Proceedings of Multi-agent systems and Agent-based Simulation (MABS'98), 10-25, Springer, LNAI 1534, 1998.

Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users' Guide

H. Van Dyke Parunak*, Robert Savit**, Rick L. Riolo**

*Center for Electronic Commerce, ERIM, 2901 Hubbard Road, Ann Arbor, MI 48105 USA
**Program for the Study of Complex Systems, Randall Laboratory, Univ. of Michigan, Ann Arbor, MI 48109-1120 USA
vparunak@erim.org, savit@umich.edu, rlriolo@umich.edu

Abstract. In many domains, agent-based system modeling competes with equation-based approaches that identify system variables and evaluate or integrate sets of equations relating these variables. The distinction has been of great interest in a project that applies agent-based modeling to industrial supply networks, since virtually all computer-based modeling of such networks up to this point has used system dynamics, an approach based on ordinary differential equations (ODE's). This paper summarizes the domain of supply networks and illustrates how they can be modeled both with agents and with equations. It summarizes the similarities and differences of these two classes of models, and develops criteria for selecting one or the other approach.

1. Introduction

In many domains, agent-based modeling competes with equation-based approaches that identify system variables and evaluate or integrate sets of equations relating these variables. Both approaches simulate the system by constructing a model and executing it on a computer. The differences are in the form of the model and how it is executed. In agent-based modeling (ABM), the model consists of a set of agents that encapsulate the behaviors of the various individuals that make up the system, and execution consists of emulating these behaviors. In equation-based modeling (EBM), the model is a set of equations, and execution consists of evaluating them.[1] Thus "simulation" is the general term that applies to both methods, which are distinguished as (agent-based) emulation and (equation-based) evaluation.

Understanding the relative capabilities of these two approaches is of great ethical and practical interest to system modelers and simulators. The question is important ethically because the duty of simulators ought to be first of all to the domain being simulated, not to a given simulation technology, and the choice of technology should be driven by its adequacy for the modeling task as well as its intrinsic interest to the modeler. The question is important practically because most funding sources are driven by domain-dependent agendas and want to put their resources behind the simulation technology that will provide the best results.

[1] When "ABM" and "EBM" are arthrous or plural, 'M' means "model" rather than "modeling."

This paper explores the question in the problem domain of manufacturing supply networks, and giving examples of both ABM's and EBM's (Section 2). It discusses the relation between these two approaches at a high level (Section 3), and then compares their practical performance in three specific areas (Section 4). A concluding section includes recommendations for advancing and propagating ABM's.

2. The DASCh Experience

In our laboratory, the contrast between the two broad categories of models arose in the context of the DASCh project (Dynamical Analysis of Supply Chains) [12, 13], which explores the dynamical behavior of a manufacturing supply network. This section describes the application area, summarizes the structure and behavior of the agent-based model that was the focus of our research, and exhibits a system dynamics model of the same system to exemplify an equation-based approach.

2.1 What is a Supply Chain?

Modern industrial strategists are developing the vision of the "virtual enterprise," formed for a particular market opportunity from independent firms with well-defined core competencies [10]. The manufacturer of a complex product (the original equipment manufacturer, or "OEM") may purchase half or even more of the content in the product from other firms. For example, an automotive manufacturer might buy seats from one company, brake systems from another, air conditioning from a third, and electrical systems from a fourth, and manufacture only the chassis, body, and powertrain in its own facilities. The suppliers of major subsystems (such as seats) in turn purchase much of their content from still other companies. As a result, the

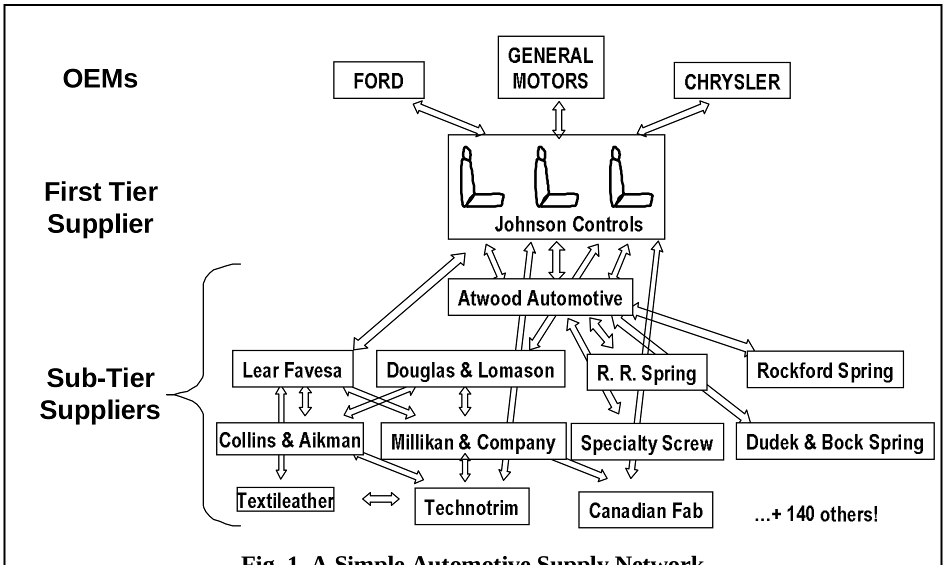

Fig. 1. A Simple Automotive Supply Network

"production line" that turns raw materials into a vehicle is a "supply network" (more commonly though less precisely called a "supply chain") of many different firms.

