SimEd: Simulating Education as a Multi Agent System

Elizabeth Sklar	Mathew Davies	Min San Tan Co
Dept of Computer Science	Dept of Computer Science	Dept of Computer Science
Columbia University	Columbia University	Columbia University
New York, NY 10027 USA	New York, NY 10027 USA	New York, NY 10027 USA
sklar@cs.columbia.edu	mdavies@cs.columbia.edu	mtc38@cs.columbia.edu

Abstract

This paper describes our efforts in creating SimEd, *a simulation of the education system. The longterm aim of this work is to be able to model the types of interactions and interplays that occur between students, teachers and administrators, resulting in a toolkit and a methodology that will allow policy makers on various levels to experiment with their decisions, examining effects over time and across levels. Here, we outline our reasoning for selecting the multi agent system paradigm and the particular techniques chosen. We present details of two components within the SimEd toolkit and show results of experimental simulations executed with each. In constructing these tools, we have begun to identify special attributes of the education environment that set it aside from more typical agent-based application areas, such as e-commerce. We highlight these aspects and discuss current and future work extending agent techniques to accommodate these types of environments.*

1. Introduction

We are constructing a simulation of the education system which we call *SimEd*. Our longterm aim is to model the types of interactions and interplays that occur between students, teachers and administrators, in order to allow policy makers on various levels to experiment with their decisions, examining effects, over time and across levels. We have built an agent-based, 3-tiered electronic institution consisting of classrooms, schoolhouses[1] and school districts[2]. At each level, agents interact directly with each other; and their actions affect agents at the same level, and at the same time indirectly affect agents at all levels.

1 We use the term "schoolhouse" to refer to one school.

2 We use the term "school district" to refer to multiple schools that fall under the same administrative and funding umbrella. In the US, school district definitions vary from state to state. A single school district typically refers to all the schools in a single city or county.

We have taken a multi agent-based approach because the environment we are modeling is complex and dynamic, consisting of independent, self-interested entities that conform to certain roles within an organized hierarchy, but decentralized when it comes to individual behaviors. This work is different from work in the area of pedagogical agents [23, 26, 18] and work on intelligent tutoring systems (ITS) [11, 7, 12, 19, 27, 31, 20] because we are not constructing agents that will interact directly with humans and act as tutors. SimEd will not model specific interactions in specific domains but rather more generalized types of interactions that occur. Yet, our model is informed by this work, because our goal is to replicate the behavior of human teachers and learners; and the previous work done to understand human learners in order to create tutoring agents to interact with them is a valuable resource here. Our work is also different from typical multi agent simulation work, which most frequently has been used to simulate e-commerce systems (such as auctions [13]) or robot systems (such as robot soccer [22] and rescue robots [21]). Again, our model is informed by this work, because our goal is to create a realistic simulation system; and previous work done to develop simulations of dynamic environments with many heterogeneous interacting agents provides a valuable resource.

Our eventual goal is the implementation of a complete simulation toolkit that will enable educators and researchers to pose "what-if" questions about policy decisions so that they might be able to predict the effect of their decisions on students and teachers, as well as on the education system as a whole. By far, the most common way we measure whether the education system is succeeding, i.e., whether our children are learning, is by examining test scores. But do these values truly reflect learning? Or are we merely measuring how good our children are at taking tests?

If we step back and examine the education system, it is clear that the system can be viewed as a dynamic environment influenced by many factors [17]. We believe that factors other than test scores can be measured to evaluate the education system. The advantage of using a multi agent simulation to investigate this notion is the inherent ability of such a system to model multiple, interleaving factors that are not typically (or easily) compared (or measured), such as the attitudes of learners, the classroom dynamics of various combinations of abilities, personalities and student-teacher ratios, and the interplay between teachers and administrators. From a theoretical standpoint, educational researchers will be able to use SimEd to test and modify their theories, running experiments over and over again while changing a variety of assumptions and conditions — without having to use real schools as laboratories.

