Explanation Exploration: Exploring Emergent Behavior

Ross Gore, Paul F. Reynolds Jr., Lingjia Tang, and David C. Brogan
University of Virginia
{rjg7v, reynolds, lt8f, dbrogan}@virginia.edu

Abstract

Understanding emergent behavior(s) exhibited in simulations poses an interesting challenge. Emergence can represent a valid behavior arising from seemingly unrelated phenomena, or it can reflect an error in a model or its implementation. We propose a new method for gathering insight into emergent behavior in simulations using the model adaptation technique, COERCE. COERCE allows a user to efficiently adapt a model to meet new requirements and can be employed to explore emergent behavior exhibited in a simulation. A subject matter expert (SME) can coerce a simulation to gather insight into characteristics of the emergent behavior as the simulated phenomenon is driven toward conditions of interest.

1. Introduction

Emergent behavior manifests itself in all sorts of settings but is of particular, growing interest in the agent-based simulation community. Behavior is *emergent* if the behavior is unexpected and stems from interactions of the underlying components of the model [10]. Traffic jams, a phenomenon observed in automobile freeways, serve as a simple example of emergent behavior. In principle, one expects a line of cars on a highway to maintain even spacing and a constant flow rate. However, when vehicles are individually powered and controlled, they may maintain their separation by sensing their relative separation and adjusting speed accordingly. This feedback amplifies inevitable spatial fluctuations to form congested regions, separated by regions of lower than normal vehicle density [14].

Validation of simulations exhibiting emergent behavior poses an interesting challenge [5]. Emergent behavior can be beneficial, for example, if the unexpected behavior allows users to adapt the model to support tasks designers never intended. Emergent behavior can be harmful if it reflects an error in the construction of a model or its implementation [10].

There is a difference between validating a simulation and validating an emergent behavior that arises in a simulation. The former represents an effort to demonstrate expected behavior. The latter is a demonstration of the validity of behavior that was unexpected for a given set of conditions, or experimental frames [17]. Validation of emergent behavior requires accumulation of insight into the behavior and the conditions under which it arises. Then the problem is reframed so that the emergent behavior becomes a part of a set of behaviors one considers valid.

The need to validate emergent behaviors in a given set of experimental frames requires an exploration capability that extends a model beyond its original intended use. Such a need in turn requires adaptation of models to meet new requirements. We are familiar with model adaptation, and propose an approach using our COERCE technology to accomplish it. COERCE provides a semi-automatic path to adapt a simulation to new requirements [2]. As a result COERCE can be employed to test subject matter expert (SME) hypotheses about the characteristics of the emergent behavior.

We propose a capability that allows a SME to observe characteristics of emergent behavior as a simulated phenomenon is driven towards conditions of interest. Due to the complexity of simulations where emergent behaviors frequently occur, the SME often does not know how to drive the simulation to conditions of interest directly. We use the term *conditions of interest* to mean when a specific condition of the simulated phenomenon is maximized, minimized or targeted to an exact point.

The process of applying COERCE to emergent behavior exploration is called Explanation Exploration (EE). EE provides new capabilities to users presented with the problem of explaining emergent behavior in the simulation domain. The exploration considers both parametric and structural model alternatives in the model space creating a more flexible and robust setting for exploration than existing techniques. Semi-automated adaptation of a simulation into new instances that facilitate exploration of parametric and structural alternatives in the model space is the major contribution of our work.

2. Previous Work

The process of exploring possible behaviors for a given simulation is not new. Others within the simulation community have encountered the problem and prescribe different solutions. Design of Experiments (DoE) focuses on determining which factors within a process affect outputs. This section reviews approaches to exploring possible behaviors in the simulation and DoE communities and summarizes our model adaptation technique, COERCE.

2.1 Simulation

Sensitivity analysis has been proposed as a methodology to explore the robustness of the behavior in a simulation [8]. The principle behind sensitivity analysis is to vary the initial parameters of the model by a small amount and rerun the simulation. This allows the SME to understand how sensitive the model is to the initial parameters. However, even with a small number of parameters, the number of combinations of parameter values quickly becomes large and the resources required to perform the analysis can be excessive [8].