Fig. 1 illustrates part of a simple supply network [1, 8]. Johnson Controls supplies seating systems to Ford, General Motors, and Chrysler, and purchases the components and subassemblies of seats either directly or indirectly from over one hundred fifty other companies, some of which also supply one another. Issues of product design and production schedule must be managed across all these firms in order to produce quality vehicles on time and at reasonable cost.

In general, supply networks form an hourglass (Fig. 2), with an OEM at the center. Raw materials, parts, and subassemblies move up through the lower half of the hourglass to reach the OEM, and finished goods make their way through the upper half to the final consumer.

Fig. 2 is over-simplified. Linkages among firms are not restricted to the hierarchical patterns shown, and a single firm may appear in both halves of the hourglass. Still, the general pattern is one of convergence of materials to the OEM, and then distribution of the product out to end users. The two halves of the hourglass exhibit different dynamical behavior, driven by different mechanisms for forecasting demand through time. In the input (lower) half of the hourglass, the manufacturer can distribute a forecast of demand to its suppliers. There is no such centralized source of demand information in the distribution (upper) half of the hourglass, and the manufacturer must estimate demand from statistical analysis of the orders received over time. A typical algorithm, and the one that we use, is a weighted average of past orders over some time horizon, giving higher weight to more recent orders.

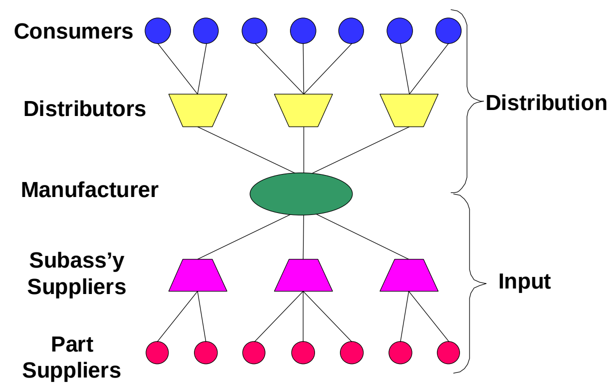

Fig. 2. Supply networks form an hourglass shape.

Supply networks, like most systems composed of interacting components, exhibit a wide range of dynamical behavior that can interfere with scheduling and control at the enterprise level. Data analytic approaches based on assumptions such as stationarity are not generally effective in understanding these dynamics, because the commercial environment changes too rapidly to permit the collection of consistent data series long enough to support statistical requirements.

2.2 An Agent-Based Model

DASCh explores the dynamics of a supply network by constructing and experimenting with an ABM that can maintain a given set of conditions as long as desired, permitting the collection of statistically relevant time series. Though artificial, this environment allows us to explore the dynamical nature of the supply network and can lead to important insights of great practical significance.

Model Structure. DASCh includes three species of agents. *Company agents* represent the different firms that trade with one another in a supply network. They consume inputs from their suppliers and transform them into outputs that they send to their customers. *PPIC agents* model the Production Planning and Inventory Control algorithms used by company agents to determine what inputs to order from their suppliers, based on the orders they have received from their customers. These PPIC agents currently support a simple material requirements planning (MRP) model.[2] *Shipping agents* model the delay and uncertainty involved in the movement of both material and information between trading partners.

The initial DASCh experiments involve a supply chain with four company agents (Fig. 3: a boundary supplier, a boundary consumer, and two intermediate firms producing a product with neither assembly nor disassembly). Each intermediate company agent has a PPIC agent. Shipping agents move both material and information among company agents.

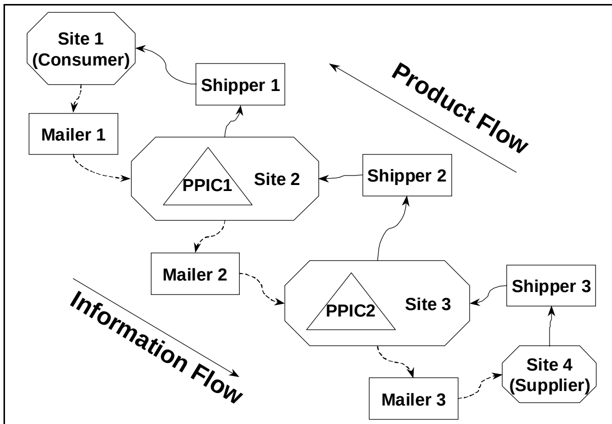

Fig. 3. The DASCh Supply Chain.

This simple structure was intended as a starting point. We expected it to exhibit relatively uninteresting behavior, on which the impact of successive modifications could be studied. In fact, it shows a range of interesting behaviors in terms of the variability in orders and inventories of the various company agents: amplification, correlation, persistence, and generation of variation in the orders and inventory levels in the system. In general, these phenomena introduce strong structural distortions into the order stream. Such disturbances obscure the suppliers' view of the top-level consumer's demand.

---

[2] The basic MRP algorithm includes developing a forecast of future demand based either on past demand or on customer forecast (depending on location in the hourglass), estimating inventory changes through time due to processing, deliveries, and shipments, determining when inventory is in danger of falling below specified levels, and placing orders to replenish inventory early enough to allow for estimated delivery times of suppliers.