One example of a situation which SimEd would be able to address occurred in California in 1996. Earlier (1985-1989), the *Student-Teacher Achievement Ratio (STAR)* experiment, conducted in real schools in Tennessee, showed that student achievement increased as student-teacher ratio fell [5, 24]. Policy makers in California were anxious to replicate these results, and so they applied the same conditions in California schools; however the results were poor [28, 4]. Decreasing student-teacher ratios led to increased hiring of teachers, which led to severe financial stress within the system as well as a shortage of qualified teachers in California. The overall conditions of the education system in each state were significantly different so that what worked well in Tennessee failed in California. This is exactly the type of situation that can be modeled in SimEd.

We build on the suggestions of Abernathy and Mackie that multi agent simulations can be used to explore policy models and expose factors and assumptions that were not originally considered [6]. Like Abernathy and Mackie, we do not think that multi agent systems should replace standard and accepted empirical methods in educational policy. Rather, we share Robert Axelrod's view of using multi agent systems [9]: SimEd offers a "Third Way" of performing educational policy research. Since agents in multi agent systems are capable of obtaining information, maintaining state, interacting and reacting accordingly, we believe that SimEd will provide insights into factors that cannot be explained or that have been overlooked by widely accepted formal methods in educational research. Such insight could then guide researchers in future educational policy studies.

This paper begins by describing the overall SimEd framework. Then we present two components (tools), along with results of simulation experiments conducted using these tools. We have started to identify attributes of the education environment that set it aside from typical agent-based simulation areas, such as financial markets and robotics. We highlght these aspects and discuss current and future work extending multi agent techniques to support more general multi-disciplinary applications.

2. Framework

SimEd is a toolkit, internally structured as an *electronic institution* [16, 15], which provides a method of organizing (or creating institutional structure around) a group of agents in a multi agent system. Following [15], we use the terms *dialogic framework* to refer to the types of locutions each agent can utter, *scene* to refer to a series of locutions between two (or more) agents (i.e., a formal dialogue game), *performative structure* to refer to a group of scenes (and the rules governing how scenes progress and who can participate in them) and *norm* to refer to the "commitments, obligations and rights of participating agents" [15]. In the work described here, we focus on scenes and norms and use a highly abstracted model to cover the dialogic framework.

The SimEd electronic institution is populated by a set of agents, each of which takes on one of four *roles*: student (S), teacher (T), principal (P) or superintendent (D). In this schema, an agent can only take on one of these roles; i.e., these roles are mutually exclusive.

2.1. Norms

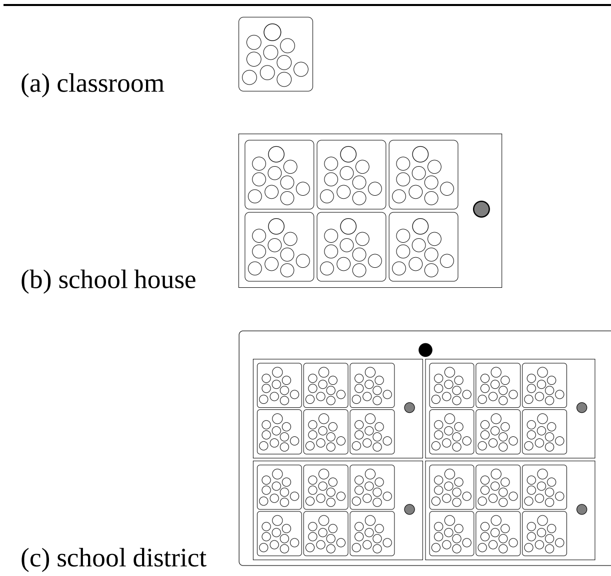

Figure 1. SimEd sub-institutions.