Sensitivity analysis has been refined to exploratory analysis. Exploratory analysis can be viewed as sensitivity analysis done efficiently [4]. Exploratory analysis relies on user insight to limit the number of parameters that need to be explored to gain a broad understanding of the potential behaviors of a model. The approach is characterized in part by parametric exploration. Parametric exploration involves conducting model runs across cases defined by discrete values of the parameters within their plausible domains. The differences in the outcomes of the model runs are examined to determine the parameters or inputs that affect the behavior [4].

EE has a different goal from exploratory analysis and sensitivity analysis. The latter are concerned with gaining a broad understanding of many possible model behaviors. EE focuses on gaining understanding to one unexpected model behavior. Different goals result in different approaches. EE broadens exploration to include parametric and structural alternatives in the model, in order to gather more insight into the unexpected behavior. Exploration of a structural alternative entails examining alternatives for a discrete decision such as the selection of an equation, to represent a smaller phenomenon within the larger model. An example of exploration of a structural alternative is the consideration of a simplified 2-step chemical reaction model and a more complex chemical reaction model consisting of 17 chemical species in 39 elementary reactions in a larger combustion model [2].

Simulation Cloning, a scheme for interactively testing what-if and alternative scenarios in parallel simulations, has also been developed. Potential choices are specified as decision points in the simulation, once a decision point is encountered different outcomes are computed by cloning the simulation and executing the simulation clones concurrently [9]. Simulation cloning is orthogonal to COERCE and EE. It can be employed to improve the performance of COERCE and thus EE, but the goals are different. Exploring the outcomes of interactive simulations based on interactive decision points is different from exploring unexpected behaviors by semi-automatically driving the simulation towards conditions of interest the SME may not know how to create directly.

2.2 Design of Experiments (DoE)

Design of Experiments (DoE) refers to experimental methods used to quantify indeterminate measurements of factors and interactions between factors statistically through observance of forced changes made methodically as directed by mathematically systematic tables [13]. Well-established DoE models exist. Recently, the RRS algorithm for efficient exploration of the parameter space to improve DoE performance was published [16]. RRS was effectively used to study how network protocols interact [1]. EE can be employed with the same goal as DoE: to determine the factors which affect an output. When used in this manner, EE provides a more robust and flexible setting for experimentation than standard DoE models. EE allows both structural and parametric alternatives in the model space to be explored, standard DoE models only allow parametric alternatives to be explored.

Furthermore, EE allows a SME to observe characteristics of emergent behavior as a simulated phenomenon is driven towards conditions of interest. Due to the complexity of simulations where emergent behaviors frequently occur, the SME often does not know how to create conditions of interest directly. Standard DoE models assume that the SME knows how to create conditions of interest; EE does not require this assumption.

2.3 COERCE

Exploring emergent behavior requires an iterative process of careful modification to the original experiment. In the realm of simulation, experiment modification translates to model adaptation. We have developed a model adaptation technique, COERCE, which we explain next.

When constructing a model, abstractions inevitably must be selected in order to reduce complexity, improve performance, or provide estimations for unknown information. When developing *coercible* simulations a SME identifies a set of *abstraction opportunities and alternatives* for each model abstraction. A flexible point of a simulation reflects one model abstraction opportunity and the corresponding bindings for the flexible point reflect abstraction alternatives. According to [2]: "The coercion process involves two roles, the SME and the software developer. COERCION proceeds as follows:

1. When a new requirement arises, the SME identifies model abstraction opportunities that relate to this new requirement.
2. The SME also identifies model abstraction alternatives for each opportunity.
3. The developer identifies the flexible points that reflect these model abstraction opportunities.
4. The developer and the SME use optimization (automatic function minimization) and/or manual modification to find new bindings for the selected flexible points. The SME may interrupt this step if it becomes apparent that a satisfactory set of bindings will not be found.
5. The SME evaluates the behavior of the simulation with the best flexible point bindings found so far. If the behavior still does not meet the requirement, the SME identifies the abstraction opportunities and alternatives that relate to the remaining differences between the behavior and the new requirement. The SME also collects additional insight about model abstraction opportunities by observing how the optimization and modification steps did (or did not) affect the simulation's behavior.
6. The process repeats until the new requirement is met."

