). When the model runs, it represents the effects of the herd's foraging for food for some number of hours by signaling to the pasture object that a given amount of standing biomass should be removed (eaten) and a different amount should be converted to surface litter (trampled), and that a specified amount of manure should be deposited across the surface of the pasture. Meanwhile, an ENKIMDU landscape evolution model operates throughout the simulations to compute daily changes in soils, moisture and vegetation across all landscape partitions in the region of interest. The inputs of the models include, among other things, the initial standing biomass and surface litter densities, as well as any changes in soil chemistry occasioned by additions of natural or chemical fertilizer. As a result, when the landscape evolution model is triggered for its next daily update, the subsequent state of the pasture partition will be influenced by that day's visit from the herd; this effect will propagate ahead in time so that any future grazing *or farming* attempts in the landscape partition will be affected as well.

MASS Software Representation of the Ancient Mesopotamian Settled Landscape

The data sources, assumptions, and approaches used to populate the MASS simulations with appropriate data were agreed upon jointly by the MASS Group membership, typically in monthly meetings, as noted above.

The MASS Domain Object Model

The term, 'domain object model', or DOM, is commonly used to refer to *'the specification of (logical) objects with direct (physical) counterparts in the work space or problem domain.* (Villers & Sommerville 2000). Defining the DOM is a pivotal step in any object-oriented simulation software project.

In discussing the MASS DOM and related topics, we employ the standard Java language object modeling nomenclature (Sun Microsystems, 1999) that is used in the ENKIMDU simulation software system:

- *Object class names* are nouns or noun phrases, in mixed case with the first letter uppercase and the first letter of each internal word capitalized.

  ENKIMDU examples*: Person, Household, LandscapePartition, GrainLoanAgreement.* Under the DIAS modeling paradigm, simulation models are also represented as software objects – thus, for example, *FieldCropManagementModel, DisposeHarvestModel*.

- *Object class attributes* are nouns or noun phrases, in mixed case with the first letter lowercase and the first letter of each internal word capitalized. Attributes of specific object classes can be denoted by concatenating the class name with a period and then the attribute name:

  ENKIMDU examples*: StateGranary.grainStores, Settlement.annals, Device.owner.*

- *Object class methods* (the standard means by which objects express their behaviors) are generally verbs or verb phrases, in mixed case with the first letter lowercase and the first letter of each internal word capitalized.

  ENKIMDU examples*: Community.assignTodaysPastures, Household.addFoodToLarder, Fauna.canBreed.*

- In the Java language, object classes can be grouped into 'packages.' Classes grouped inside the same package have an additional special level of access to each other. Object classes residing in specific packages can be uniquely identified by concatenating the package name (all in lowercase) with a period and then the class name. In addition, packages can be hierarchically nested; this additional organizational context is denoted by concatenating package names together, top to bottom, separated by periods

  ENKIMDU examples*: Enkimdu.entities.Person, Enkimdu.domain.transport.Route, Enkimdu.model.ManageHerdLocalModel.*

In Java, names do not allow blanks to separate words, though underscore characters are sometimes used to achieve the same result. In the preferred standard Java naming convention, individual words within a longer, compound name are capitalized to make it easier to visually parse them.

We will follow these naming conventions throughout the chapter, to help make it clear when we are discussing a simulation software objects as opposed to their real-life counterparts. Thus a 'Person' is a simulation object that represents an individual person. When we need to describe the characteristics common to an entire *class* of objects, we will refer to, for example, 'the Person class.'

Structure and Dynamics of the MASS Domain Object Model

The Domain Object Model for the MASS project attempts to encompass the key structural components and dynamic elements of the settled landscape. Our software representation of this domain is comprised of software objects representing the elements that make up the real-world landscape. In addition, they include objects that encapsulate the simulation model functionality that expresses the dynamic behaviors of these objects in the ENKIMDU simulations. In the following discussion we identify and describe the salient object classes in our ancient Mesopotamian modeling domain, and discuss the principal modeled behaviors of these objects and the simulation models that have been built or adapted to represent these object behaviors within the ENKIMDU framework.

The simulation system prototype as it currently stands contains formal Java class definitions for over 400 Java language object classes specific to the ENKIMDU application. These object classes represent the entities that comprise the DOM: a fine-grained view of ancient Mesopotamian society and its infrastructure, embedded in its natural environment. This collection of domain objects represents both concrete instances (e.g., fields and plow teams) and abstractions (e.g., perceptions and plans). In addition, ENKIMDU is supported by the hundreds of software object class definitions that make up Argonne National Laboratory DIAS and FACET system software libraries and the relevant Argonne general-purpose modeling, simulation and spatial analysis utility software libraries.

Figure 10.1 presents a simplified view of the entire ancient Mesopotamian settled landscape of interest, broken down into conceptual entities that are implemented in software object forms in the simulation framework's DOM. Many of the key inter-object linkages are also shown: required connections are depicted by heavy blue lines, and incidental though often extremely important connectivity among objects is indicated by lighter green lines. Figure 10.1 can be thought of as a triptych made up of the dominant themes of *Person*, *Community*, and *Landscape*, running left to right across the figure. We will discuss the components of the DOM and their behaviors and interactions in the framework of this triptych of themes.

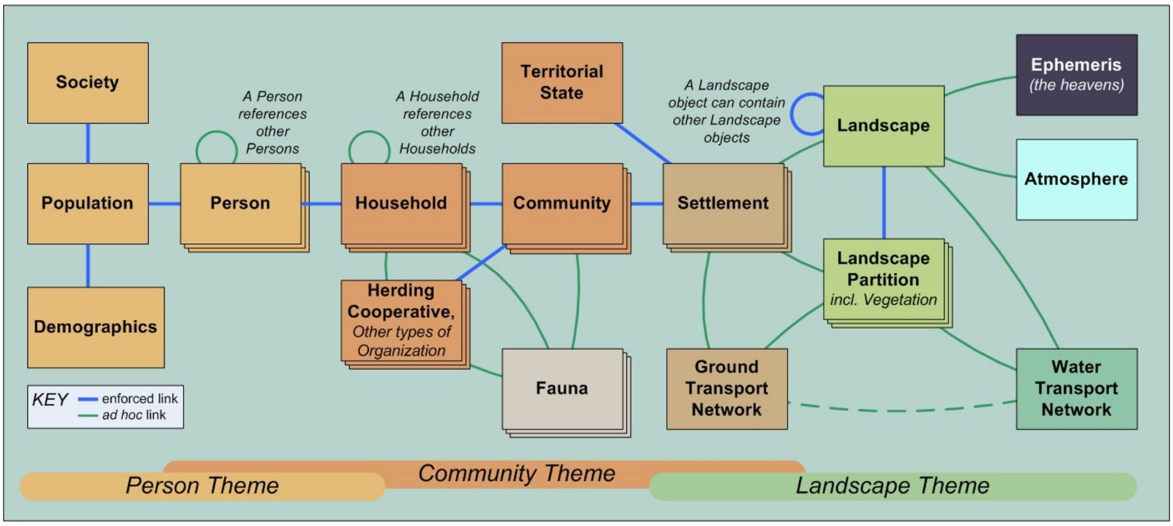

Fig. 10.1 Key software object classes in the MASS ENKIMDU domain object model.

Figure 10.1 calls out domain entity types and hints at cases for which there may be multiple instances of a given entity object class in the simulation (shown as *stacks* of objects in the figure). However, since it is essentially a structure diagram it cannot paint a clear picture of the scale and scope encompassed in ENKIMDU simulations, in terms of the sheer numbers of software entities involved. Thus it may be helpful here to briefly characterize the simulations that have been undertaken for ancient Mesopotamian studies in terms of the numbers of domain objects of each class that were employed. For a simulation of a single medium-sized settlement, the following roster of objects is typical:

The Upper Khabur-Type Model Settlement

Data to characterize this model settlement are drawn from excavations and surveys from a range of sites in the Upper Khabur basin in NE Syria. These include, the 22.5[2] hectare Tell Beydar site and the surrounding landscape, as well as other sites in the region. The MASS simulations for the model site have generally assumed an initial population of 500 or more, supported by a halo of agricultural fields extending to a distance of roughly two kilometers from the site center, as shown in Figure 10.2. In the figure, the settlement's agricultural lands present a mosaic of over 300 crop field patches, each around 3 ha in area, which avoids the relatively barren basalt plateau to the west of the settlement site.

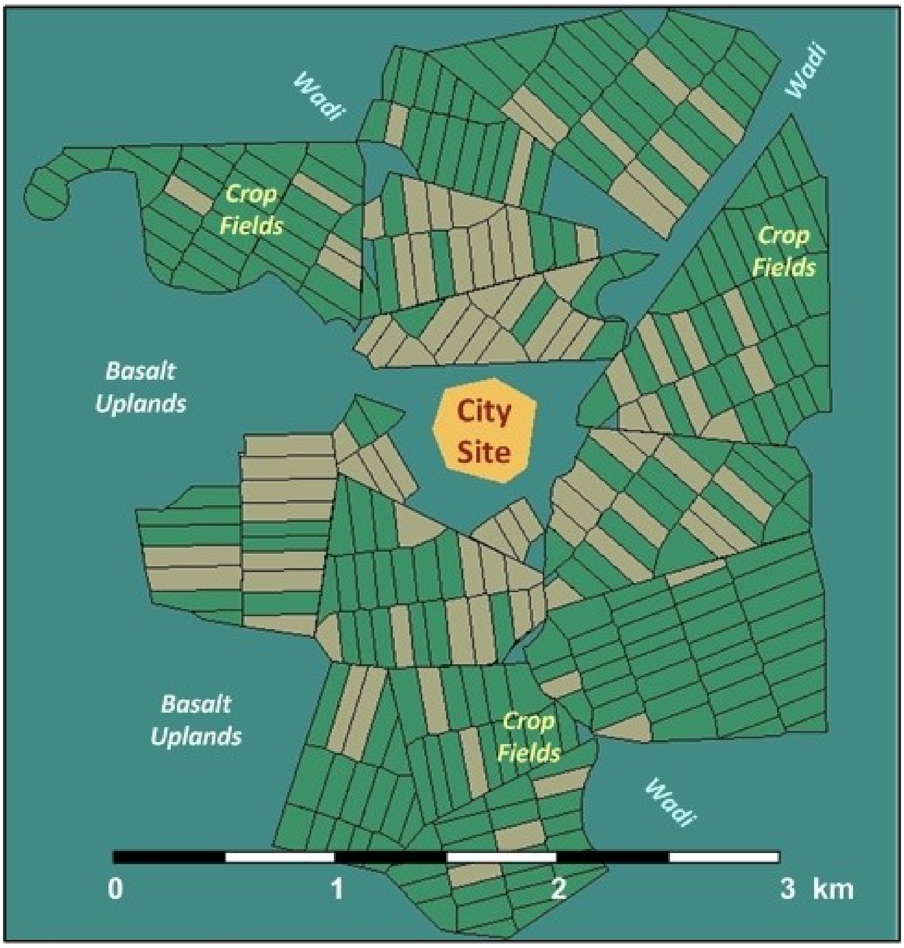

Fig. 10.2 The Upper Khabur model settlement layout.

The field mosaic was synthesized by the Landscape Partition Engine (LPE), an ENKIMDU system preprocessor, based on:

[2] In the case of Tell Beydar, 22.5 ha is the total area of the Early Bronze Age site; the settled area is approximately 17 ha.

(a) a full specification of the geometrical map projection scheme, in this case the Universal Transverse Mercator system (Snyder 1987), to be used in mapping site area locations;

(b) the extent of the field area to be subdivided, expressed as a mean radius from the settlement center to the outer boundary of the field system;

(c) the locations of significant local wadis;

(d) a soil map for the site area based on regional survey data from the Upper Khabur valley (van Liere 2003); and

(e) a handful of 'spatial texture' parameters governing the palette of field sizes, geometries, and orientations to be generated.

Spatial texture metrics were selected in order to approximate the appearance of field mosaics for existing archetype settlements for example at the village of Qaraqosh in Iraq, and to reflect constraints imposed by agricultural techniques and technologies – e.g., relatively long, narrow fields to allow room to turn plow teams at the ends of furrows (see Chapter 4).