Although each type of agent could, in general, interact with all other types, the most common sets of interactions (and the corresponding *norms* circumscribing those interactions) are specified hierarchically:

1. A classroom Ω consists of one teacher T and a set of students $\{S_0..S_{z-1}\}$ (figure 1a).

2. A schoolhouse Λ consists of a set of classrooms $\{\Omega_0..\Omega_{c-1}\}$ headed by one principal P (figure 1b).
3. A school district Θ consists of a set of school houses $\{\Lambda_0..\Lambda_{l-1}\}$ and one superintendent D (figure 1c).

We refer to classrooms (Ω), schoolhouses (Λ) and school districts (Θ) as *sub-institutions*, and express the hierarchical relationship between them as: $\Theta \succ \Lambda \succ \Omega$. Each of these entities could be specified as independent electronic institutions; however, the behavior of individual agents can have affects on, and be affected by, not only other agents in the sub-institution in which they are defined, but also agents in other sub-institutions. Thus there exists interconnectivity between these sub-institutions which governs the relationships among agents across sub-institutions, both directly and indirectly.

2.2. Dialogic framework

The dialogic framework is based upon a very simple yet rich model of agent interaction derived from the classic Iterated Prisonner's Dilemma (IPD) [8, 10]. Developed in previous work by Sklar and collaborators [25, 29], the *Meta-Game of Learning (MGL)* is essentially a restatement of the IPD within the context of education and is used to describe a student-teacher relationship and, broadly speaking, the types of interactions that can take place between these two types of agents. In the MGL, one agent, the $Teacher$, provides either *easy* or *hard* questions to another agent, the $Student$; and the student responds with either *right* or *wrong* answers (see figure 2). The goal is *student learning*, which theoretically only occurs when the teacher asks hard questions and the student provides the right answers, i.e., when both agents $Cooperate$.

$Student$:	RIGHT $Cooperate$	WRONG $Defect$
$Teacher$: HARD $Cooperate$	*learning*	*frustration*
EASY $Defect$	*verification*	*boredom*

Figure 2. The Meta-Game of Learning (MGL).

It may seem unwarranted to call such an interaction scheme a "dialogic framework" as it is implemented at present, since individual (and sets of) locutions are abstracted away, leaving only their semantic results represented as $Cooperation$ or $Defection$. However, this abstraction is quite purposeful, allowing us to take advantage of a large body of results based on the IPD paradigm. At the same time, we are carrying out related work exploring formal dialogues for education [30]. Nonetheless, even considering the future implementation of a full dialogue system in the traditional sense, the end result of any dialogue game here would still, in the final analysis, be represented abstractly as either cooperation or defection.

2.3. Application

In this paper, we present the application of the SimEd toolkit to two current issues in education and demonstrate how agent-based simulation allows exploration of approaches to each issue from a policy perspective. These issues correspond to two simple interaction models, each based on a sub-institution of SimEd: (1) a classroom Ω with a set of students $\{S_0,..,S_{z-1}\}$ and a single teacher T; and (2) a school district θ with a set of school houses $\{\Lambda_0,..,\Lambda_{l-1}\}$, an implicit principal P for each Λ, and a set of students that move between school houses.

The first issue is *absenteeism*; this is a critical problem, particularly in urban schools, and refers generally to the circumstance of students who are absent so frequently that they fail to matriculate and end up repeating grades multiple times. In section 3, we describe how our classroom tool, by extending the MGL, can be used to evaluate various approaches to the problem of absenteeism.

The second issue is the "No Child Left Behind" Act [1], a national policy in the U.S. that, in essence, allows parents with children enrolled in public schools to transfer their children to another school if their default (typically, the closest proximity) school "needs improvement" according to some standard metric. In section 4, we describe how our school district tool, by utilizing another variation of the MGL, can demonstrate various affects of No Child Left Behind as students move between school houses.

3. Absenteeism and the Classroom Tool

The classroom tool centers around the classroom sub-institution Ω. All activity in Ω is structured in terms of *scenes*, each of which may be thought of as corresponding to a single instruction period. Each scene is composed of a set of interactions among the agents in the classroom:

$$\Omega = T \cup \{S_0..S_{z-1}\}$$

The classroom tool facilitates development of a generic model of classroom learning that can simulate the impact of various factors and can be used to help determine best responses by teachers and/or administrators.