COERCE is a semi-automatic process. The process of finding the best flexible point bindings to satisfy the objective function describing the new requirements has two parts: optimization and manual modification. The optimization portion is automatic. Manual modification requires changing simulation source code; it is not automatic. Each optimization or manual modification results in a simulation instance. Observation of each of the simulation instances allows the SME to iteratively gather insight into the behavior of the simulation and the chosen flexible points as COERCE moves the instances toward meeting the new requirements.

The iterative nature of COERCE is described by: $S_0 \rightarrow_p S_1 \rightarrow_p \ldots \rightarrow_p S_{n-1} \rightarrow_p S_n$ where $p = o \mid m$ Here S_0 is the original simulation, S_n is the simulation instance meeting the specified new requirements, and $\rightarrow_p$ is the process of performing an optimization o, or a manual modification to the source code m [15].

3. Explanation Exploration (EE)

EE is a method for increasing insight into unexpected –emergent– behaviors in simulations, and providing a path to validation of those behaviors that are valid. It is important to note that we offer a process for increasing confidence and insight, not a complete validation method or strategy. Our method incorporates semi-automated exploration of conditions in which a user can test hypotheses about emergent behaviors, and thereby increase confidence about their assessment of the meaning of the emergent behaviors. Validation is a goal, but not necessarily an outcome of our method.

When testing a hypothesis about emergent behavior characteristics, a SME may want to observe the emergent behavior as it relates to some other simulated phenomenon as the latter is driven towards a specified set of target behaviors. However, when there are non-linearities in simulated behaviors, the SME may not know how to adapt the simulation to achieve the related target behaviors directly. In general, there are no known techniques for validating emergent behavior efficiently. We advocate the application of COERCE technology for exploration. COERCE enables observation of characteristics of an emergent behavior under controlled conditions of interest.

In EE, COERCE flexible points are utilized to capture both parametric and structural alternatives, and user-guided optimization methods are employed to achieve target simulation behaviors that enlighten a SME about the validity of emergent behaviors. EE offers new capabilities to SMEs needing to explain emergent behavior in their simulations.

3.1 The EE Process

The EE process utilizes COERCE technology. Multiple advantages arise as a result: 1) COERCE flexible points enable capture of a broader range of model abstraction alternatives (both structural and parametric [2]) than a typical parameterized approach supports, 2) because COERCE employs semi-automated search methods, users can explore questions they might not have otherwise investigated, and 3) users can explore relationships between simulation behaviors they understand, but do not necessarily know how to induce, directly or indirectly, and emergent behaviors.

What constitutes a behavior can vary. We do not concern ourselves too deeply with the characteristics of individual behaviors. Of greater importance is how a user relates choices about flexible points to behaviors, and how she relates behaviors to each other. A user will identify either a direct coupling between a set of flexible points and a set of simulation behaviors, or an indirect coupling (or a combination of the two, or neither). A *direct coupling* occurs when the user believes she can explain a set of behaviors as a function of selected choices over a selected set of flexible points. An *indirect coupling* occurs when behaviors are explained as a function of intermediate simulation behaviors, which themselves can be directly coupled to flexible points. The EE process is designed to support both couplings, but is particularly useful for indirect couplings.

In section 3.2 we present a consolidated example of emergent behavior relating to sailing and demonstrate application of our technology to it. It is possible for certain types of sailing craft to attain a forward velocity that exceeds true wind speed. Such craft are capable of exploiting apparent wind, which is a combination of the true wind speed and the craft's own forward velocity [3]. In a sailing simulation, faster than true wind speed sailing could qualify as emergent behavior. We use aspects of the phenomena related to this behavior to elucidate the process description that follows.