Amplification and Correlation of Order Variation. As the demand generated by the top-level consumer propagates to lower levels, its variance increases, so that lower-level suppliers experience much more variability than higher-level ones. This amplification phenomenon is widely discussed in the literature. Not as well recognized is the correlation imposed on an originally uncorrelated series of random orders by the PPIC algorithms in the supply network.

To explore this dynamic we set all batch sizes to one, so the economic order quantity does not introduce a nonlinearity. The consumer generates Gaussian random IID (Independent, Identically Distributed) orders with a mean of 100 per week and variance of 10. Capacity at Sites 2 and 3 is set at 10,000 per week, virtually infinite in comparison with the order levels, again avoiding a threshhold nonlinearity. The forecast algorithm is the weighted average mechanism appropriate to the distribution half of the supply network hourglass. We examine the results using time delay plots, in which each element in a time series is plotted on the Y-axis against the previous element on the X-axis.

Fig. 4 shows the delay plot for the consumer orders. As expected for IID data, they form a circular blob, with no apparent structure. Fig. 5 and Fig. 6 show the orders issued by sites 2 and 3, respectively, in response to the IID consumer orders. These plots show two interesting features. First, although plotted to the same scale, the clouds of points are larger, reflecting amplification of order variation in successive tiers of the supply chain. Second, the clouds are no longer circular in shape, but are stretched along the line X = Y. This stretching indicates that these sites are more likely to follow a large order with another large one, and a small order with another small one. In other words, their orders have become correlated in time, and increasingly so as we go deeper in the supply chain.

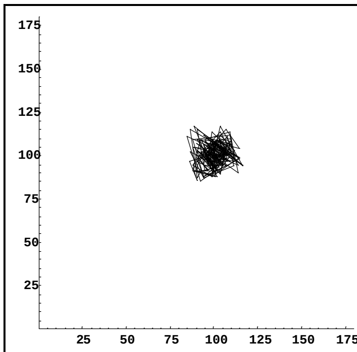

Fig. 4. Consumer Orders

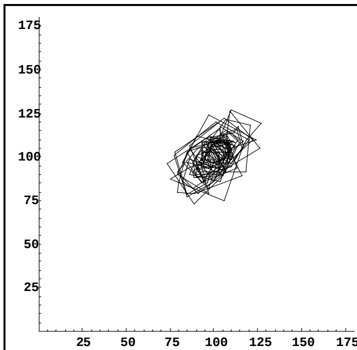

Fig. 5. Site 2 Orders

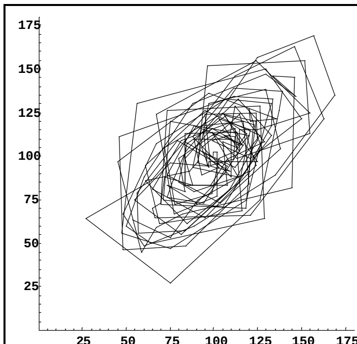

Fig. 6. Site 3 Orders

Persistence of Order Variation. A single modest change at the top of the chain generates disturbances in the order sequences of lower tier suppliers that persist long after the original change. Fig. 7 shows the effect of two successive step functions in consumer orders (the solid line) on the orders issued by site 3 to the supplier (the dashed line), using weighted average forecasting. In both cases, the consumer increases its order level by 10 orders per time period. Though the change in consumer orders is a one-time phenomenon, its effect persists in the orders that site 3 issues to the supplier. The persistence time is of the same order as the forecast window over which the manufacturer averages past orders to estimate future demand.

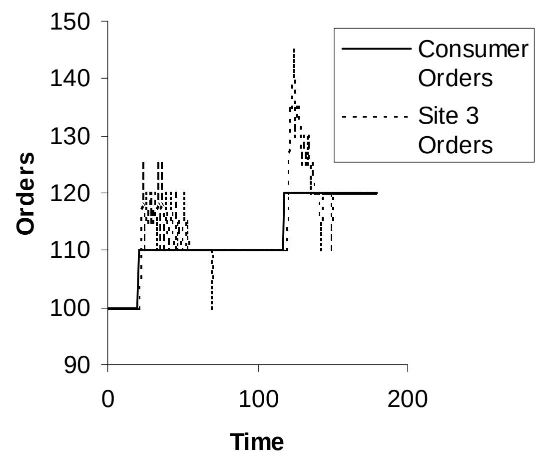

Fig. 7. Persistence of a One-Time Disturbance

For the first step increase in consumer orders, the forecast window is 39 weeks and the disturbance in site 3 orders persists for between 31 weeks (to the last upward spike over the new demand level) and 47 weeks (to the downward spike). The amplitude of the variability in site 3 orders ranges from a high of 125 to a low of 100, or a total range of 25.

Before the second increase, we reduce the forecast window in both PPIC modules from 39 to 20. The period of variability lasted fewer time steps (between 22 to the last order above 120, and 29 to the final downward spike). But shortening the forecast window has the effect of increasing the amplification. Thus the second set of peaks is taller than the first (ranging from 110 to 145, or a total range of 35).

Thus the weighted forecasting algorithm has the effect of imposing a memory on the system. The longer the forecasting period, the longer the memory, but the lower the amplitude of the variations generated.

Generation of Inventory Variation. Even when top-level demand is constant and bottom-level supply is completely reliable, intermediate sites can generate complex oscillations in inventory levels, including phase locking and period doubling, as a result of capacity limitations.