For each student S, we include both external (social) factors and internal (cognitive) factors, which jointly affect perception and learning. These cognitive factors partly comprise the student's internal *state*, which is represented by four abstract values[3]:

- *ability* (A) — indicates the relative ease with which the student can learn a new concept, specifically the maximum the student can learn in an optimal interaction;
- *motivation* (M) — reflects the likelihood of cooperation (in general, if motivation is high, then the student will make an effort to learn; if motivation is low, then the student will not make an effort);
- *emotion* (E) — represents the student's emotional or affective state as a valence (sign) and activation level (intensity), corresponding to some general indicator of mood (e.g., happy);
- *belief set* (Σ) — represents the student's knowledge base or set of beliefs about the world which includes the particular knowledge domain being learned.

External influences on the student arise from the dynamics of multi agent interaction. These dynamics are partly governed by the implementation of the dialogic system in terms of the MGL. The student's interpretation of the teacher's action in a scene (did the teacher cooperate or defect?) depends on not only on the difficulty, or *hardness* (H), of the question the teacher selected, but also on the student's own ability A: as currently implemented, the student's perception of each question as easy or hard depends on the ratio of the student's ability A to the difficulty H. Learning takes place when the student cooperates, or tries to answer correctly; how much is learned depends on both the student's motivation and emotion. Presentation of the particular cognitive model of learning used is beyond the scope of this paper, but specifying how cognitive factors and external factors relate to learning is an important prerequisite whose best expression must be determined experimentally, in the context of the simulation.

The classroom tool iterates over a fixed number of timesteps (for simplicity, each timestep corresponds to a scene in a performative framework of the institution). At each timestep, i, the teacher T asks each student S_j a question. The student's response depends on its internal state and on the external factor H. The student's action also consists of updating its internal state $\{E, M, \Sigma\}$. In related work, we are exploring rules for updating E and M [14].

At time i, a student's behavior is governed by the 5-tuple $(A_i, M_i, E_i, \Sigma_i, H_i)$, where A, M, E and Σ are internal attributes unique to each student. The evolution of the values of internal factors is determined individually for each student by a common (parameterized) update procedure, while the progression of question difficulties is determined by the teacher. Although in theory, all five of these factors might be considered continuous, real, possibly vector-valued, functions of time, for the present we discretize the values. Ability, emotion and motivation are binary or n-ary scalar values. These values range between a minimum (0) and a maximum of $degree - 1$, where $degree$ is a parameter of the simulation that the user can select.

For the simulations described here, we used the following parameter constants: $A.degree = 20$, $M.degree = 4$, $E.degree = 4$, $H.degree = 10$. (Earlier work established that varying these degree values from 2 to 10 affects the precise outcome, but not the important characteristics, of classroom simulations [14].) We assume that ability is innate and hence constant over the course of each simulation. In general, however, ability might change over time. For all students in Ω, the value of A is set randomly using a normal distribution centered around $A.degree/2$. The initial values of M and E are set randomly using a uniform distribution. Σ is initially the empty set (at present, we only model knowledge obtained during the course of a simulation run) and accumulates elements as the student learns. H is deliberately initialized to $A.degree/2$ (e.g., targeting concepts toward the class mean). This implementation is designed to be flexible, allowing for a variety of initialization and update methods, as well as definition of other factors, both internal and external.

We represent the knowledge domain that the students are learning as an ordered set of concepts $\{C\}$. This is the set from which the teacher chooses elements to present to each student, one element per time step. The teacher's presentation ordering could in general be quite complex and interdependent, but in the present model a concept C_k is always presented following C_{k-1}, where C_0 is the first concept presented. If a student "learns" concept k, $S.\Sigma$ updates according to

$$S.\Sigma \leftarrow S.\Sigma \cup (C_k, \gamma)$$

where γ is a value between 0 and 1 that represents the amount of concept C_k that the student actually learned; we refer to this value as *grasp*. We define *progress* to be the average of all γ values corresponding to each elements in Σ.