We define a set of behaviors of interest for a given execution of a simulation – a "simulation trial" – as:

$$\boldsymbol{\beta_{st(i)} = \{ B_{1,st(i)}, B_{2,st(i)}, \dots B_{n,st(i)}, \},}$$
$$\textbf{st(i): the } \boldsymbol{i^{th}} \textbf{ simulation trial, } \boldsymbol{j = 1,\dots,n} \textbf{ behaviors}$$

It is important that a behavior, $B_{i,st(i)}$, is observable. Some behaviors may not be directly observable in a simulation, if behaviors are instrumentable to a degree satisfactory to the user, they are observable by our definition.

An emergent behavior, E, occurs when some subset of the $\beta_{j,st(i)}$, let it be $B_{j(1),st(i)}$, $B_{j(2),st(i)},\dots$, $B_{j(m),st(i)}$, exhibits a pattern of unexpected behavior(s) across a set, i= 1,…t, of simulation trials.

$$\boldsymbol{E = F(B_{j(1),st(i)}, B_{j(2),st(i)},\dots, B_{j(m),st(i)}),}$$
$$\boldsymbol{i= 1,\dots,t} \textbf{ trials, } \boldsymbol{j = 1,\dots, m} \textbf{ behaviors}$$

where F maps to descriptions of behavioral patterns observed in a given set of simulation trials, e.g. "the velocity of the sailboat is sometimes greater than the true wind speed when the sailboat's orientation is near perpendicular to the true wind direction."

Given an emergent behavior, a user must establish if expectations regarding simulation behaviors need to be modified to include the emergent behavior. Alternatively the user may decide the emergent behavior is an error and not valid. EE facilitates this decision process. The user generally needs to formulate hypotheses about the relationship between flexible points and variations of E, manifested as a function of bindings chosen for the flexible points. Informally a direct coupling hypothesis is:

Direct Coupling Hypothesis: Within selected sets of bindings for a selected set of flexible points, predictable behavior E_{dc} related to E will be manifested in accordance with user expectations.

Formally, U_{dc}= selected flexible point sets, V_{dc}= selected flexible point bindings for selected flexible point sets, and $FP_{set(x),\ bindings(y)}$ = the set of flexible points x with bindings y. E_{dc} is an expected set of behaviors, related to emergent behaviors E. E_{dc} may be a function of x and y, namely, flexible points and their bindings, however we have chosen not to show it as $E_{dc}(x,y)$ because there are cases where it may be independent.

$$\boldsymbol{H_{direct}\text{:}\quad \forall x \in U_{dc},\ \forall y \in V_{dc},\ FP_{set(x),\ bindings(y)} \rightarrow E_{dc}}$$

In the sailing example, a candidate direct coupling hypothesis may be that boat speed is directly related to angle of incidence between true wind direction and boat hull orientation. A flexible point chosen for the direct coupling hypothesis may be hull orientation. While hull orientation is an important factor, and a user may be able to observe a relationship between hull orientation and boat velocity, there are cases where the prediction would break down because sail orientation is also important. The situation may call for the use of an indirect coupling hypothesis.

A user may not be able to hypothesize a direct link between flexible point choices, bindings for those flexible points and expectations about an emergent behavior. However, it may be possible to identify instrumentable conditions within the simulation that can be related directly to emergent behavior expectations. If the user can then identify flexible points that relate directly to the intermediate conditions then a composition of the direct relationships yields a direct relationship between the flexible points and the emergent behavior. However, it is often the case that the user does not know how to make the intermediate conditions occur directly. If she can offer possible relevant sets of flexible points and bindings then a hypothesis may be testable with the support of search methods, including COERCE. Informally an indirect coupling hypothesis is:

Indirect Coupling Hypothesis: For a range of allowable sets of bindings for a range of allowable flexible points, there are cases when intermediate condition C arises. When C arises, behaviors E_{ic} related to emergent behaviors E will be manifested in accordance with user expectations. Because the user does not know which specific flexible point sets or bindings will cause condition C to arise, search will be employed.