Approximate Number of Objects in the Simulation: Upper Khabur Model Settlement

- One *Ephemeris* object, used to express the relevant behaviors of the solar system (times of sunrise and sunset, etc.).
- One *Atmosphere* object, providing daily surface weather for the region of interest.
- One *Landscape* object, representing the settlement and its agricultural surrounds as a single entity.
- 300 to 500 *LandscapePartition* objects within the Landscape, the vast majority of which represent the lone settlement's agricultural fields. These partitions are small, each on the order of 3 ha in area.
- One *Hydrosphere* object, representing the local ground water and surface water state (drainage networks and ephemeral surface ponding, both inadvertent and due to irrigation impoundment, if appropriate). The Hydrosphere's WaterTransportNetwork object may or may not be present, depending on the importance of major surface water channels in the specific case investigated. For the simulations employed in this volume (see also Christiansen & Altaweel 2006b, Wilkinson *et al.* 2007), explicit surface water channels were neglected. Surface water transport is implicitly dealt with in the framework, as discussed later in this chapter under the *Landscape* theme.
- One *Settlement* object, e.g. the nucleated Upper Khabur type settlement. The Settlement contains hundreds of additional objects, as discussed later under the *Community* theme.
- One *Community* object representing the ensemble of persons, ranged in households and other types of organization, that populates the Settlement. (In ENKIMDU, Settlements are inanimate, infrastructure objects, that is simply piles of mud-bricks; their resident Community objects, which represent purely societal constructs, give them life.)
- One *GroundTransportNetwork* object, comprising a complete movement network for foot, mounted and vehicular traffic within the village, and to and from all of the agricultural fields and pastures of the village. Such networks contain on the order of 3,000 nodes and links.
- Optionally, one *TerritorialState* object, to permit representation of interactions, such as tax assessments and public works projects, between the settlement and a higher-order, controlling polity.
- 100 to 150 *Household* objects representing the households in the village. Household objects have been the central actors, or cognitive agents, in the MASS simulations, serving as the hubs for the bulk of the simulations' resource allocation activities, interactions with the dynamic landscape, as well as social interactions of all types.
- 500 to 800 Person objects, each belonging to a Household and connected via kinship and other relationships to Persons in that and other Households.

- 600-1,500 Fauna objects. The preponderance of these are sheep and goats belonging to individual Households; others represent draft and pack animals.
- Four to twelve *HerdingCooperative* objects, each representing a joint effort among a number of households to combine their livestock holdings into larger flocks for purposes of shepherding the animals to and from local pasturage each day.
- One *Population* object, linked with every Person in the simulation, with one associated *Society* object and one associated *Demographics* object. This represents a single assumed sub-population of people with common ethnic background and cultural outlook. This design supports convenient differentiation of the modeled population into sub-populations with varying demographic and ethnic characteristics and cultural norms. Population, Society and Demographics class objects are discussed later in this chapter under the *Person* theme.
- Optionally, one *NomadCommunity* object, presenting opportunities for the settlement residents to interact intermittently with another community that in general has different production specialties, stocks of commodities, wants and needs.

Scaling up to a larger modeling region, we consider the North Jazira Survey Region, extending across parts of modern northern Syria and Iraq (Wilkinson & Tucker 1995).

The North Jazira Modeling Sub-Region

The MASS Group chose to simulate a 40 x 20 km. area within this region, containing some 45 settlements ranging from small hamlets to significant towns. To obtain a plausible distribution of settlement sites and sizes, the modeled settlements were chosen by the MASS Group from among Ubaid-era sites in the region that were considered likely to have supported active settlements at the same time. Figure 10.3 shows the modeling representation of the sub-region.

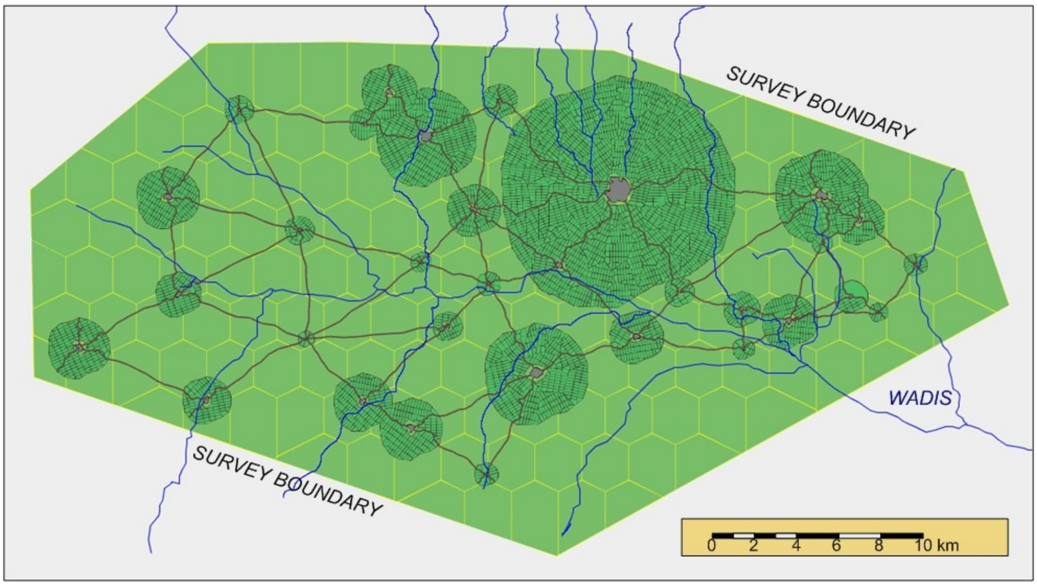

Fig. 10.3 North Jazira sub-region model layout.

As was the case with the modeled Upper Khabur landscape, the North Jazira sub-region model landscape geometries, as depicted in Figure 10.3, are generated automatically by ENKIMDU's LPE preprocessor. In this instance the inputs to the LPE consists of:

(a) map projection specifications, as above;

(b) a polygon representing the perimeter of the overall sub-region to be addressed;

(c) the geometries of the network of permanent and ephemeral watercourses (wadis) in the sub-region;

(d) polygons representing the measured extents of the formerly settled areas for the 45 selected sites;

(e) initial modeled population (optional input for each settlement);

(f) exploratory decisions regarding the likely density of the inter-settlement ground transport network in the sub-region; and

(g) spatial texture parameter settings as above.

The initial extent of the halo of agricultural land surrounding each modeled settlement is computed based on the initial population of the settlement. This is (necessarily, given the paucity of data) a crude approximation based on carrying capacity considerations. The radius of the 'agricultural catchment,' the farming land that sustains the settlement, is computed as follows.

- Assume biennial fallowing practice is rigorously observed
- Assume local land is 80 percent arable (includes an exclusion zone for the settlement site itself)
- Assume mean grain (barley) yield is 750 kg/ha (rounded up from FAO data of 740 kg /ha[3]: Deckers & Riehl 2008 tables 1 & 2), and conservatively assume that grain provides 100 percent of the settlement residents' sustenance.
- Conservatively assume per capita grain consumption is 250 kg/person (Wilkinson *et al.* 2007, footnote 3).
- Allow for 50 percent population expansion without adding crop field area.

Then on average each person will require two thirds of a hectare of arable cropland for sustenance. Multiplying by a factor of 1.25 to account for the fraction of land that is arable, and a factor of 1.5 to accommodate settlement expansion, each person in the initial population will need 1.25 ha of land. Using the simple formula for the area of a circle, we find that the radius, r_s, of a settlement's agricultural catchment is:

$$\mathrm{r(km)} = 1.5 \times \sqrt{0.00469 \times (\text{initial population}) \div \pi}$$

The LPE preprocessor constructs agricultural catchments for each settlement as 'lumpy' circles of mean radius r_s computed for that settlement, and fills each of them with individual fields using its spatial texture parameter values. The settlement sites themselves, plus an automatically derived buffer zone around them, are excluded from the agricultural landscape partitioning process.

Unless an initial population is explicitly specified for a settlement in the simulation initialization data, a settlement's initial population is computed automatically by assuming a constant population density in persons per hectare of site area, and multiplying that by the site area in hectares. Unfortunately, because there is no single figure that can be used for site population densities, past estimates have been based on a plausible range of figures. For most of the MASS simulations to date, nominal population densities of 100 to 200 persons per hectare of settlement site area have been assumed (Wilkinson *et al.* 2007: 56). These figures compare with those offered by Modelski (1997) although for larger cities higher figures of 200 persons per hectare and higher are more likely (further discussion is found in Chapter 6). The MASS Landscape Partition Engine does more than subdivide landscapes into field mosaics for the modeled settlements. The semi-automatically generated site and landscape layouts depicted in Figures 10.2 and 10.3 also embody complete, highly detailed transportation networks that are generated concurrently with the partitioning of the landscape. As is further discussed later in this chapter, the LPE creates a constellation of interlinked ground movement networks that operate across a number of spatial scales:

[3] Note that in earlier simulations we used different figures, e.g. yields of 500 kg/ha in Wilkinson et al. 2007 footnote 3. The above figure is just an approximate figure used for the laying out of the fields

- *Regional Landscape Scale Network*: generally sparse network of regional, inter-settlement roadways.
- *Settlement Landscape Scale Networks*: extensive networks made up of all crop field boundaries for the mosaic of crop fields surrounding each settlement, interpenetrated by the regional roadways.
- *Settlement Scale Networks*: network of connector pathways linking each settlement's urban network with the surrounding network, through a 'suburban' belt comprised of variable land use plots situated between the settled area and the crop field system.
- *Urban Neighborhood Scale Networks*: dense networks of pathways within the built-up areas of each settlement, linking individual dwellings and other structures along urban lanes and interior courtyards.

Taken together, this multi-scale system of networks connects the front doors of every household's dwelling in the settlement, and can support highly detailed social behavior models.

Water transport networks are also supported in the MASS simulations, though they are not significant factors for the upper Khabur and North Jazira landscapes, which are both set in northern Mesopotamia. However, waterborne transport is a crucial component of southern Mesopotamian modeling scenarios.

Approximate Number of Objects in the Simulation: North Jazira Modeling Sub-Region

- One *Ephemeris* object.
- One *Atmosphere* object, providing spatially as well as temporally varying surface weather across the area.
- Roughly five to 30 *Landscape* objects, each representing a swath of land encompassing one or more settlements and their agricultural and pastoral lands. In this regard a Landscape may be equated in hydrological terms to a watershed or basin, or a portion thereof. The decomposition of a large modeled area into several Landscape objects serves two distinct purposes, one related to the domain science and one owing to strictly computational concerns:
    - Recourse to use of multiple Landscapes acknowledges and permits explicit simulation of different hydrologic divisions (watersheds, or basins) extant in the modeled region.
    - Multiple Landscapes can provide a convenient decomposition of the modeled hydrosphere so that different processors in a cluster can be assigned to computation of water balance for different Landscapes, in order to improve overall speed of simulation.

  These issues are elaborated upon in a detailed discussion of Landscape object structure and dynamics to follow.
- 25,000 *Landscape Partition* objects, divided up among the defined Landscapes.
- One *Hydrosphere* object, addressing ground and surface water state and transport for all of the modeled Landscapes. One responsibility of the Hydrosphere objects is to deal with ground and surface water exchange between adjacent Landscapes.
- 45 *Settlement* objects, each with a *Community* object representing the inhabitants of the settlement and their local social organization.
- 46 or more *GroundTransportNetwork* objects, comprising a complete movement network for foot, mounted and vehicular traffic within each village, to and from all of the village's agricultural fields and pastures, and between different settlements across the region. These multiple networks, if connected via common nodes (as they almost always are) function precisely as would a single unified network. The main advantage of portraying the movement networks as a patchwork of local systems is that it supports agent reasoning on the concepts of local versus regional transport options. The ground transport network taken as a whole contains close to 100,000 nodes and links.
- *Water Transport Network* objects do not come into play for North Jazira simulation scenarios. The MASS Group did not consider waterborne transport likely to have been a major factor in this relatively

dry patch of northern Mesopotamia. For southern Mesopotamian simulation scenarios, on the other hand, waterborne transport is vital (see Chapter 9), and the suite of WaterTransportNetwork object classes in the ENKIMDU simulation framework assume great importance.