The consumer has a steady demand with no superimposed noise. The bottom-level supplier makes every shipment exactly when promised, exactly in the amount promised. Batch sizes are still 1, but now we impose a capacity threshold on sites 2 and 3: in each time step they can only process 100 parts, a threshhold nonlinearity. As long as the consumer’s demand is below the capacity of the producers, the system quickly stabilizes to constant ordering levels and inventory throughout the chain. When the consumer demand exceeds the capacity of the producers, inventory levels in those sites begin to oscillate.

Fig. 8 shows the inventory oscillation that arises when demand exceeds capacity by 10%. Site inventories oscillate out of phase with one another, in the form of a sawtooth that rises rapidly and then drops off gradually. The inventory variation ranges from near-zero to the level of demand, much greater than the excess of demand over capacity

Fig. 9 shows the dynamics after increasing consumer demand to 150. The inventories settle to a sawtooth with a shorter period. Now one cycle's production of 100 can support only two orders, leading to a period-three oscillation. The inventories of sites 2 and 3, out of synch when Demand/Capacity = 110/100, are now synchronized and in phase.

The transition period is actually longer than appears from Fig. 9. The increase from 110 to 150 takes place at time 133, but the first evidence of it in site 2's dynamics appears at time 145. The delay is due to the backlog of over-capacity orders at the 110 level, which must be cleared before the new larger orders can be processed.

Fig. 10 shows the result of increasing the overload even further. (Because of the increased detail in the dynamics, we show only the inventory level for site 2.)

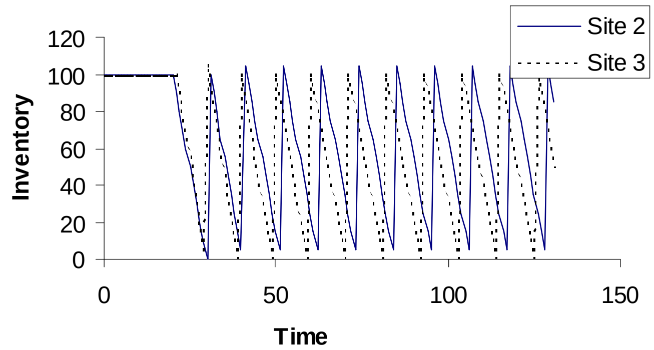

Fig. 8. Demand/Capacity = 110/100

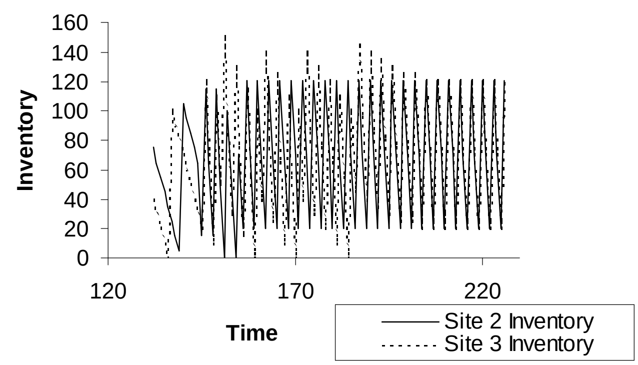

Fig. 9. Demand/Capacity = 150/100

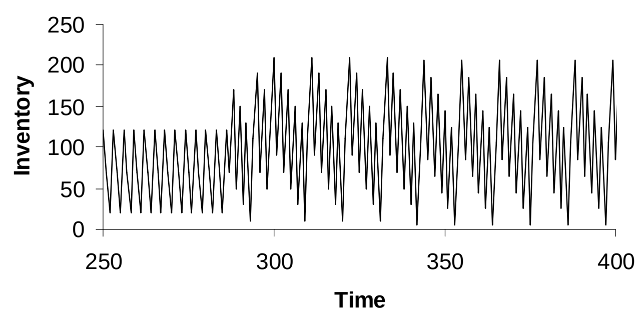

Fig. 10. Demand/Capacity = 220/100 (Site 2)

Now the consumer is ordering 220 units per time period. Again, backlogged orders at the previous level delay the appearance of the new dynamics; demand changes at time 228, but appears in the dynamics first at time 288, and the dynamics finally stabilize at time 300.

This degree of overload generates qualitatively new dynamical behavior. Instead of a single sawtooth, the inventories at sites 2 and 3 exhibit biperiodic oscillation, a broad sawtooth with a period of eleven, modulated with a period-two oscillation. This behavior is phenomenologically similar to bifurcations observed in nonlinear systems such as the logistic map, but does not lead to chaos in our model with the parameter settings used here. The occurrence of multiple frequencies is stimulated not by the absolute difference of demand over capacity, but by their incommensurability.

Details of these behaviors are discussed in [13].

2.3 An Equation-Based Model

Following the pioneering work of Jay Forrester and the System Dynamics movement [5], virtually all simulation work to date on supply chains integrates a set of ordinary differential equations (ODE’s) over time. It is customary in this community to represent these models graphically, using a notation that suggests a series of tanks connected by pipes with valves. Fig. 11 gives a simple example of the flow of material in the DASCh network as it appears in the VenSim® simulation environment [15].

The *rectangular boxes* (“Finished2,” “WIP2,” “Finished3,” “WIP3”) are “levels,” or variables whose assignments change over time. In this particular model, they represent the four critical inventory levels in DASCh sites 2 and 3, a work-in-process inventory and a finished goods inventory for each of the sites.