Formally, for U_{ic}= allowable flexible point sets, V_{ic}= allowable flexible point bindings for allowable flexible point sets, the indirect coupling hypothesis, $H_{indirect}$, applies to predicate C(x,y) which is true when intermediate condition C arises for flexible points x taking on bindings y, and predicate E_{ic} which is true when predictable behaviors related to emergent behaviors E arise.

$$\mathbf{H_{indirect}} : (\exists x \in U_{ic}, \exists y \in V_{ic}, C(x,y)) \wedge (\forall x \in U_{ic}, \forall y \in V_{ic}, C(x,y) \rightarrow E_{ic})$$

The conditions under which hypothesis $H_{indirect}$ is true would be established using COERCE. The relationship between C and E_{ic} is conjecture on the part of the user, to be established by the outcome of testing the indirect coupling hypothesis.

In the sailing example, C may be the condition when the sail is *full of wind*. In sailing *full of wind* describes when wind flow is smooth over both surfaces of the sail. Maintaining a sail full of wind is a condition that depends partly on boat speed. As the boat accelerates the sail must be oriented increasingly towards the boat's forward direction. A user may be able to hypothesize a direct relationship between sail fullness and faster than wind speed hull velocity, but the user may not be aware of how to directly control sail fullness. If the user can identify the flexible points that can be explored to achieve sail fullness for a given set of conditions, then the user can form the hypothesis linking sail fullness and hull speed, and let search establish cases where the sail is full [3].

3.2 EE and Apparent Wind

We suppose the effect of apparent wind on sailboat velocity is unexpected and observed in a model of a sailboat traveling across a body of water.

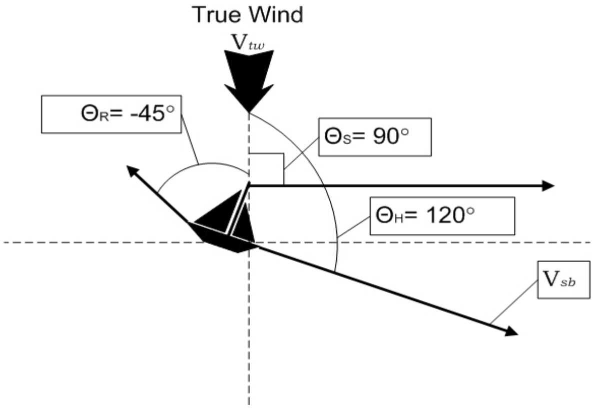

Figure 1: SME Hypothesis $H_{indirect}$

V_{sb}= velocity of the sailboat.

V_{tw}= velocity of the true wind.

Θ_R= angle between true wind and the sailboat rudder.

Θ_H= angle between true wind and the sailboat hull.

Θ_S = angle between true wind and the sailboat's sail.

The EE Process is applied as follows :

1. The SME identifies emergent behavior, E, which occurs when $V_{sb} > V_{tw}$.
2. The SME speculates E will exhibit a particular characteristic, E_{ic}, when a condition, C, arises. C = "the sail is full of wind." The SME speculates that Θ_R, Θ_H, Θ_S and V_{sb} determine when C is achieved. This hypothesis is $H_{indirect}$, displayed in Figure 1. Θ_R, Θ_H, Θ_S and V_{sb} are chosen based on the SME's understanding of sailing and wind. If E is correct, it must possess the characteristics described in E_{ic}, when the total effect of the wind is maximized as described in C. If the SME does not observe E_{ic} she gains insight into the emergent behavior to contribute to its invalidation.
3. The SME identifies model abstraction opportunities related to Θ_R, Θ_H, Θ_S and V_{sb}.

4. The SME also identifies model abstraction alternatives for each model abstraction opportunity.
5. The developer identifies the flexible points, FP_1, FP_2,… FP_k, that reflect these model abstraction opportunities.
6. Formally, the SME forms the hypothesis to test. The hypothesis is:
   $H_{indirect}$: $(\exists x \in FP_1 \ldots FP_k, \exists y \in \text{bindings for } FP_1 \ldots FP_k, C(x,y)) \wedge (\forall x \in FP_1 \ldots FP_k, \forall y \in \text{bindings for } FP_1 \ldots FP_k\ C(x,y) \rightarrow E_{ic})$ The SME believes E_{ic} will be observed when predicate C is true, and C will be true for some subset of possible bindings to the FP_i.
7. The developer and SME use optimization and/or manual modification to find new bindings for the selected flexible points FP_1, FP_2,… FP_k, to discover cases where C is true.
8. The SME observes each simulation instance as it is produced through COERCE. With each observation the SME attempts to confirm or refute $H_{indirect}$.