- Optionally, one to perhaps four or five *TerritorialState* objects. These create opportunities for modeling to represent diplomacy, warfare, and other modes of interaction among higher-order polities, in addition to internal measures of the state, such as shifting food reserves to deal with localized famines.
- 6,000 *Household* objects.
- 30,000 *Person* objects.
- Upwards of 50,000 *Fauna* objects.
- 300 *Herding Cooperative* objects, each representing a joint effort among a number of households to combine their livestock holdings into larger flocks for purposes of shepherding the animals to and from local pasturage each day.
- One or more *Population* objects, each representing a single subpopulation of people with common ethnic background and cultural outlook.
- Optionally, one or more *NomadCommunity* objects.

Our overview of software object classes has been couched in terms of the static structure of object classes, and the relationship of each class to other object classes in the DOM. Figure 10.4 attempts to summarize, in a single image, the range of diverse dynamic behaviors that this mix of objects, and the simulation infrastructure that supports them, can produce.

Figure 10.4 is a simplified logarithmic scale chart depicting the spectrum of modeled domain object behaviors in MASS simulations, and the range of temporal scales at which each type of behavior is manifested. For example, it can be seen from the figure that each Person object in a simulation may participate in various communal activities (on behalf of its household, kin-group or community) over timescales of from a few minutes to several weeks. At the high end of the temporal scale, a simulated person might be committed for months to the ongoing task of cultivating a barley crop in one of his household's fields. In the course of meeting that commitment, he may be called on by his Household to perform specific sub-tasks, such as creating an opening in a dike to allow irrigation water to enter a field, that can be accomplished in a matter of minutes.

The left-hand side of the figure acknowledges a key DIAS-supported simulation design feature touched upon earlier: behaviors of domain objects in ENKIMDU are implemented by simulation models that can exist independently outside of the software framework.

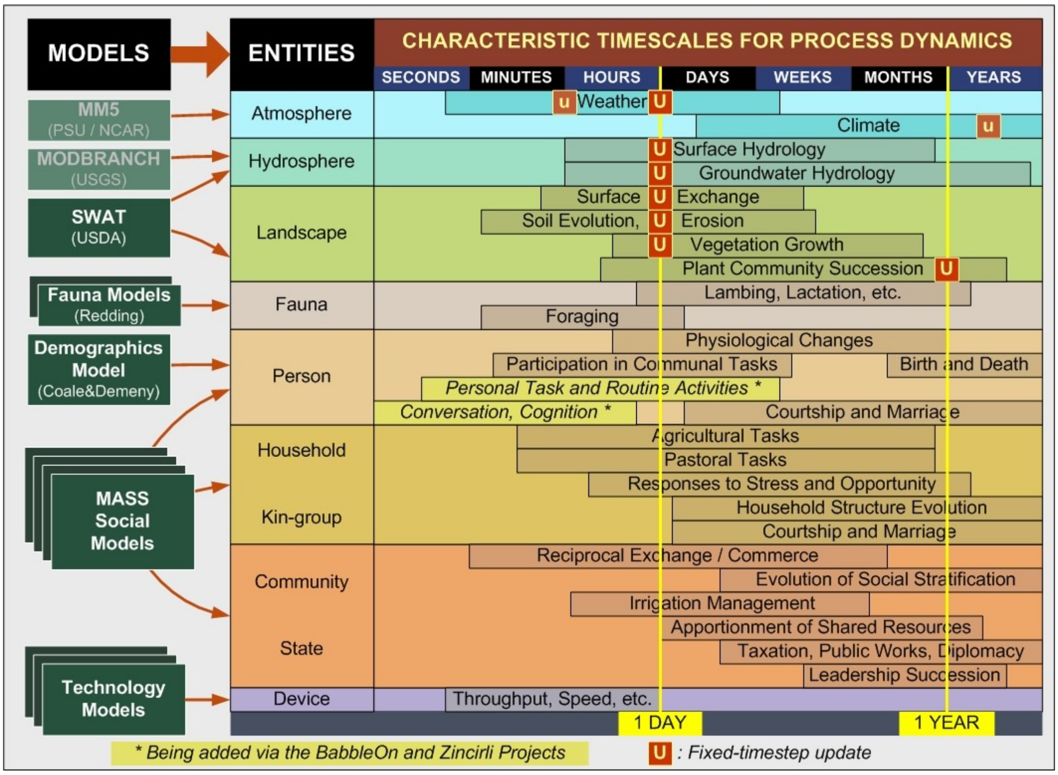

Fig. 10.4 MASS ENKIMDU simulation's domain object dynamic behavior.

The extremely broad range of dynamic process timescales that the ENKIMDU system can address is largely due to its employment of a discrete event simulation approach, in which modeled processes can begin, update, and end at whatever times are appropriate to them. The majority of the dynamic behaviors represented in the system are modeled in this way. However, certain domain behaviors are better represented through the use of well-established, existing models that operate and report results for fixed time-steps. For such models, occurrences of the letter 'U' enclosed in a square in Figure 10.4 point out fixed update time intervals for certain processes: usually these are one day, but in some cases can be one year or some other interval. These include the models for weather and climate, hydrology, as well as vegetation and soil evolution. Many of these aspects are currently modeled using the U.S. Department of Agriculture's Soil and Water Assessment Tool (SWAT) model (Arnold *et al.* 1998), which operates on a fixed daily time-step. ENKIMDU emulates SWAT's timing protocol by scheduling update events for soil state and so on that are spaced precisely 24 hours apart. SWAT and other specific models identified in Figure 10.4 will be more fully discussed later in this chapter.

With a high-level perspective on the structure and dynamics of the MASS simulation framework established, it is time to dig deeper into the characteristics and capabilities of the logical elements that make up the ENKIMDU simulation software framework. We will discuss in turn each of the 'themes' identified in Figure 10.1: Person, Community, and Landscape, along with important ancillary simulation elements, such as those dealing with technology and infrastructure, and will attempt to situate these diverse components in the overall context of an holistic simulation.

MASS ENKIMDU Domain Object Model: The Person Theme

Figure 10.5 highlights the portion of the full MASS ENKIMDU domain object model (Figure 10.1) that focuses on individual persons in the simulations. Only the most important object classes are shown here.

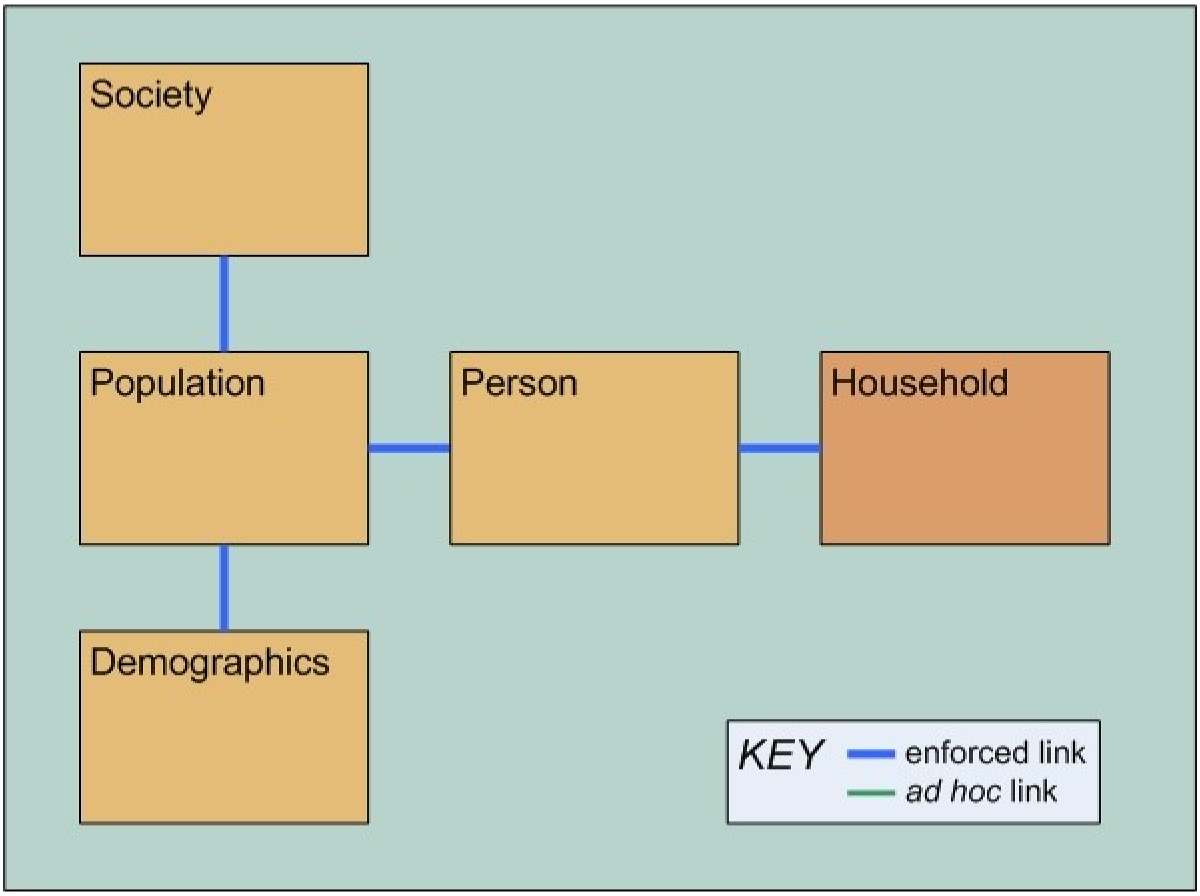

Fig. 10.5 MASS Domain Object Model: person and related object classes.

The Person Object Class

Each Person object represents one individual. Each Person carries links to parents, all spouses/partners, and all offspring, living or dead. This makes it possible to trace all kin relationships among modeled Persons, even when some of the Persons in a chain of kinship linkages are deceased. The simulation maintains basic genealogical information on each Person even after their sim-deaths, until such time as no living Person has a link to them anymore. Person objects presently modeled by ENKIMDU carry roughly 50 attributes, of which the most important are as follows:

- Demographic information:
  - Date of birth (from which age can be derived), gender.
  - Population or micropopulation to which the Person belongs, from which may be determined the Person's ethnic and cultural background, as well as fertility and mortality metrics.
  - Kinship information: household; lineage; father, mother, list of siblings; list of offspring.
- Cognitive information:
  - Perceptions, operations (ongoing and planned activities and commitments), agenda and schedule of activities; personal strategies[4].
- Social information:
  - Personal resources (owned or controlled) in detail.
  - List of social roles that the Person may take on; Person's *current* role. Roles are encapsulated in a list of Role objects maintained by each Person. Roles tend to depend on age group and

[4] Note that in most ENKIMDU simulations to date, the bulk of the cognitive behavior representation has been vested with heads of households, and have been maintained within Household objects, rather than with individual Persons; as a result, the cognitive gear associated with individual Persons is lightweight compared with that of Households.

gender. For ancient Mesopotamian simulations, the MASS Group selected five age groups for the models to recognize: Infant (age 0-1 yr.), Child (age 2-9), Adolescent (age 10-13), Adult (age 14-59), and Elder (age 60 and up).
  - Principal language spoken (used only in computational linguistics extensions of ENKIMDU).
  - Person's social status in a stratified society.
- Basic simulation information: Person's current location, activity; and physiological state.

In the MASS simulations, Persons can be born, die, eat, become fatigued or ill, recover from fatigue or illness, give birth, engage in courtship and be married, travel, perform a wide variety of tasks, particularly subsistence tasks, and so on.

The Population Object Class

An ENKIMDU Population object represents a single micropopulation of Persons. Attributes of Population objects include a list of all Persons belonging to the Population, along with references to instances of the following object classes:

- Society (see below)
- Demographics (see below)
- Dominant Language (for recent computational linguistics extensions of ENKIMDU)

Instances of the Population object class have no dynamic behaviors of their own, but their attributes, society and demographics, can have a profound impact on simulation dynamics.

The inclusion of a complex Population attribute for each simulated person allows our simulations to support multicultural, socially heterogeneous simulations. The DOM's ability to articulate micropopulation membership for each individual in the simulation, with accompanying overtones of ethnicity and culture, as well as physiological/ demographic implications and linguistic differentiation, provides a rich environment for social agent modeling.