We conducted two sets of experiments with the classroom tool. In the first set of experiments, γ is binary, so that a concept is either learned (1) or not learned (0). Two different styles of instruction were evaluated for a simulated class of one teacher and 100 students, by comparing the distribution of progress after 180 timesteps, averaged over 10 runs. The results are shown in figure 3. Figure 3a shows what happens when the teacher conducts the class in a traditional lecture format: at each timestep $k = i$, moving from concept C_k to concept C_{k+1}, regardless of each student's γ value. Any student S_j who did not learn C_k at timestep k cannot learn any subsequent concept, and hence experi-

3 Note that the first three values are scalars, while Σ is a set.

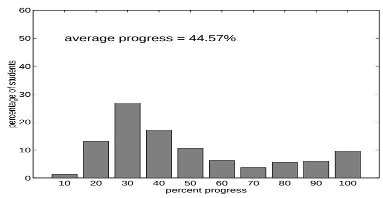

(a) group lecture

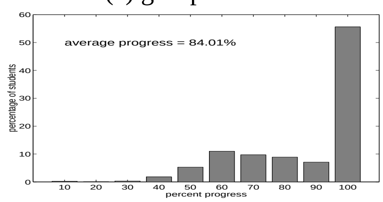

(b) individualized tutor

Figure 3. Different teacher behaviors.

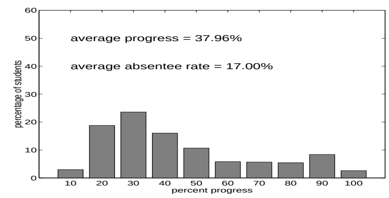

(a) group lecture

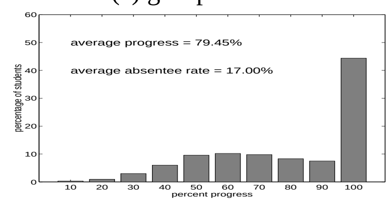

(b) individualized tutor

Figure 4. Effect of absenteeism.

ences no change in progress, until C_k is learned. The graph shows the distribution of students at the end of the lesson sequence, according to their progress, representing the proportion of concepts learned.

In figure 3b, the teacher personalizes her behavior through an abstractly modeled tutoring mechanism that allows her to interact with each student as though on an individual basis, stepping back to easier concepts when frustration sets in, and jumping ahead to more difficult concepts when boredom threatens. In this case, each concept index $S_j.k$ is independent of the timestep i and of other students. The difference between the two graphs is marked, illustrating how effective personalized tutoring can be.

Using these experiments as a control, we can estimate the impact of absenteeism on the mean progress over all the students in the classroom. At each timestep, each student has a chance of being absent, as determined by the absentee rate (set to 17%)[4]. Figure 4a shows the effect of 17% absenteeism on the traditional group lecture classroom. The proportional decrease in average progress is an expected consequence of the simple model of learning used. Teachers and school administrators might hope that some kind of tutoring mechanism would reduce the impact of absenteeism, and just such a result is indicated in figure 4b, where in the classroom employing individualized instruction the 17% absentee rate results in a mere 4.56% reduction in performance. This demonstrates the types of effects SimEd can highlight.

We are currently expanding the classroom tool to account for incomplete or partial learning of concepts (where γ can take on continuous real values between 0 and 1), variable concept difficulty and logical interdependence or "connectedness" between concepts (where γ for C_k depends not only on cognitive factors, but also on the values $\gamma_{0..i-1}$ associated with $C_{0..i-1}$). Since concepts are learned to some degree, tutoring strategies must refer to some *threshold* value at which the concept is considered to be learned well enough. There are many more parameters to take into account in seeking a best strategy or policy. An example of how this might be done, figure 5 plots the classroom mean progress at the end of a set of lessons, as a function of threshold value τ and logical connectivity κ, for three values of κ, showing the nonlinear relationship between κ, τ and progress. When $\kappa = 0$, concepts are not dependent on each other; this was effectively the value used for the set of experiments shown in figures 3 and 4. When $\kappa = 1$, each concept is highly dependent on what came before it, as would be the case in subjects like geometry, where the theorems introduced in class one day are used in proofs the following day.