The SME gathers insight from the hypothesis test in Step 8. She repeats the EE process until she has gathered sufficient insight to validate or refute the emergent behavior, E, in the simulation.

The previous example is not meant to be rigorous, but to illustrate how the EE process is applied to emergent behavior occurring in a fictional model. In Section 4 we present a case study where EE is rigorously applied to an actual simulation we have constructed.

3.3 Possible Issues

Exploration could incur significant execution costs as the process of discovering where condition C is true is carried out. A large number of candidate flexible points may have to be considered, and/or a large number of candidate bindings to currently chosen flexible points may have to be considered. Also, C may never become true. The iterative nature of COERCE supports SME intervention, which can expedite exploration. The expert may halt COERCE and reject her hypothesis before C has been achieved if she observes a sufficient number of iterations of COERCE that she believes that either the truth of C will not be achieved, or that E_{ic} will not be observed when C is true. Similarly, the SME could halt COERCE to accept $H_{indirect}$ before C becomes true, if she believes it will be come true and that E_{ic} will be observed. A SME need not have perfect knowledge of the relationship between achieving C and observing E_{ic} to take advantage of this capability. Knowledge of an important property of the relationship may be sufficient, e.g. the relationship is monotonic.

COERCE must address issues of correctness. Each simulation instance produced by coercion should be analyzed for correctness to the user's satisfaction. Efficient correctness analysis in coercion is addressed with automated lightweight validation [11]. In automated lightweight validation an important subset of requirements, represented as correctness properties, is identified, employing the Pareto Principle: 80% of behavior important to a user is captured in 20% of the correctness properties. Automated lightweight validation is a replacement for traditionally expensive, full validation or even regression testing methods [11].

4. Case Study

To evaluate EE we conducted a case study inspired by combining several of the models considered in [6]. The case study allows agents to interact on a landscape of 2 commodities: sugar and spice. Agents have variable finite lifespans, are able to reproduce, and accumulate wealth through harvesting sugar and spice on the landscape. Agents are allowed to engage in trade of sugar and spice if it increases the welfare of each agent. We have incorporated a separate model to govern the growth of sugar and spice. The growth model is composed of 5 ordinary differential equations, and 5 unknowns. The model uses 3 environmental data inputs and 13 parameters [7].

4.1 Emergent Supply and Demand Curves

In microeconomic theory, the partial equilibrium supply and demand economic model attempts to explain changes in the price and quantity of goods sold in markets. The model describes how prices vary as a result of a balance between product availability at each price, supply, and demand [12].

Supply and demand is calculated in our model by querying the individual agents as to the quantity of sugar each agent is willing to supply or demand at a given price. Price is the amount of spice per unit of sugar an agent is willing to supply or demand. Summing the aggregate quantities of sugar gives the supply and demand plot. The result is shown in Figure 2 by the 'Original Supply' and 'Original Demand' curves. The 'Original Supply' curve and 'Original Demand' curve approach the partial equilibrium supply and demand economic model. This behavior is unexpected in our model. The agents are not