The Society Object Class

The Society object class provides access to models of normative social behaviors pertaining to a specific cultural grouping. Persons and households (via their household heads) refer to the Society associated with the Population to which they belong (see Figure 10.5) for guidance when participating in such a social behavior. Thus in a simulation a Person may be performing an action on behalf of her Household, and according to a cultural pattern established by her Society. A few of the most important of these Society-specific behavior models include:

Courtship and Marriage:

- EligibleMale, EligibleFemale. These models determine socially acceptable matches. Note that in the MASS simulations to date we have conservatively assumed monogamous relationships only.
- AccommodateNewlyweds. Defines mechanisms for introducing new household members (generally the new wife and possibly some of her kin, for ancient Mesopotamian simulations) and associated resources into an existing household.
- CanStartNewHousehold / EstablishNewHousehold: The CanStartNewHouseohld model determines whether a new, independent household can be created from an established household. In order for this to occur, the 'parent' household must be sufficiently well off, in terms of current and projected future labor, food and other resources, to be willing to risk parting with a substantial fraction of its wealth to start the new household off with a decent chance of success. The EstablishNewHousehold model executes any household fissions dictated by the CanStartNewHousehold model.

Kinship-Based Coping Mechanisms:

- AskCloseKinToShareFood / OfferFoodToNeedyKin. These models regulate nonreciprocal exchanges of grain at need along kinship lines.
- AskCloseKinToShareWorkers / OfferLaborToNeedyKin / ExecuteGiftLaborAgreement. These models regulate nonreciprocal exchanges of agricultural labor at times of need along kinship lines.

Household Evolution:

- ViableHousehold / DissolvingHousehold. The ViableHousehold model is activated periodically and at any demographic change in a Household's membership. For ancient Mesopotamian society, the MASS group decided that a household would be considered socially viable if there is either at least one adult (not elder) male present, or, under a more relaxed alternative criterion, at least one adult, either male or female, present. The basis for this is primarily one of fitness to perform agricultural tasks, such as working the land behind a plow team. From a purely social standpoint, elders can head households, but for relatively simple agrarian simulation milieus there is a problem with farm labor if no adults are present to take on the bulk of the workload. If the ViableHousehold model determines that a Household is *not* viable, the DissolvingHousehold model is invoked to guide any surviving household members in appropriate coping activities. This will involve cueing additional models, such as TakePersonIn, CastPersonOut, and FindLocalHavenWithFamily, to attempt to relocate members of the defunct household with other local households. These models operate based on, for example, strength of kinship bonds and the ability of other households to take on additional members in terms of food reserves and subsistence task labor.
- DividingHousehold. This model is activated periodically and at any demographic change in a Household's membership. It assesses social stress within the household in terms of total size of household and number of eligible household heads (heads of Basic Family Units, or BFUs, within the Household; see Community Theme discussion below). If the stress levels are sufficiently high, the CanStartNewHousehold model described above is invoked to check whether household resources are sufficient to allow one of more BFUs to split off and form new independent households.

Death:

- Resolve Inheritance / Determine Legacy. These models are called upon whenever a simulated Person dies. The deceased Person's Society object is consulted to determine what portion of personal and household resources (land owned or field shares in a communal land tenure system, livestock, durable goods and stored grain, for example) are considered inheritable, and to whom such resources should be conveyed. These models represent a synthesis of four different ancient textual sources (Tenney 2003). These include evidence from Old Babylonian Sippar (Harris 1976) and from Mesopotamian legal codes such as the Laws of Lipit-Ishtar; the Code of Hammurabi, and Middle Assyrian Laws from Assur (Brinkman 1985; Roth 1995). The models recognize four different inheritance patterns, depending on the beneficiaries involved. The ResolveInheritance model thus calls upon one of four sub-models:
  - AssignInheritanceNoOffspring
  - AssignInheritanceSonsOnly
  - AssignInheritanceDaughtersNoSons
  - AssignInheritanceSonsDaughtersWives

The Demographics Object Class

A Demographics class object encapsulates probabilistic information regarding the fertility and mortality of the individual persons in a modeled population as a function of the age and gender of the individual. The object's demographic algorithms incorporate Coale and Demeny's model life tables (Model West Level 2 and 4 for females and males respectively) to create individual-level demographic data suitable for the ancient Near East (Coale & Demeny 1966). This creates a population with relatively high death rates at very young and old ages, and a life expectancy at birth in the mid to early twenties for males and females respectively. The probability of male births is slightly higher than female births (51 versus 49 percent). Overall, high birth rates – an average of 6 children per woman – are predicted, though the very high proportion of children who die either die at birth or in childhood counteracts the high fertility rate.

In a simulation, each modeled individual (a Person object) automatically checks its Demographics object (accessed via its Population object; see Figure 10.5) annually to see if death will occur some time in the coming year (this is of course *not* recorded in the Person's modeled *perceptions*). If death is predicted in the coming year, a random number draw is used to determine the precise time of death under the assumption that all death dates during the year are equally likely.

A similar approach is used for childbirth, for modeled females who are within the age range for fertility as set forth in the data tables in the Demographics object, *and who are socially in a position to bear children*. This last stipulation is highly significant. The Coale and Demeny fertility tables provide *aggregated* fertility values, listing, e.g., live births per year per thousand women in each one-year age interval. They do <u>not</u> account for physically or social based constraints on fertility that could affect individual women differently. In our initial simulations using the Coale and Demeny demographic data, we consistently saw initial settlement populations quickly dwindle away. The reason for this anomaly, as it turned out, was that in the simulations we were checking women for upcoming childbirth events only if they were in a situation in which childbearing would be socially condoned – in other words, only if they were married. We determined that, given the prevailing demographic climate, fertile-age females are in this state only about 45 percent of the time, so we were underestimating birthrates by more than a factor of two by using aggregated birthrates in a socially disaggregated simulation. We adjusted the birthrates in the ancient Mesopotamian Demographics object to account for this factor and immediately saw population dynamics fall right into line with the expected norms. In the foregoing it should be noted that the MASS Group was not so naïve as to be unable to imagine potential pathways to births out of wedlock; it was decided not to attempt to divert effort to drive the ENKIMDU social modeling representation to the fine level of detail that would be necessary to properly unravel such mechanisms. In any case it seems likely that a great many such occurrences might have been legitimized by an *ex post facto* marriage.

This sort of issue can be expected to arise repeatedly in constructing simulations that are informed by highly aggregated data. For example, aggregated death rates include such causes of death as warfare, homicide and disease. When such mechanisms are also *explicitly* modeled, it is important to guard against overestimating overall mortality by double-counting: in this case reflecting a given cause of death both in the aggregated death rates and in situation-specific death events. There is no universally correct path to take in such cases. This is a topic area that begs further research in its own right, since as modeling and simulation techniques become ever more sophisticated and fine-grained such complications will become increasingly frequent.

The ENKIMDU Populator Preprocessor

ENKIMDU simulations typically involve from hundreds to tens of thousands of simulated individuals, each with appropriate demographic characteristics and detailed kinship linkages. It would be an enormous, not to mention error-prone, task to generate such model populations manually by specifying data for one person at a time. To obviate this concern, a special program, the Populator, was built for construction of the initial population of simulated individuals and their initial assignments to simulated households. The Populator's present demographic algorithms incorporate Coale and Demeny's model life tables (Model West Level 2 and 4 for females and males respectively) to create individuals' demographic data suitable for the ancient Near East (Coale & Demeny 1966). Each Person object generated by the Populator is given a random, gender-appropriate individual name and a reference number to help in tracing family histories throughout a simulation run.

When the Populator generates a micropopulation of individual Person objects, the age and gender mix of Persons is demographically representative of the population archetype, but the Persons are unrelated to each other. The Populator draws selectively form the modeled population to assemble new Household objects within which the members are in appropriate kin, age and gender relationship to each other. For the Populator, rural household census data from Ptolemaic Egypt (Bagnall & Frier, 1994) were first used to reconstruct the initial percentages of various household types encountered in a rural society. These data, although problematic in some aspects, do match well with historical records from other similar regions and periods. Future simulations may employ other data sets, particularly if detailed census data from Mesopotamia becomes available. The basic household types found in the ancient Near East are present in the Bagnall and Frier data: single-person, non-kin related, nuclear family, extended family, and multiple family (two or more nuclear families) households (see also Chapter 6). Although multiple family households may have been the societally preferred form, it is assumed that social stress and mortality rates may have prevented many households from achieving their ideal patrilocal type. The simulation runs can help to answer to what degree this assumption is valid. Once all the individual Persons have been assigned to Households, and linked via appropriate kin relationships to others in their Households, the Populator has one further task. It will attempt to establish demographically and socially plausible inter-Household kinship links to an extent specified in the simulation scenario input data, with the density of such kin linkages determined by an input parameter value.

It is important to note that the household and kinship structures described above relate only to the establishment of initial conditions for a MASS simulation. Once a simulation is underway, the demographic and social texture of the simulated society is endogenously maintained and evolved, as a result of the cumulative actions of the individuals, households, and other social organizations represented.

The Household Object Class is a key element within both the Person Theme and the Community Theme. Many aspects of Household objects have already been touched upon. We will focus more closely on specific attributes and behaviors of the simulation objects of the Household class in the *Community Theme* discussion below.

The Domain Object Model: The Community Theme

Figure 10.6 moves the social focus up in scale from individual persons to communities and their related entities in the settled landscape: thus to the right in the domain object model schema (Figure 10.1). Again, only the most significant object classes are represented.

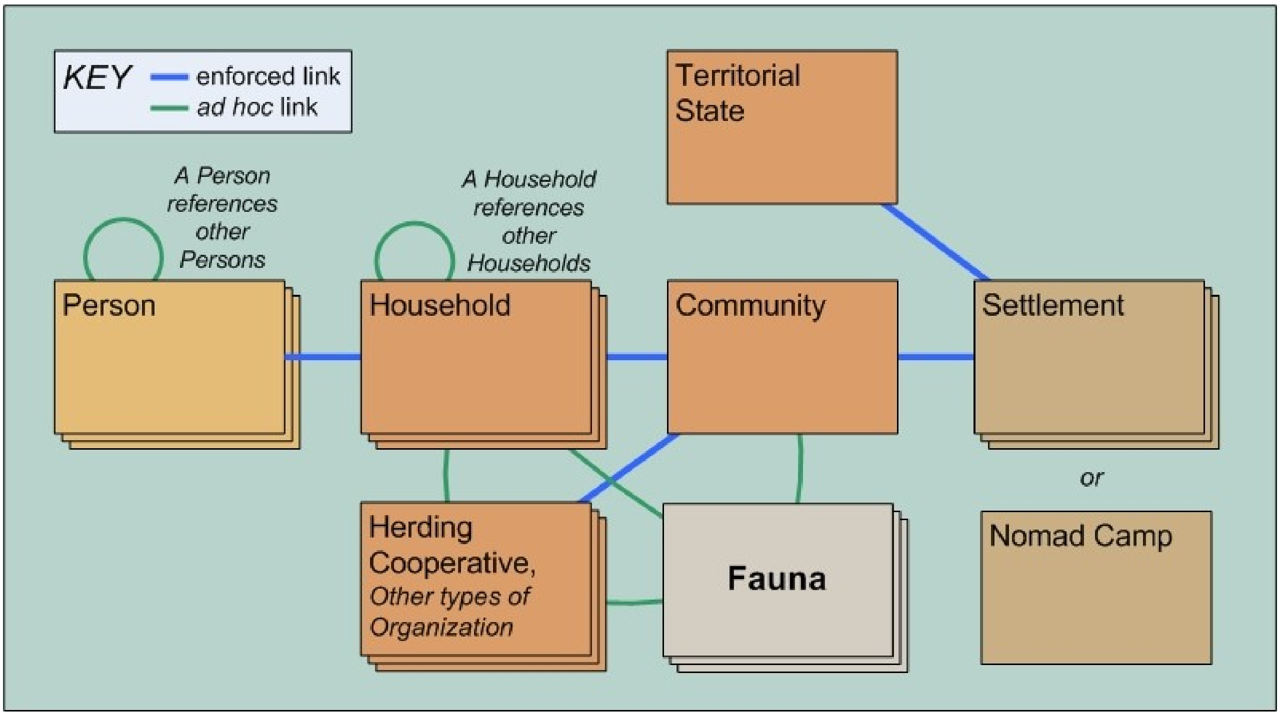

Fig. 10.6 MASS Domain Object Model: community and related object classes.