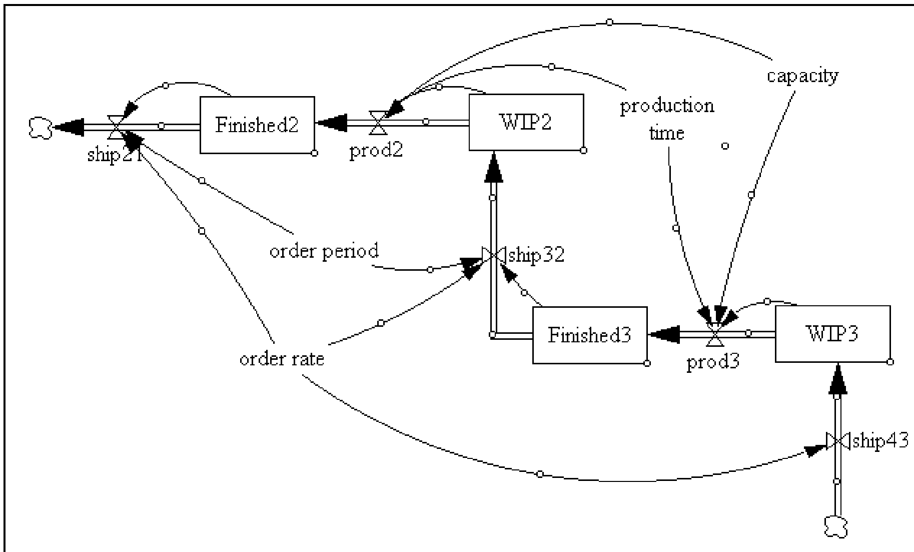

Fig. 11. System Dynamics (ODE) Model

The *arrows with valve symbols* (shaped like hour-glasses: "ship21," "prod2," "ship32," "prod3," "ship43") are flows between the levels that they connect, and the valves are "rates," variables that determine the rate of flow. For example, the rate "prod2" is the rate at which site 2 converts work-in-process to finished goods.

The *cloud shapes* at the upper-left and lower-right corners represent external sources and sinks. In this case, the upper-left corner represents the end consumer (DASCh site 1), while the lower-right corner represents the supplier (DASCh site 4).

The *legends* associated with neither a box nor a valve ("order rate," "order period," "production time," "capacity") are auxiliary variables.

Single-bodied arrows show the dependency of rates on other variables (both levels and auxiliaries) in the model. The exact mathematical form of the dependency is not shown in the graphical representation, but is encoded in the rate. For example, the arrows show that the rate "prod2" depends on the level "WIP2" and the auxiliaries "production time" and "capacity." The actual dependency is prod2 = min(capacity, WIP2/production_time).

This particular model shows the interplay between site capacity and order rate. When the order rate exceeds the site capacity, it demonstrates oscillations comparable to those in the DASCh model. For example, Fig. 12 shows the biperiodic oscillations for Demand/Capacity = 220/100. The system dynamics model shows the same periodicities as the agent-based model. It does not show many of the effects that we observe in the ABM and in real supply networks, including the memory effect of backlogged orders, transition effects, or the amplification of order variation. Such effects require the explicit representation of levels and flows for orders as well as parts. In particular they require a model of PPIC in the system dynamics formalism, which is (as we shall see) not easy to produce.

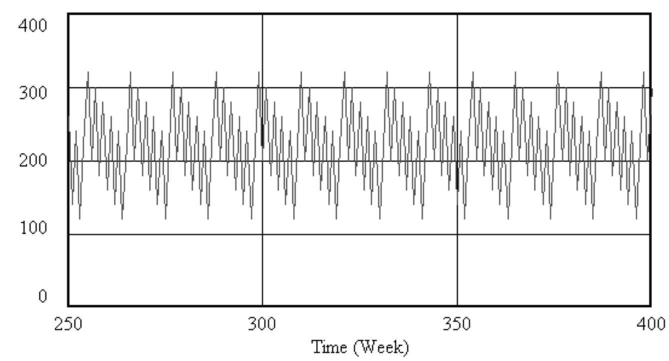

Fig. 12. Inventory Oscillation in an ODE Model

System dynamics models of this nature are widely used in studying organizational dynamics, business processes, environmental and ecological systems, policy implications, and a wide range of similar domains. In principle, ABM can be applied to all of these domains, often in a way that seems more natural.

3. Agents vs. Equations: A High-Level View

ABM and EBM share some common concerns, but differ in two ways: the fundamental relationships among entities that they model, and the level at which they focus their attention.

Both approaches recognize that the world includes two kinds of entities: individuals and observables, each with a temporal aspect.

Individuals are bounded active regions of the domain. In some domains, the boundaries that set individuals apart are physical, as when we are studying ants, or bees, or people. In other domains, the boundaries may be more abstract, as in the case of DASCh’s sites, each representing a business firm. In any case, the boundaries are such that those who are interested in the domain recognize the individuals as distinct from one another. They are “active regions” because those interested in the domain conceive of the individuals as having behaviors. Individuals “do things” as time passes.

Observables are measurable characteristics of interest. They may be associated either with separate individuals (e.g., the velocity of gas particles in a box) or with the collection of individuals as a whole (the pressure in the box). In general, the values of these observables change over time. In both kinds of models, these observables are represented as variables that take on assignments.

Each of these sets of entities invites us to articulate the relationships that unify it and show how those relationships predict the behavior of the overall system through time. The first fundamental difference between ABM and EBM is in the relationships on which one focuses attention.