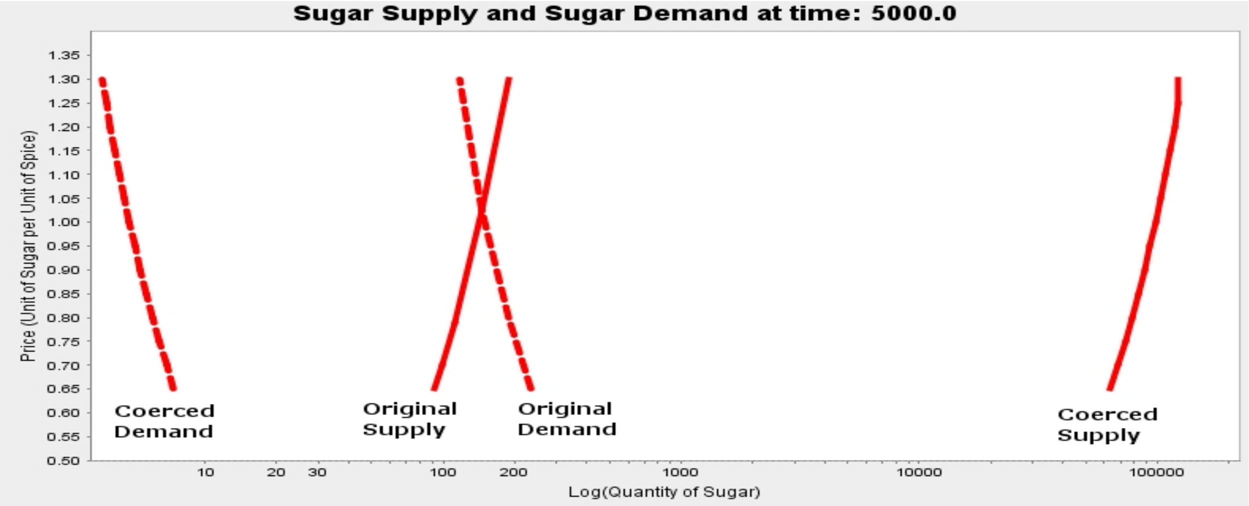

Figure 2: Supply and Demand under original conditions and under coerced maximization of sugar wealth.

programmed to consider supply or demand in their behavior [6].

4.2 Applying EE

To determine if the partial equilibrium economic model is a valid characterization of the supply and demand curves in our model we apply EE. The partial equilibrium supply and demand economic model describes a relative relationship between supply and demand. Specifically, as the supply of a commodity is increased the demand for the commodity will decrease relative to the supply. This relation becomes our indirect coupling hypothesis $H_{indirect}$, where C: sugar supply is maximized E_{ic}: demand is minimized. Because we are uncertain about how to maximize sugar supply, our hypothesis is necessarily an indirect coupling hypothesis. The maximization of sugar supply occurs when the average sugar wealth of each agent in the population is maximized. We use the term *sugar wealth* to mean the amount of sugar each agent owns. Sugar wealth varies in time, so the simulation is executed until a steady state is achieved. The maximization of the sugar wealth for each agent is the condition of interest under which we expect to observe that the demand for sugar will decrease relative to the supply of sugar. An equilibrating point occurs where at a price of P, producers are willing to supply Q units and buyers demand the same quantity. We coerce our case study simulation using 7 parametric flexible points and 2 structural flexible points to achieve C, the condition of interest. The parametric flexible points we coerce across are:

- Maximum agent strength of vision
- Maximum agent sugar metabolism.
- Minimum agent sugar metabolism.
- Maximum sugar wealth of the initial agents.
- Minimum sugar wealth of the initial agents.
- Percent of edible of the sugar growth model.
- Scaling factor of the sugar growth model.

The structural flexible points we coerce across are:

- Alternative models of agent reproduction varying in the complexity of reproductive factors they include.
- Alternative models of endowment parents supply to their children at birth varying in terms of parenting philosophy.

The comparison of supply and demand under the original simulation conditions and under condition C, the maximization of sugar wealth, is displayed in Figure 2. Our hypothesis, $H_{indirect}$, is correct. When sugar wealth is maximized there is a significant decrease in the demand for sugar, resulting in the absence of an equilibrium point. The observed emergent behavior under the maximization of sugar wealth (E_{ic}) corresponds with the partial equilibrium supply and demand economic model. From this we gain confidence that the partial equilibrium supply and demand characterization of the emergent behavior is valid. Further investigations with EE into the emergent behavior that corresponds to the partial equilibrium supply and demand model can increase our confidence enough to pronounce the emergent behavior valid.

5. Conclusion

We have proposed a new method for gathering insight into emergent behavior in the simulation domain using our model adaptation technique, COERCE. A SME can coerce a simulation to gather insight into characteristics of the emergent behavior as the simulated phenomenon is driven towards conditions of interest. The SME is not required to know how to directly create the conditions of interest. The process of applying COERCE to emergent behavior exploration is called Explanation Exploration (EE). We have tested EE on a case study combining a differential equation growth model [7] and several of the canonical Sugarscape models [6]. The results are encouraging; EE semi-automatically creates new conditions under which the SME can observe the emergent behavior.