Instances of the *Person Object Class* provides an important context at the community level, especially as they perform diverse activities as agents of their Household objects. The Person class was discussed in detail in the previous subsection (*The Person Theme*) of this chapter.

The Household Object Class

The MASS Group agreed upon the household as the central decision making social unit for ancient Mesopotamian simulation-based analyses; thus the Household object class has been the principal social agent in our simulations. Most modeled social behaviors undertaken by simulated Persons are defined and modulated by their Households, drawing on normative cultural behavioral templates defined by the Society object associated with the head (Person) of each Household. It is at the household level that most resource allocation decisions are made, and it is at the household level that most social behaviors represented in the simulations are expressed. Real households can at times show a remarkable degree of adaptability and resilience under stress. Accordingly, the design that we have adopted for ENKIMDU Household objects emphasizes mutable, opportunistic coping mechanisms at several levels. The static structure of the Household object class is shown in Figure 10.7.

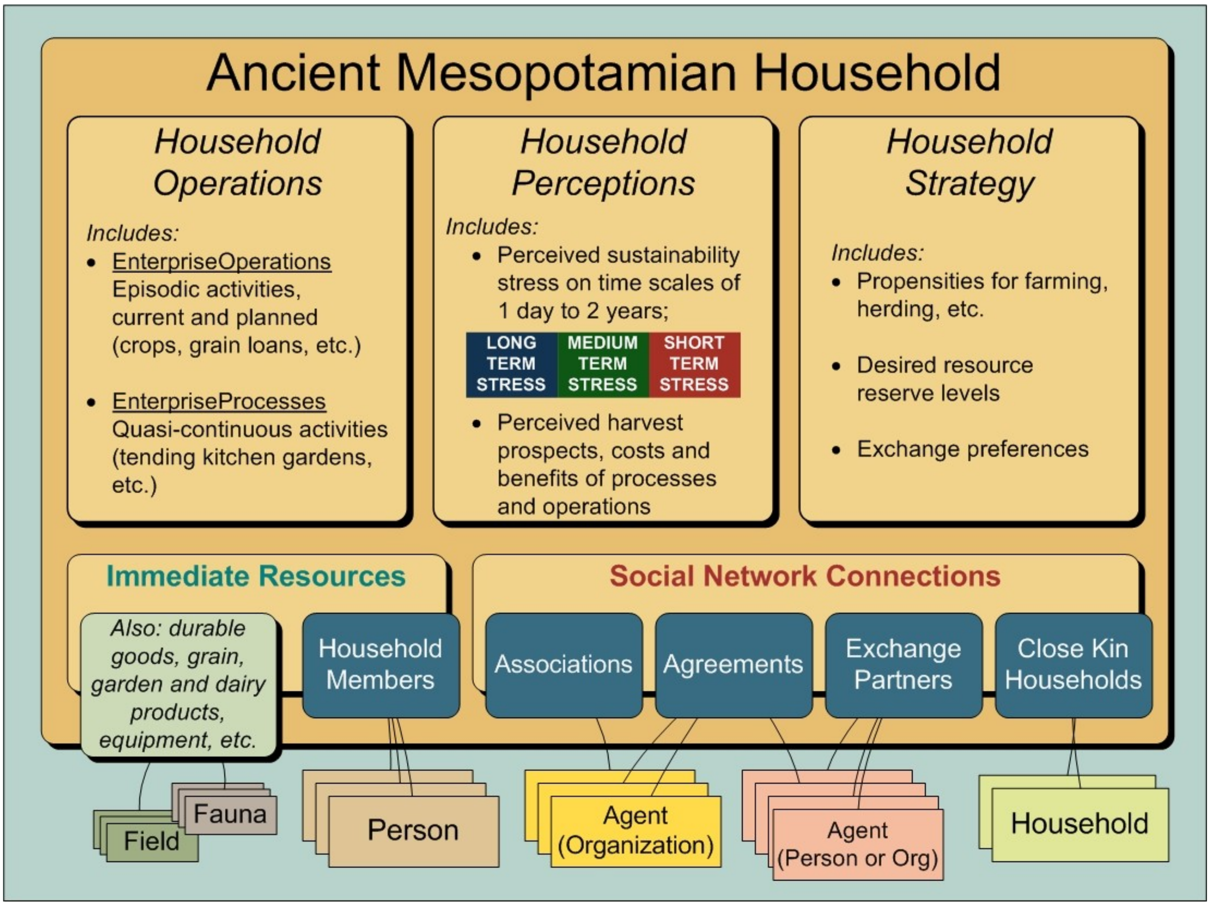

Fig. 10.7 Simplified schematic view of the Household Object Class.

Households and all other types of formal and *ad hoc* organizations modeled in ENKIMDU are represented as subclasses of a generic Organization class. All Organizations carry Operations, Strategy and Perceptions component objects. The first features of the Household object class portrayed in Figure 10.7 to be discussed are the Household's sub-classed versions of these classes: HouseholdOperations, HouseholdPerceptions, and HouseholdStrategy.

Household Operations

A HouseholdOperations object encapsulates the Household's information about its current and planned activities. Its main components are sorted collections of 'EnterpriseOperation' and 'EnterpriseProcess' objects. In the sense that we are using them here 'operations' refer to specific, episodic household activities of varying duration, such as planting and tending an annual grain crop or taking out and later repaying a grain loan. 'Processes' here refer to ongoing, quasi-continuous household activities such as tending livestock or working as an artisan or state official. These component objects describe the activities that the Household has in progress; those not yet underway that it has committed to pursue, and those simply being considered. Each of these has its own probability weights and error estimates based on each Household agent's perceptions.

Household Perceptions

A HouseholdPerceptions object encapsulates a Household's projections of its own social position and sustainability over a spectrum of time scales. The ability of the Household to meet its obligations, chief among them being feeding its members, is evaluated in the short term (days to weeks), middle-term (months, seasons) and long-term forecast periods. The long term estimate looks at the adequacy of the food supply from the current simulation time up to a point just short of the *harvest after next*. With this choice of timeframes, the household is forced to look beyond the upcoming harvest of already-planted crops, for which no strategic decision is either necessary or effective, to try to decide whether it appears advisable to intensify cultivation or apply other coping mechanisms (e.g. shift emphasis to pastoralism, etc.).

Household Strategy

A HouseholdStrategy object carries parameters intended to encode a Household's present strategy for sustaining (and perhaps aggrandizing) itself. In terms of relative importance these are weighted for general categories of gainful activity such as farming, herding, hunting, artisanship, and so on. Also included are parameters that scale the Household's degree of resistance to altering its strategy in favor of each of the various alternatives.

The Household sustainability estimate is based on resources on hand, and on the probability-weighted net gain or loss in negotiable household resources[5] to be accrued during the forecast period due to any of the household's enterprise operations and processes that will be active over the period of interest. Such estimates are modulated by the perceptions of the Household agents and their degrees of optimism, uncertainty, etc. with respect to these forecasts.

Household Sustainability

The portfolio of potential coping mechanisms for sustainability stress compiled by the MASS Group varies according to the time scale of the stress forecast, as can be seen in Figure 10.8. If a food supply shortfall in three days is projected, the potentially efficacious remedies are almost exclusively social (namely borrowing, liquidating assets, etc.), and are radically different than those that would mitigate a chronic shortfall predicted to begin, say, 18 months in the future.

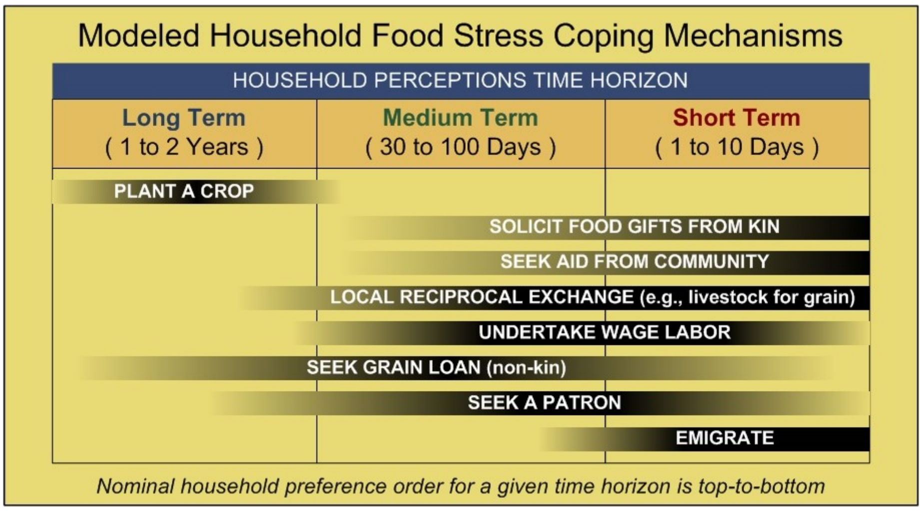

Fig. 10.8 Household food stress coping mechanisms for different time horizons.

[5] We assume that food is the fundamental item required in such negotiations.

Each Household in the simulation independently checks its sustainability (food) stress levels periodically at adaptively varying intervals (e.g. much more frequently if they are close to the edge of subsistence). All coping mechanisms available to the household and relevant to the time scale of the problem (Figure 10.8) are considered. Sustainability checks are also triggered:

- upon any change in the composition of a Household due to births, deaths, marriages, arrivals or departures;
- upon any change in a Household's sustainability status due to transactions it has made with other social agents, including entering or leaving a patronage arrangement with another Household, either as patron or as client;
- upon any change in sustainability status due to losses sustained by the Household due to diverse environmental or social causes, including losses in *expected* food production due to failure of a crop;
- immediately after the Household has engaged in a stress coping response, in order to reset the baseline stress levels and the stress check update intervals.

If perceived sustainability stress is significant, a Household will attempt to take action. If stress is perceived for multiple time horizons, the shortest-range problems are addressed first and stress is rechecked, since short-range solutions can sometimes mitigate longer-term problems as well; the reverse is seldom true.

If a Household chooses a new course of action as a result of the stress check, it will activate a corresponding aspect of its behavior that will cause new social pattern models to be executed. In ENKIMDU, growing field crops represents just one of a spectrum of sustainability strategies, that may be adopted or rejected as each Household deems appropriate. Thus, modeled households are not consigned to robotically set out crops every year. If they do set out crops, it is after weighing their perceived pros and cons of investing resources in setting out, maintaining and harvesting the crop in the context of all other sustaining options. Of course, all options are not equally attractive nor are they equally realizable for a given Household; these will depend on the demographic characteristics, material assets, and social standing of each Household.

The coping responses identified in Figure 10.8 can further described as follows:

- Plant a Crop (Agriculture). Depending on the modeling scenario conditions, Households may have control of cropland either by direct tenure or through the community-level mechanism of *musha'* (see also Chapter 4). In *musha'*, a settlement community can hold crop fields as a general resource, assigning them annually to households by lottery, with each household receiving the use of cropland in proportion to the each household's number of (inheritable) field shares within the community (Granott 1952). The agricultural and pastoral modeling schemes employed are discussed later in this chapter, in the *Agriculture and Pastoralism* subsection.
- Solicit Food Gifts from Kin. Households can rely on kinship links between Households, particularly those between household heads from the same nuclear family, for nonreciprocal gifts of food or labor at need. Gifts are made in proportion to the strength of the kinship link, modulated by the levels of food or labor reserve of the prospective donor Household.
- Seek Aid From Community. This modality is complementary to the kin-gift mechanism. Households may request nonreciprocal gifts or interest-free loans of resources from the village community at large. The strength of this mechanism can be varied in the simulation specification input, across a range that extends from no non-kin community assistance whatever to a fully communal arrangement in which all agricultural labor and its fruits are freely shared across the community.
- Local Reciprocal Exchange. In the present simulation, Households may exchange commodities and labor, opportunistically or in response to stress. In either case the mechanism is as described below under *Household Opportunistic Exchange*.
- Undertake Wage Labor. Household members may be employed in full-time positions by a TerritorialState or a patron Household, or may take on day labor to assist other Households with agricultural tasks. Wages are generally assumed to be in the form of a grain ration in simulations run to date.