EBM begins with a set of equations that express relationships among observables. The evaluation of these equations produces the evolution of the observables over time. These equations may be algebraic, or they may capture variability over time (ODE’s, as used in system dynamics) or over time and space (partial differential equations, or PDE’s). The modeler may recognize that these relationships result from the interlocking behaviors of the individuals, but those behaviors have no explicit representation in EBM.

ABM begins, not with equations that relate observables to one another, but with behaviors through which individuals interact with one another. These behaviors may involve multiple individuals directly (foxes eating rabbits) or indirectly through a shared environment (horses and cows competing for grass). The modeler pays close attention to the observables as the model runs, and may value a parsimonious account of the relations among those observables, but such an account is the result of the modeling and simulation activity, not its starting point. The modeler begins by representing the behaviors of each individual, then turns them loose to interact. Direct relationships among the observables are an output of the process, not its input.

Fig. 13 summarizes the critical relationships:

- Individuals are characterized, separately or in aggregate, by observables, and affect the values of these observables by their actions.
- Observables are related to one another by equations.
- Individuals interact with one another through their behaviors.

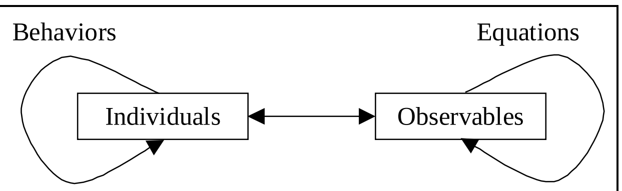

Fig. 13. Unifying Multiplicities

A second fundamental difference between ABM and EBM is the level at which the model focuses. A system is made up of a set of interacting individuals. Some of the observables of interest may be defined only at the system level (e.g., the pressure of an enclosed gas), while others may be expressed either at the individual level or as an aggregate at the system level (e.g., location of an organism vs. the density of organisms per unit space of habitat). EBM tends to make extensive use of system-level observables, since it is often easier to formulate parsimonious closed-form equations using such quantities. In contrast, the natural tendency in ABM is to define agent behaviors in terms of observables accessible to the individual agent, which leads away from reliance on system-level information. In other words, the evolution of system-level observables does emerge from an agent-based model, but the modeler is not as likely use these observables explicitly to drive the model’s dynamics as in equation-based modeling.

These two distinctions are tendencies, not hard and fast rules. The two approaches can be combined [4]: within an individual agent in an ABM, behavioral decisions may be driven by the evaluation of equations over particular observables, and one could implement an agent with global view whose task is to access system-level observables and make them visible to local agents, thus driving an ABM with system-level information. Furthermore, while agents can embody arbitrary computational processes, some equation-based systems (those based on PDE’s, but not the simple ODE’s used in system dynamics) are also computationally complete [11]. The decision between the two approaches must be made case by case on the basis of practical considerations.

4. Agents vs. Equations: Practical Considerations

A practitioner is concerned with the underlying *structure* of a model, the naturalness of its *representation* of a system, and the *verisimilitude* of a straightforward representation. This section discusses these considerations with special reference to modeling supply networks. Some of these issues have been discussed by others in the domains of social science [2, 3] and ecology [9, 16] (where ABM’s are usually called “Individual-Based Models”).

4.1 Model Structure

The difference in representational focus between ABM and EBM has consequences for how models are modularized. EBM’s represent the system as a set of equations that relate observables to one another. The basic unit of the model, the equation, typically relates observables whose values are affected by the actions of multiple individuals, so the natural modularization often crosses boundaries among individuals. ABM’s represent the internal behavior of each individual. One agent’s behavior may depend on observables generated by other individuals, but does not directly access the representation of those individuals’ behaviors, so the natural modularization follows boundaries among individuals.

This fundamental difference in model structure gives ABM a significant advantage in commercial applications such as supply network modeling, in two ways.

1. In an ABM, each firm has its own agent or agents. An agent’s internal behaviors are not required to be visible to the rest of the system, so firms can maintain proprietary information about their internal operations. Groups of firms can conduct joint modeling exercises while keeping their individual agents on their own computers, maintaining whatever controls are needed. Construction of an EBM requires disclosure of the relationships that each firm maintains on observables so that the equations can be formulated and evaluated. Distributed execution of EBM’s is not impossible, but does not naturally respect commercially important boundaries among the individuals.
2. In many cases, simulation of a system is part of a larger project whose desired outcome is a control scheme that more or less automatically regulates the behavior of the entire system. The agents in an ABM correspond one-to-one with the individuals (e.g., firms or divisions of firms) in the system being modeled, and their behaviors are analogs of the real behaviors. These two characteristics make agents a natural locus for the application of adaptive techniques that can modify their behaviors as the agents execute, so as to control the emergent behavior of the overall system. The migration from simulation model to adaptive control model is much more straightforward in ABM than in EBM. One can easily imagine a member of a supply network using its simulation agent as the basis for an automated control agent that handles routine interactions with trading partners. It is much less likely that such a firm would submit aspects of its operation to an external “equation manager” that maintains specified relationships among observables from several firms.

More generally, ABM’s are better suited to domains where the natural unit of decomposition is the individual rather than the observable or the equation, and where physical distribution of the computation across multiple processors is desirable. EBM’s may be better suited to domains where the natural unit of decomposition is the observable or equation rather than the individual.