ACKNOWLEDGEMENTS

We gratefully acknowledge support from the DDDAS program at the National Science Foundation (ITR 0426971), as well as from our colleagues in the Modeling and Simulation Technology Research Initiative (MaSTRI) at the University of Virginia.

6. References

[1] D. Bauer, M. Yuksel, C. Carothers, S. Kalyanaraman, "A Case Study in OSPF and BGP Interactions Using Efficient Experimental Design.", *Proceedings of the 2006 Conference on Principles of Advanced and Distributed Simulation*, IEEE Computer Society, Los Alamitos, CA, 2006.

[2] J.C. Carnahan, "Language Support for the Coercible Software Domain", *A Dissertation Proposal*, University of Virginia: School of Engineering and Applied Science, Charlottesville, VA, 2006.

[3] Colgate, S., *Fundamentals of Sailing, Cruising, and Racing*, W.W. Norton & Company, New York City, NY, 1996.

[4] P.K. Davis, "Dealing with complexity: exploratory analysis enabled by multiresolultion, multiperspective modeling", *Proceedings of the 2000 Winter Simulation Conference*, Institute of Electrical and Electronic Engineering, Piscataway, NJ, 2000, pp. 293-302.

[5] P.K. Davis, "New Paradigms and Challenges", *Proceedings of the 2005 Winter Simulation Conference*, Institute of Electrical and Electronic Engineering, Piscataway, NJ, 2005, pp. 293-302.

[6] Epstein, J.M. and R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*, The MIT Press, Cambridge, MA, 1996.

[7] France, J., *Mathematical Models in Agriculture: A Quanitative Approach to Problems in Agriculture and Related Sciences*, Buttersworths Publishing, Boston, MA, 1984.

[8] Gilbert, Nigel and K. Troitzsch, *Simulation for the Social Scientist*, Open University Press, Philadelphia, PA, 1999.

[9] M. Hybinette, R. Fujimoto, "Cloning parallel simulations.", *ACM Transactions on Modeling and Computer Simulation vol. 11 num. 4*, ACM Press, New York City, NY, 2001, pp 378-407.

[10] C.W. Johnson, "What are emergent properties and how do they affect the engineering of complex systems?", *Reliability Engineering and System Safety vol. 91 issue 12.*, Elsevier Ltd., New York City, NY, 2005, pp. 1475-1481.

[11] X. Liu, P. F. Reynolds and D. C. Brogan, "Using abstraction verification of simulation coercion.", *Proceedings of the 2006 Conference on Principles of Advanced and Distributed Simulation,* IEEE Computer Society, Los Alamitos, CA, 2006.

[12] Marshall, A, *Principles of Economics*, Macmillan and Co., London, England, 1891.

[13] Montgomery, D.C., *Design and Analysis of Experiments 6th Edition*, Wiley & Sons, Indianapolis, IN, 2004.

[14] H. V. D. Parunak, R. S. Vanderbok, "Managing Emergent Behavior in Distrubuted Control Systems.", *Instrument Society of America Technology Expo*, ISA Transactions, Anaheim , CA, 1997.

[15] S. Waziruddin, P.F. Reynolds and D.C. Brogan, "The Process for Coercing Simulations", *Proceedings of the Fall 2003 Simulation Interoperability Workshop*, Simulation Interoperability Standards Organization, Orlando, FL, 2003.

[16] T. Ye, S. Kalyanaarman, "A recursive random search algorithm for large-scale network parameter configuration.", *Proceedings of the 2003 ACM SIGMETRICS International conference on Measurement and modeling of computer systems,* ACM Press, New York City, NY, 2003, pp. 196-205.

[17] Zeigler, B.P., H. Praehofer, and T.G. Kim *Theory of Modeling and Simulation 2nd Edition*, Academic Press, Burlington, MA, 2000.