- Seek Grain Loan. A Household may contract with another modeled Household for a loan of grain. Loan repayment with interest is expected after the next grain harvest. Interest has been taken to be one third of the principal (Garfinkle 2004). The simulation's GrainLoanAgreement objects allow for varying the interest rate and provides for various alternatives for modeling a creditor's response to delays or defaults in repayment.
- Seek a Patron. Households may opt to seek protection as clients of a stronger household. Patron and client privileges and obligations are set forth in an agreement (represented in the simulation by a PatronageAgreement object). The terms may vary depending on simulation scenario specifications input by the user. In general, patrons are responsible for assuring an uninterrupted food supply for their clients, and clients generally must provide some portion of their labor and their harvest output to their patron. A Household's retinue of client Households is taken onto account in reckoning the social status of each household; this comes into play in, for example, determining leadership succession in a modeled TerritorialState. Depending on the detailed terms of the agreement, the patron may legally become the son of the head of the client household, so that he may inherit the client's land upon the client's death.
- Emigrate. Households that can find no local remedy to sustainability problems will attempt to emigrate to another village (preferably where there are kin already residing or where a patron can offer help) or to leave the modeled settlement system altogether.

Household Opportunistic Exchange

Households respond to sustainability stresses in a largely reactive way; they are invoked only in the context of a perceived problem. It could be argued that when a modeled Household evaluates and formulates responses to perceived *longer-term* food stress levels, it is prudently planning ahead and is thus behaving proactively. There is another mechanism available to modeled Households that is explicitly proactive. This mechanism is provided via adaptively periodic opportunity checks that are the proactive complement to a Household's adaptively periodic sustainability checks. The thrust of the agent opportunity mechanism is to support reasoning by a Household regarding exchanges that it could make with other agents that would result in improvement in the net well-being of the Household. Households then 'shop around' within the social scope available to them, which could be the Household's local village community, or perhaps an ephemeral meeting-ground for exchanges with a visiting nomadic community. The Household tentatively offers its own beneficial exchanges of commodities in search of other Households for which the offered exchange would be beneficial to them. This process, in essence, invokes an assessment of marginal utilities (Jevons 1871), in the microeconomics sense, for all possible exchanges a Household could make. In the current simulation framework, the commodities available for exchange include grain, livestock, textiles, labor, and silver (as surrogate for durable goods of various sorts). Each Household ranks and evaluates potential exchanges from an idiosyncratic, less than fully rational perspective more in keeping with the approach of behavioral economics (Arielly 2008). For example, our model Households can tend to 'overvalue' livestock as a high-status item, above and beyond their worth in strictly utilitarian terms.[6] Household reciprocal exchanges can be driven by social, rather than strictly economic needs, as with the willingness of model Households to tolerate highly disadvantageous exchanges in order to have fresh meat on hand for a wedding feast.

The Community Object Class

The Community class is intended to represent the social fabric and dynamics of an individual settlement. It carries references to all of the Households and other types of organization or institution (such as a community or state granary) that reside or operate in the community. A subclass, NomadCommunity, deals with such mobile communities. Community objects also maintain a small cache of local knowledge and normative local practices and beliefs. Relative values of commodities, expressed in terms of equivalent kilograms of wheat, are tracked and maintained at the Community level, as a reference for reciprocal exchanges among community members. Community-level social processes that are represented by Community-class objects include the following:

[6] This practice has been noted by observers in traditional Syrian communities.

- Cropland Allocation. In the *musha'* communal land tenure system described earlier, the Community (presumably via a council of elders) conducts an annual lottery to assign specific fields to specific households. If the community as a whole has committed to biennial fallowing, the Community object enforces this restriction by only allocating fields that were left fallow the previous year. In individual ownership land tenure systems, Communities do not enforce fallowing, as that becomes an individual household decision.
- Daily Pasture Allocation. In the simulations, flocks of sheep and goats comprising livestock pooled from several cooperating households (a HerdingCooperative object) in a settlement community are pastured locally each day and brought back to the settlement in the evening. Naturally, some pasture areas are better than others due to their proximity and greater abundance of biomass. Each day, the Community adjudicates the assignment of specific pastures to specific cooperatives, to assure fair allocation of this resource.
- Communal Risk-Sharing. The Community regulates the collection of agricultural surpluses and their allocation to a communal reserve if called for in the simulation scenario specifications. It also receives and acts upon requests by needy Households for access to communal resources.

The NomadCommunity object class differs from the Community class in some respects, but it performs the same types of function. In the present simulation system prototype, nomad communities are assumed for simplicity to be non-hierarchical, so that decisions on such matters as travel routes and schedules are made at the NomadCommunity level.

The Settlement Object Class

Each Settlement-class object represents the built environment of a human settlement. Settlements maintain directories of all of the property lots and structures of all types within them, and of the roads and pathways that run through them. The urban transportation network is represented by a GroundTransportNetwork linked to other networks extending into the settlement's field system and to regional roadways that connect to other settlements, as has been discussed in Chapter 9. Settlements have no dynamic behaviors of their own; it is the communities they harbor that provide all the various aspects of urban dynamics.

Village and town spatial layouts for MASS simulations can be automatically generated by the Kabta urban texture engine (named for the ancient Mesopotamian god of bricks), a component of the Landscape Partition Engine (LPE) discussed earlier. Input parameters to the Kabta engine can be varied to produce arbitrary urban morphologies ranging from regular, rectilinear street grids to convoluted, irregular layouts replete with meandering lanes, oddly shaped blocks or insulae and cul-de-sacs. Lot sizes and layouts within each block / insula can be similarly varied. Figure 10.9 illustrates a layout for one of the North Jazira settlements generated by the Kabta engine.

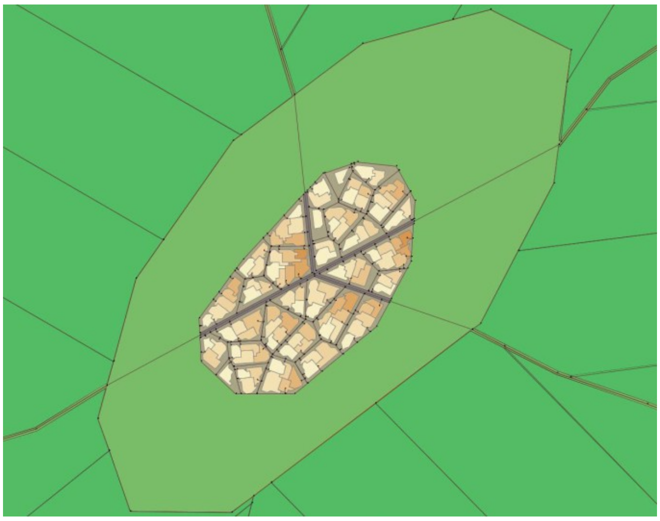

Fig. 10.9 Example of settlement layout generated by the MASS Kabta preprocessor program.

The TerritorialState Object Class

Territorial State objects are intended to provide crude representations of polities of higher order than the individual settled or nomadic communities just described: petty kingdoms typically controlling several settlements, up to perhaps a dozen or more. A baseline set of State-level social behavior patterns has been implemented, addressing the following aspects of state-level social dynamics.

Taxation. The StateTaxationModel simulates a state grain tax imposed at harvest time. States collect and maintain grain surpluses in State granaries in the settlements they control, to support their own extended bureaucracies and to provide a buffer against bad harvests.

Specialization. As a user-selectable option, we are modeling the tendency of states to encourage an increasing proportion of non-primary producer households (NPPs) in their settlements. NPPs can for present purposes be defined as households in which the principal livelihood is gained by members engaging in activities other than farming or herding, such as trades, crafts or services. NPPs are largely undifferentiated in the current ENKIMDU design, but they can act as reasonable proxies for elements of state bureaucracy and vocational specialization, especially in that they require settlement systems to boost agricultural production to generate sufficient surplus to feed the additional 'non-producing' segment of the population. State workers receive a living allowance (in grain) from the State reserves in return for their services. Other NPPs such as free artisans of all types (they remain undifferentiated in the present prototype simulation system) must also receive the means to secure their living needs. In the absence of a complete economic model, state workers receive a periodic income in 'silver' (a proxy for tradable durable goods, in some cases perhaps actual silver) that can be exchanged for grain as long as surplus grain is available in the community. This is an area of the MASS simulation that will benefit from future research and development.

State Internal Security. The StateResponseToLocalizedFamineModel represents a modality of state-level response to internal subsistence crises (e.g. localized famines). Leaders of territorial states monitor the settlements in their kingdoms for food shortages that may lead to instability and unrest, and may as a result either release portions of the state grain reserves in those settlements for famine relief, or may mobilize relief caravans to transport state grain reserves from other, less affected, settlements to the afflicted villages. The effectiveness of the state in mobilizing grain reserves is expected to vary widely between northern and southern simulation scenarios due to long-range bulk transport mechanisms that are more efficient in the south (waterborne transport by barge), than in the north (pack animals and, terrain and, road conditions permitting, carts). This model will be extended to address grain transport by water, to support sub-regional scale simulations of southern Mesopotamian scenarios.

Leadership Succession. The SuccessionModel provides mechanisms for determining succession upon the death of a ruler, as well as the option for forcible removal of unpopular or under supported ruling households by better-favored members of the elite in certain situations.

The Fauna Object Class

Individual animals are represented as Fauna objects. To date, the Fauna class has been used to represent domesticated animals: sheep, goats, and draft and pack animals such as donkeys and oxen. The Fauna class accesses detailed models of physiology and reproductive attributes for the species of interest, including nursing and suckling, onset and duration of lactation and amount of milk produced daily, body mass, meat quantity at time of slaughter, and rates of biomass consumption and rates of manure excretion (Redding 1981; Blaxter 1967). Instances of the Fauna class reference a FaunaType object for data concerning the fundamental characteristics of their particular species. If the need arises, wild creatures at large in the ecosystem can also be modeled as individuals using the Fauna object class, though there would have to be substantial enhancements to the Fauna behavior models to address the wider repertoire of behaviors for undomesticated animals.

MASS -ENKIMDU Domain Object Model: The Infrastructure and Technology Theme

A number of domain object classes provide support for the societal activities that the MASS simulations attempt to portray. These constitute a collection of technological artifacts: tools, devices, vehicles, and so on. Table 10.1 identifies some of the more important of these object classes.

Table 10.1 Important MASS ENKIMDU infrastructure and technology objects.

Object Class	**Description; Examples**
Device	An implement such as a plow or harrow; a harness or yoke for a draft team.
Vehicle	LandVehicle subclass: wagons, carts. Vessel subclass: boats, rafts, barges.
Coupling	An appurtenance attached to a Device to allow it to connect to another Device in a specific way.
Traction (interface)	Any entity with the ability to pull a load over a land surface, for example harnessed draft animals.
Apparatus	A compound device assembled from coupled Devices, capable of doing a specific set of tasks. Example: draft team + harness + plow is a plowing apparatus.
Building	A dwelling, warehouse, palace, temple, etc.
TransportNetwork	A network of links connecting nodes used to coordinate land movement (GroundTransportNetwork) or navigation (WaterMovementNetwork).
HydroStructure	A dam, weir, spillway, etc. Levees and berms are not included here; they form portions of the boundaries of LandscapePartition objects (e.g., fields).

The Domain Object Model: Landscape Theme

Discussion in this chapter has thus far centered on the societal component of the MASS simulations. The environmental component of the simulation is every bit as important, and will now brought into the spotlight. The simulation software object classes most relevant to the simulation system's dynamic representation of the natural environment are shown in Figure 10.10.