4.2 System Representation

The variety of EBM with which we have experimented (ODE’s) most naturally represents the process being analyzed as a set of flow rates and levels. ABM most naturally represents the process as a set of behaviors, which may include features difficult to represent as rates and levels, such as step-by-step processes and conditional decisions. ODE’s are well-suited to represent purely physical processes. However, business processes are dominated by discrete decision-making. This is only one example of representational advantages of ABM’s over EBM’s. More generally:

1. ABM’s are easier to construct. Certain behaviors are difficult to translate into a consistent rate-and-level formalism. PPIC algorithms are an important example. In our attempts to duplicate DASCh results using VenSim®, we were unable to construct a credible PPIC algorithm using the rate-and-level formalism. [17] comments on the complexity of such models, and we have been unable to find an actual example of such a model in the system dynamics literature. Recent enhancements to ithink® reflect such difficulties. The most recent release of this popular system dynamics package includes “black boxes” for specific entities such

as conveyors or ovens whose behavior is difficult to represent in a pure rate-and-level system [6]. One suspects that the only realistic way to incorporate complex decision algorithms such as PPIC in system dynamics models will be by implementing such black boxes, thus incorporating elements of ABM in the spirit of [4].

2. ABM's make it easier to distinguish physical space from interaction space. In many applications, physical space helps define which individuals can interact with one another. Customer-supplier relationships a century ago were dominated by physical space, leading to the development of regional industries, such as the automotive industry in southeast Michigan. Advances in telecommunications and transportation enable companies that are physically separate from one another to interact relatively easily, so that automotive suppliers in Michigan now find themselves in competition with suppliers based in Mexico or the Pacific rim. Such examples show that physical space is an increasingly poor surrogate for interaction space in applications such as commerce. ODE methods such as system dynamics have no intrinsic model of space at all. PDE's provide a parsimonious model of physical space, but not of interaction space. ABM's permit the definition of arbitrary topologies for the interaction of agents.
3. ABM's offer an additional level of validation. Both ABM's and EBM's can be validated at the system level, by comparing model output with real system behavior. In addition, ABM's can be validated at the individual level, since the behaviors encoded for each agent can be compared with local observations on the actual behavior of the domain individuals. (A balancing consideration is that the code needed to represent an agent's behavior in ABM is often longer and more complex than a typical equation in an EBM, and thus potentially more susceptible to representational error.)
4. ABM's support more direct experimentation. Managers playing "what-if" games with the model can think directly in terms of familiar business processes, rather than having to translate them into equations relating observables.
5. ABM's are easier to translate back into practice. One purpose of "what-if" experiments with a model is to identify improved business practices that can then be implemented in the company. If the model is expressed and modified directly in terms of behaviors, implementation of its recommendations is simply a matter of transcribing the modified behaviors of the agents into task descriptions for the underlying physical entities in the real world.

4.3 Verisimilitude

In many domains, ABM's give more realistic results than EBM's, for manageable levels of representational detail. The qualification about level of detail is important. Since PDE's are computationally complete, one can in principle construct a set of PDE's that completely mimics the behavior of any ABM, and thus produce the same results. However, the PDE model may be much too complex for reasonable manipulation and comprehension. EBM's (like system dynamics) based on simpler formalisms than PDE's may yield less realistic results regardless of the level of detail in the representation.

One example in the case of extremely simple agents is the Ising model of ferromagnetic phase transitions in statistical physics. The agent in this model is a single atom in an N-dimensional square lattice of similar agents. Its behavior is to change the orientation of its spin to minimize the energy in its environment. One common and generally useful approach to such systems employs mean field theory, analyzing the behavior of a representative atom under statistical averages over the states of neighboring atoms [14, pp. 430-434]. In some dimensions, this mean field EBM approach may miss the order of the phase transition, predict a phase transition where there is none, or yield an inaccurate temperature for the transition. (In one and two dimensions, the equations defining the Ising model can be solved exactly and analytically without the homogeneity assumptions that lead to the errors of the mean field approach, but such solutions are intractable in higher dimensions.) ABM models that emulate the behavior of individual atoms can be developed for arbitrary dimensions, and are more accurate both qualitatively and quantitatively than the mean field approximation.

In a more complex domain, researchers in the dynamics of traffic networks have achieved more realistic results from traffic models that emulate the behaviors of individual drivers and vehicles, compared with the previous generation of models that simulate traffic as the flow of a fluid through a network [7]. The latter example bears strong similarities to the flow-and-stock approach to supply chain simulation, and encourages us to develop an agent-based approach for this application as well.

Wilson [18] offers a detailed study that compares ABM and EBM using the same system (a predator-prey model). He develops a series of EBM’s, each enhancing the previous one to rectify inconsistencies between the ABM and the EBM. The study assumes that the ABM is the more realistic model, and that the EBM is the appropriate locus for making adjustments to bring the two models into agreement. The initial ODE EBM describes reactions between the two species, but representing dispersal through space requires extending it to a set of spatio-temporal integro-differential equations. These equations, modeling both individual characteristics and dispersal using population averages, lead to qualitatively different behaviors than do ABM’s. For example, ignoring local variation in dispersal leads to limit cycles rather than the extinction scenarios that dominate ABM’s. To correct for these lumped parameter effects, the EBM is interrupted at each iteration of the integration to add a random perturbation to the population parameter at each location and to zero local population levels that fall below specified threshholds.