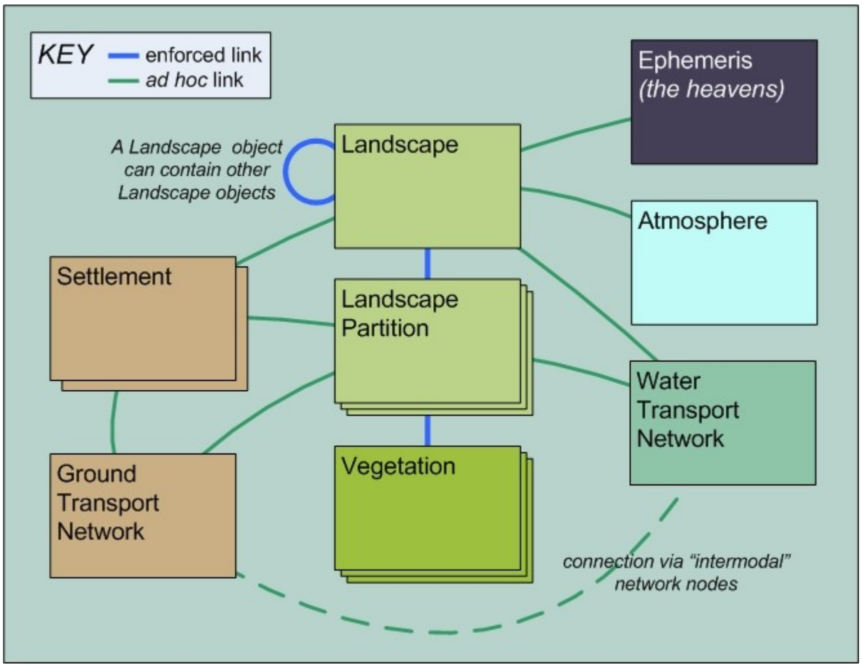

Fig. 10.10 MASS Domain Object Model: landscape and related object classes.

The key object classes here are Landscape and LandscapePartition. These classes most directly express the distributed environmental state, and the inhomogeneous changes in that state, that are experienced by the social agents in the simulations, and through which the agents impose societal impacts on the natural world.

The Landscape Object Class

The Landscape object class is the root class for all terrestrial simulation entities. A Landscape objects can represent the entire landscape of interest or a smaller subdivision, such as a drainage basin or watershed. Landscape objects are nestable in spatial hierarchies. It is at the Landscape level that information on the coordinate system transformation is specified. This is used to generate cartesian coordinates from geocentric locations. A Landscape, which is generally the root Landscape in a hierarchy, will carry a reference to a Hydrosphere object that encapsulates information on surface water (both permanent and ephemeral features), and nominally on ground water as well, although to date all of our ground water modeling has been done directly in the Landscape object itself. A Landscape's mosaic of surface cover and land use patches is carried in its collection of Landscape Partition objects.

The most important function of the Landscape class is in evolving the state of the landscape forward in time: vegetation, soil and subsoil moisture and structure, as well as ground water states, are all updated periodically. In the MASS modeling system, this Landscape behavior is currently implemented by the USDA's SWAT (Soil and Water Assessment Tool) model (Arnold *et al.* 1998; Arnold & Allen 1992).

The full list of processes addressed by the SWAT model is extensive, and includes:

- *Hydrology*: surface runoff and ponding; drainage routing at individual field to watershed scale; percolation and water table dynamics; lateral subsurface flow; evaporation; snowmelt; and precipitation recharge.
- *Meteorology*: daily agricultural weather, via a Markov Chain stochastic weather generator that is driven by data tables of climatological means and extremes.
- *Soil Evolution*: soil structure, moisture, and temperature dynamics; water and wind erosion, deflation, etc.
- *Nutrient Cycling Dynamics*: plant nitrogen and phosphorus balance, and nutrient runoff in surface water.
- *Vegetation Growth*: canopy and root development; evapotranspiration; water budget; nutrient uptake; growth constraints (by moisture, temperature, nutrient levels, etc.); harvest yield, with constraints as above; dormancy; effects of grazing and browsing by livestock.
- *Human Interventions*: tillage effects (leveling, plowing, planting, harrowing, harvesting, etc.); irrigation and fertilizer. (Pesticides are not relevant for ancient studies but come into play in modern extensions of the simulation system.)

The SWAT[7] model's native domain data elements have been mapped 1-to-1 with ENKIMDU domain objects (Table 10.2). This is convenient, of course, but is also quite natural, since both simulation systems share a commonly held view of landscape structure and dynamics.

[7] The SWAT model version 2000 accessed by ENKIMDU is a stand-alone Fortran 90 program. To allow it to communicate with the DOM, it is 'wrapped' inside a Java language object (Christiansen & Altaweel 2006a) that can access (with the aid of some C language code for direct access to SWAT's large internal common data block) the internals of the model and can send and receive messages to and from other remote processes such as ENKIMDU.

Table 10.2 Logical mapping from SWAT data structures to ENKIMDU object classes.

SWAT Model Data Structure	ENKIMDU Object Class
Basin	Landscape
Sub-Basin	Nested Landscape
Hydrologic Response Unit (HRU)	LandscapePartition (e.g., a pasture or crop field)

In general, a separate execution thread of the SWAT model will be running concurrently for each dynamic Landscape object as defined in a MASS simulation scenario specification.

SWAT is capable of handing water routing among modeled sub-basins within a basin. For scenarios involving hydrologically connected basins, each is modeled in a separate execution thread; another mechanism is needed to move water between basins in a realistic fashion. It was anticipated that the U.S. Geologic Survey's MODBRANCH model (Swain and Wexler 1996) could be used to implement water flow routing among drainage basins in a linked network.

The LandscapePartition Object Class

LandscapePartitions are spatial subdivisions of a Landscape object. They may be of arbitrary size and shape as long as they represent a single, contiguous area within the boundaries of a Landscape. Each LandscapePartition is assumed to have a single uniform soil type and a single uniform vegetation type. These last constraints, imposed by the formalisms underlying the SWAT model, will to some degree dictate how a model Landscape is partitioned. LandscapePartition soil characteristics are carried in the LandscapePartition objects themselves, while vegetation characteristics reside in a Vegetation object referenced by each vegetated LandscapePartition.

The Field object class, a special subclass of LandscapePartition, carries additional attributes and methods related specifically to agriculture, such as area plowed, sown, harvested, etc., crop rotation year number, dates when the field was last in fallow, and so on.

The Vegetation Object Class

The Vegetation class represents the current vegetation cover of a single LandscapePartition. Each Vegetation object carries information on the current condition of the vegetation, including such parameters as phenological state, dormancy, above-ground and total biomass, leaf area index, and rooting depth. It references a VegetationType object for data concerning the fundamental characteristics of the plant species being grown. Vegetation types can be annual or perennial, herbaceous or woody, and can to a limited degree include a mix of plant species, a long as a single set of VegetationType growth parameters can adequately represent the growth of the LandscapePartition's plant population as a whole.

The Hydrosphere and Atmosphere Object Classes

For complex simulation scenarios that include a network of multiple hydrologically connected Landscape objects (each representing drainage basin), a Hydrosphere object is needed to periodically update and record ground water and surface water flows between the connected Landscapes and into and out of the network as a whole. The Hydrosphere object calls upon a model designated to implement this behavior, such as the MODBRANCH model (Swain & Wexler 1996).

Similarly, for large and complex scenarios, particularly geographically extensive scenarios involving more than one Landscape object, an Atmosphere object must be assigned to advance the spatially varying atmospheric state forward in time in order to provide periodic (daily or more frequent) updates of spatially varying surface weather conditions. An appropriate model for this purpose would be a mesoscale weather model such as MM5 (Anthes & Warner 1978).

The Ephemeris Object Class

The Ephemeris object class provides necessary information on times of sunrise and sunset and solar declination and zenith angle as a function of the latitude and longitude of an observer. This information is critical to the timing of diverse processes in ENKIMDU representations of a settled landscapes. The Ephemeris class includes a re-implementation of the relevant algorithms drawn from the NASA International Reference Ionosphere astrophysical software library (Bilitza 1990).

AGRICULTURE AND PASTORALISM

In agricultural and pastoral activities, interactions between social and environmental processes are centrally important (see also Chapter 12). Households that have access to cropland (via ownership, sharecrop agreement with a patron, or community musha' allotment) and are seeking to put it into agricultural production are faced with a complex sequence of agricultural tasks. Each task has its own distinct logistical and manpower needs (Gallant 1991; Russell 1988) along with a host of contextual requirements dealing with such issues as pre-existing field conditions, crop phenological state, and timing to coordinate with weather. For example, it takes longer to plow a field if it is necessary to first break fallow. Within the agricultural year, simulated households must clear and level each field, and then plow, sow, weed and maintain, and harvest it, and then must process the harvested crop for storage and use. Depending on household resources as well as household and community preferences and common practices, households may also manure their fields, and if appropriate and feasible, irrigate them. The sequence of tasks for raising a crop in a single field is illustrated in Figure 10.11. Not shown in Figure 10.11 is the pervasive and generally nontrivial task of walking to and from the field each day, often burdened with equipment, supplies, or harvest products, for household work crews residing in their nucleated settlements. Within ENKIMDU, the complex dynamics of this sequence of operations is represented by two models of Household behavior: the FieldCropManagementModel and the DisposeHarvestModel.

The FieldCropManagementModel for ancient Mesopotamia reflects Mesopotamian conditions. For instance, in both modern and ancient northern Mesopotamia, the rainy season begins in November and continues through April. Seed is sown in winter to take advantage of moisture available for germination and establishment of seedling grain plants, and the harvest season is in the early summer, requiring crops to be harvested before they wilt under the increasing heat, lack of rain, and changing soil conditions during the harvest action step (Buringh 1960). Crop growth computations make use of crops (e.g. barley) known to have been grown in the areas modeled, and with growth parameters (Neitsch 2002) adjusted to reflect ancient genotypes (Coyle 1998).[8]

Interactions between social agents and landscape elements are explicitly simulated in fine spatiotemporal detail. Household social agents pursuing agricultural tasks can perceive and respond to changes to the landscape that result from the daily evolution of the state of each LandscapePartition (e.g. crop field) as well as its vegetation in a Landscape as estimated by the SWAT model. SWAT computations reflect changes due to endogenous environmental processes, but are also responsive to human interventions, such as tillage operations, sowing of seed, and fertilizer application. Regarding fertilizer, household manure budgets are tracked in terms of faunal excrement deposits in and around the household compound during the time that herds and flocks are not foraging in nearby pastureland. In the simulations, this manure can be collected by the households, dried, and then carried to the fields and spread as fertilizer.

[8] Coyle, C. (1998): Personal communication with Texas A&M University, Blackland Research Center.

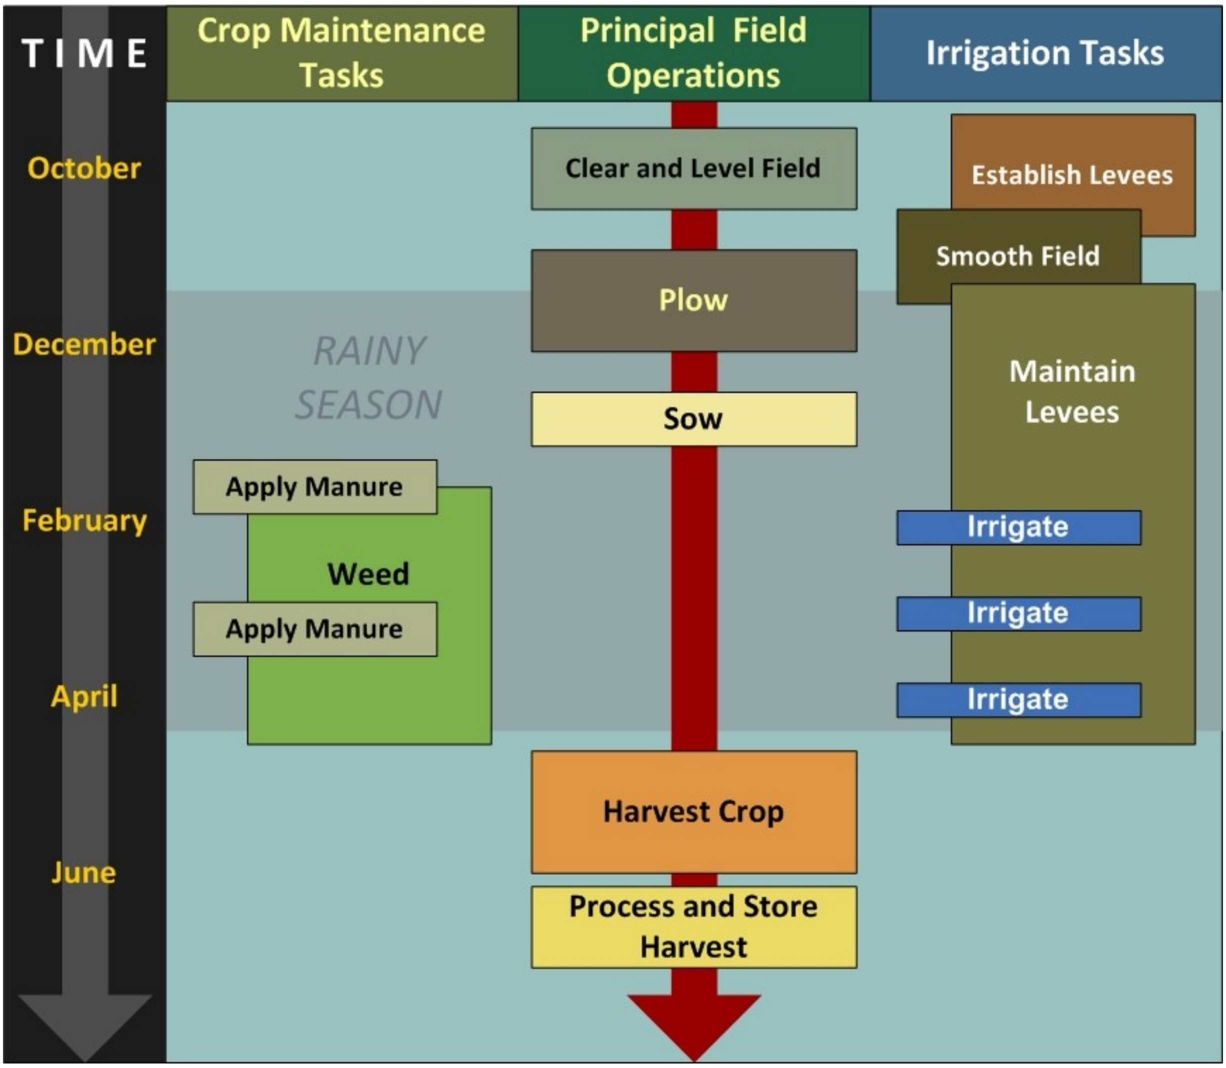

Fig. 10.11 Modeled household agricultural task sequences (note that the date of various tasks such as that of manure application can be varied to suit the model scenario).

The ManageHerdLocalModel addresses cooperative livestock management activities of a community as undertaken by groups of households (represented by HerdingCooperative objects). In turn, these keep their sheep and goats within their own household compounds overnight, but pool their livestock each day into flocks of roughly 150 to 300 animals. This pooling allows the combined herds and flocks to be efficiently taken out each day by a small work force to forage pastures that are near enough to allow them to return to the settlement before nightfall. The model is largely based on an ethnographic study that describes a pastoral system that may have also existed in past Mesopotamian societies (Sweet 1974).

As with modeled agricultural processes, simulated decisions regarding herding practices are dynamic and continuously evolving as circumstances change for each day. Flocks can be pastured on fallow crop fields or on pastureland beyond a settlement's crop field system. As discussed earlier in the chapter under the Community Theme, factors such as standing biomass in fields and distance from the settlement determine which areas are the most desirable as pasture, and each day the Community as a whole decides which cooperative's flock gets which pasture for the day. Livestock flocks affect the foraging value of the pastures by their trampling and consumption of standing biomass, and in the slightly longer term by and excrement deposition that enhances plant growth on the fields they traverse and graze.

Sheep and goats enhance their households' resource base by providing wool and hair for textiles, milk for dairy products, and meat. They are usually their principal exchangeable liquid asset, and are thus excellent buffers against food stress. Livestock physiological processes are simulated at the level of the individual animal, so the simulation generates births of lambs and kids, as well as lactation, nursing and suckling, and the growth and eventual death of each animal, by natural causes or (more likely, especially for males) by slaughter.

In Chapters 11 and 12 we discuss the results of several diverse simulation exercises and describe the framework's results and outputs. It also allows us to expand upon our exposition of key features of the framework in the course of laying out simulation scenario specifications and to discuss the details and implications of the modeling results.

REFERENCES

Anthes, R.A. & Warner, T.T. (1978). Development of Hydrodynamic Models Suitable for Air Pollution and Other Mesometeorological Studies. *Monthly Weather Review* 106:1045-1078.

Arielly, D. (2008). *Predictably Irrational*. New York, NY: HarperCollins.

Arnold, J.G. & Allen, P.M. (1992). A comprehensive surface-groundwater flow model. *Journal of Hydrology* 142:47-69.

Arnold, J.G., Srinivasin, R., Muttiah, R.S. & Williams, J.R. (1998). Large area hydrologic modeling and assessment. Part I: Model development. *Journal of the American Water Resources Association* 34(1):73-89.

Blaxter, K.L. (1967). *The Energy Metabolism of Ruminants*. London: Garland.

Brinkman, J.A. (1985). A Note on the Middle Assyrian Laws. *Revue d'Assyriologie* 79: 88-89.

Bagnall, R.S. & Frier, B.W. (1994). *The Demography of Roman Egypt*. Cambridge, UK: Cambridge University Press.

Bilitza, D. (ed.) (1989). International Reference Ionosphere 1990, pp. 90-22. Greenbelt, MD: NSSDC.

Buringh, P. (1960). *Soils and Soil Conditions in Iraq*. Baghdad: Ministry of Agriculture.

Christiansen, J.H. (2000a). A flexible object-based software framework for modeling complex systems with interacting natural and societal processes. Banff, Alberta, Canada: Proceedings 4th International Conference on Integrating GIS & Environmental Modeling, http://www.colorado.edu/research/cires/banff/pubpapers/72/index.html

Christiansen, J.H. (2000b). FACET: a simulation software framework for modeling complex societal processes and interactions. Proceedings 2000 Summer Computer Simulation Conference of the Society for Computer Simulation, International. Vancouver, Canada.

Christiansen, J.H. & Altaweel, M. (2006a). Understanding ancient societies: a new approach using agent-based holistic modeling. *Structure and Dynamics: eJournal of Anthropological & Related Sciences*, 1 (2), Article 7. http://repositories.cdlib.org/imbs/socdyn/sdeas/vol1/iss2/art7/ .

Christiansen, J.H. & Altaweel, M. (2006b). Simulation of natural and social process interactions: an example from Bronze Age Mesopotamia. *Social Science Computer Review* 24 (2):1-17.

Coale, A.J. & Demeny, P. (1966). *Regional Model Life Tables and Stable Populations*. New York: Academic Press.

Doran, J.M., Palmer, M., Gilbert, N. & Mellars, P. (1994). The EOS project: Modeling Upper Paleolithic social change. In N. Gilbert & J. Doran (eds.) *Simulating Societies: The Computer Simulation of Social Phenomena* pp. 195-221. London: UCL Press.

Epstein, J.M. & Axtell, R. (1996). Growing Artificial Societies: Social Science from the Bottom Up. Cambridge, MA: MIT Press.

Gallant, T.W. (1991). *Risk and Survival in Ancient Greece*. Stanford: Stanford University Press.

Garfinkle, S.J. (2004). Shepherds, merchants, and credit: some observations on lending practices in Ur III Mesopotamia. Journal of the Economic & Social History of the Orient 47 (1): 1-30.

Gosling, J. & McGilton, H. (1995). *The Java Language Environment: A White Paper*. Mountain View, CA: Sun Microsystems Computer Co.

Granott, A. (1952). *The Land System in Palestine: History and Structure*. London: Eyre & Spottiswoode.

Harper, R.F. (1904). *The Code of Hammurabi*. Chicago, IL: University of Chicago Press.

Harris, R. (1976). On Kinship and Inheritance in Old Babylonian Sippar, *Iraq* 38:129-32.

Jevons, W.S. (1871). *The Theory of Political Economy*. London: MacMillan & Co.

Kohler, T.A., Kresl, J., West, C.V., Carr, E., & Wilshusen, R. (2000). Be there then: A modeling approach to settlement determinants and spatial efficiency among late ancestral Pueblo populations of the Mesa Verde region, U. S. Southwest. In T.A. Kohler & G.J. Gumerman (eds.) *Dynamics in Human and Primate Societies: Agent-Based Modeling of Social and Spatial Processes*, pp. 145-178. New York: Oxford University Press.

Lansing, J.S. (1991). *Priests and Programmers: Technologies of Power in the Engineered Landscape of Bali* Princeton, NJ: Princeton University Press.

Lansing, J.S. & Kremer, J.N. (1993). Emergent properties of Balinese water temple networks: coadaptation on a rugged fitness landscape. *American Anthropologist* 95(1): 97-114.

Lurie, G., Sydelko, P., & Taxon, T. (2002). An object-oriented GIS toolkit for web-based and dynamic decision analysis applications. Journal of Geographic Information & Decision Analysis 6(2): 108-116.

Modelski, G. (1997). Cities of the Ancient World: An Inventory (-3500 to -1200). The Evolutionary World Politics Home Page: https://faculty.washington.edu/modelski/WCITI2.html

Neitsch, S.L., Arnold, J.G., Kiniry, J.R., Srinivassan, R. & Williams, J.R. (2002). *Soil and Water Assessment Tool User's Manual Version 2000*. Texas Water Resources Institute, College Station, TX: TWRI Report TR-192.

Redding, R.W. (1981). Decision Making in Subsistence Herding of Sheep and Goats in the Middle East. PhD dissertation, Dept. of Anthropology, University of Michigan.

Roth, M.T. (1995). *Law Collections from Mesopotamia and Asia Minor*. S.B. Parker (ed.) vol. 6, Writings from the Ancient World. Atlanta: Scholars Press.

Russell, K.W. (1988). *After Eden: The Behavioral Ecology of Early Food Production in the Near East and North Africa*. BAR International Series, 391. Oxford: British Archaeological Reports; Archaeopress.

Simon, H.A. (1982). *Models of Bounded Rationality*. Cambridge, MA: MIT Press.

Snyder, J.P. (1987). Map Projections - A Working Manual. U.S. Geological Survey Professional Paper 1395. Washington DC: U.S. Government Printing Office.

Sun Microsystems (1999). Code Conventions for the Java Programming Language. http://java.sun.com/docs/codeconv/ .

Sun Microsystems (2006). Sun Developer Network: Java 2D API, http://java.sun.com/products/java-media/2D/ .

Swain, E.D. & Wexler, E.J. (1996). A coupled surface water and groundwater flow model (MODBRANCH) for simulation of stream-aquifer interaction. *Techniques of water resources investigations of the USGS*. Washington DC: U.S. Geological Survey & U.S. Government Printing Office.

Sweet, L. (1974). *Tell Toqan, a Syrian Village*. Ann Arbor, MI: University of Michigan Press.

Tenney, J, (2003). Notes on Mesopotamian Inheritance. Unpublished report, MASS archive.

Widell, M. (2004). Some observations on the administration, agriculture and animal management of Tell Beydar. *Ugarit-Forschungen* 35: 717-733.

Widell, M. (2005). Some reflections on Babylonian exchange during the end of the third millennium BC. *Journal of the Economic & Social History of the Orient* 48(3): 388-400.

Widell, M. (2007). Historical evidence for climate instability and environmental catastrophes in northern Syria and the Jazira: the chronicle of Michael the Syrian. *Environment & History* 13(1): 47-70.

Wilkinson, T.J., Christiansen, J. H., Ur, J. A., Widell, M. & Altaweel, M. (2007). Urbanization within a dynamic environment: modeling Bronze Age communities in Upper Mesopotamia. *American Anthropologist* 109 (1): 52-68.

Wilkinson, T.J. & Tucker, D.J. (1995). *Settlement Development in the North Jazira, Iraq: A Study of the Archaeological Landscape*. Warminster UK: Aris and Philips & British School of Archaeology in Iraq.

van Liere, W. (2003). Survey of Soil, Present Land Use and Land Capabilities of the Jezireh. Unpublished Report in the MASS archives.

Villers, S. & Sommerville, I. (2000). Ethnographically informed analysis for software engineers. *International Journal of Human-Computer Studies* 53:169-196.