The disadvantages of EBM in these examples result largely from the use of averages of critical system variables over time and space. They assume homogeneity among individuals, but individuals in real systems are often highly heterogeneous. When the dynamics are nonlinear, local variations from the averages can lead to significant deviations in overall system behavior. In business applications, driven by “if-then” decisions, nonlinearity is the rule. Because ABM’s are inherently local, it is natural to let each agent monitor the value of system variables locally, without averaging over time and space and thus without losing the local idiosyncrasies that can determine overall system behavior. The EBM for our toy four-firm supply network (Fig. 11) does not use averages over individuals, and so does not suffer from this disadvantage. However, real-world supply networks are much larger. The total number of shipping points in the U.S. automotive industry is on the order of 40,000, and it is difficult to see how a parsimonious EBM of such a system could avoid the use of lumped parameters.

5. Conclusion

ABM is a relatively new approach to system modeling and simulation. In many domains, it faces entrenched competition from EBM methodologies such as system dynamics. Our experience with both approaches leads to three general recommendations.

First, ABM is most appropriate for domains characterized by a high degree of localization and distribution and dominated by discrete decisions. EBM is most naturally applied to systems that can be modeled centrally, and in which the dynamics are dominated by physical laws rather than information processing.

Second, researchers in agent-based modeling should be aware of the long history of EBM, and should consider explicit case comparisons of their ABM’s with existing or potential EBM’s where relevant. Such comparisons are particularly valuable in simple systems in which one can trace the causes of divergence between the models. The sketch of the relative advantages and disadvantages of the two approaches presented in this paper is preliminary. Our ability to adopt the best modeling approach for a given problem depends on developing a collection of cases that demonstrate the respective strengths and weakness of the two approaches.

Third, the widespread popularity of EBM is due in large measure to the availability of several intuitive drag-and-drop tools for constructing and analyzing system dynamics models. Widespread realization of the advantages of agent-based modeling will depend on the availability of comparable tools for this approach, and the ABM community should encourage the development and refinement of such tools.

Acknowledgements

DASCh was funded by DARPA under contract F33615-96-C-5511, and administered through the AF ManTech program at Wright Laboratories under the direction of James Poindexter. The DASCh team includes Steve Clark and Van Parunak of ERIM’s Center for Electronic Commerce, and Robert Savit and Rick Riolo of the University of Michigan’s Program for the Study of Complex Systems. Sections 2.1 and 2.2 report on joint work of this team, drawn from [13]. This paper has benefited considerably from the detailed comments of Robert Axtell.

References

[1] AIAG. Manufacturing Assembly Pilot (MAP) Project Final Report. M-4, Automotive Industry Action Group, Southfield, MI, 1997.

[2] R. Axtell. Three Distinct Uses of Agent-Based Computational Models in the Social Sciences. Brookings Institution, Washington, DC, 1997.

[3] M. E. Epstein and R. Axtell. *Growing Artificial Societies: Social Science from the Ground Up*. Boston, MA, MIT Press, 1996.

[4] P. A. Fishwick. *Simulation Model Design and Execution: Building Digital Worlds*. Englewood Cliffs, NJ, Prentice Hall, 1995.

[5] J. W. Forrester. *Industrial Dynamics*. Cambridge, MA, MIT Press, 1961.

[6] High Performance Systems. ithink - The Premiere Business Simulation Tool from High Performance Systems, Inc. http://www.hps-inc.com/products/ithink/ithink.html, 1997.

[7] K. R. Howard. Unjamming Traffic with Computers. *Scientific American*, (October), 1997.

[8] T. Hoy. The Manufacturing Assembly Pilot (MAP): A Breakthrough in Information System Design. *EDI Forum*, 10(1):26-28, 1996.

[9] O. P. Judson. The Rise of the Individual-Based Model in Ecology. *Trends in Ecology and Evolution*, 9(1 (January 1994)):9-14, 1994.

[10] R. N. Nagel and R. Dove. *21st Century Manufacturing Enterprise Strategy*. Bethlehem, PA, Agility Forum, 1991.

[11] S. Omohundro. Modelling Cellular Automata with Partial Differential Equations. *Physica D*, 10:128-134, 1984.

[12] H. V. D. Parunak. DASCh: Dynamic Analysis of Supply Chains. http://www.iti.org/~van/dasch, 1997.

[13] H. V. D. Parunak, R. Savit, R. Riolo, and S. Clark. Dynamical Analysis of Supply Chains. ERIM/University of Michigan Preprint, 1998. To be submitted for publication. Available at http://www.erim.org/cec/projects/dasch.htm.

[14] F. Reif. *Fundamentals of Statistical and Thermal Physics*. New York, McGraw-Hill, 1965.

[15] Ventana Systems. Ventana Systems Home Page. http://www.vensim.com, 1997.

[16] F. Villa. New Computer Architectures as Tools for Ecological Thought. *Trends in Ecology and Evolution*, 7(6 (June 1992)):179-183, 1992.

[17] M. E. Warkentin. MRP and JIT: Teaching the Dynamics of Information Flows and Material Flows with System Dynamics Modeling. *Proceedings of The 1985 International Conference of the Systems Dynamics Society*, pages 1017-1028, International System Dynamics Society, 1985.

[18] W. G. Wilson. Resolving Discrepancies between Deterministic Population Models and Individual-Based Simulations. *American Naturalist*, 151(2):116-134, 1998.