4 The average absentee rate — percentage of days absent — for New York City public schools in 2001-02 was 17%.

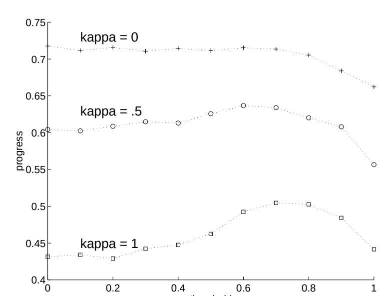

Figure 5. Concept connectedness.

In this section, we have demonstrated how teachers, policy makers and administrators could use SimEd to explore fundamental questions related to classroom management such as organization and curriculum structure and issues such as absenteeism. In the next section, we show how our school district tool can be used to simulate macro-level concerns.

4. No Child Left Behind and the School District Tool

The school district tool centers around the school district sub-institution Θ. Activity is sequenced in *scenes*, where the relationship between a scene and a unit of real time depends on the experiment being conducted; in the example outlined below, one scene equals one school year.

For each student (in each classroom, in each school house) in the school district, we expand on the feature set described in section 3 and introduce factors that are relevant at the macro, school district level:

- *satisfaction (F)* — reflects the student's happiness with his/her current school (low satisfaction indicates that the student wants to transfer to another school), where F can be thought of as a longitudinal combination of E and M described in the previous section;
- *family income (I)* — represents the economic standing of the student's household;
- *mobility (V)* — indicates the student's ability to transfer schools (which is determined not by No Child Left Behind, but rather by a combination of factors such as income and family's ability to relocate);
- *performance (ρ)* — represents the student's academic performance, where ρ can be thought of as a relative comparison of a student's $progress$ (from the previous section) to that of the other students in the same peer group.

The school house sub-institution Λ is represented by a number of scalar-valued features:

- *enrollment* — the set of all students enrolled in the school $\{S_0, .., S_{z-1}\}$, which incorporates implicitly the number of students (z) as well as the feature values associated with a student agent;
- *income* — the average income of families of all students currently enrolled in the school;
- *achievement rate* — this is essentially the number of "good" students, i.e., students who perform better than average on standardized tests;
- *staffing* — the set of all teachers working in the school $\{T_0, .., T_{y-1}\}$, which incorporates implicitly the number of teachers (y) as well as the feature values associated with a teacher agent[5];
- *student-teacher ratio* — the number of students per teacher (a low number is better);
- *capacity* — the maximum number of students the school building can physically handle;
- *funding* — the amount of financing that the school receives, expressed in dollars per student (higher is better);
- *performance* — the school's overall performance, which takes into account all of the above factors.

The school district sub-institution Ω is represented by the following features, many of which are aggregates of those at the school house level:

- *enrollment* — the set of all students enrolled in the school district $\{S_0, .., S_{z-1}\}$, which incorporates implicitly the number of students (z) as well as the feature values associated with a student agent;
- *staffing* — the set of all teachers working in the school district $\{T_0, .., T_{y-1}\}$, which incorporates implicitly the number of teachers (y) as well as the feature values associated with a teacher agent;
- *funding* — the total amount of funding allocated to the district (includes dollars from all sources, local and national, public and private);
- *schools* — the total number of schools in the district; and
- *teachers' salary* — the average salary for all teachers in the district.

We use the School District Tool to demonstrate possible effects of the "No Child Left Behind" Act. No Child Left Behind is an act of U.S. federal government policy, signed in early 2002, designed to "hold states and schools accountable for the academic achievement of all students, ensure that the teaching and paraprofessional staff is highly qualified, and provide parents with access to information and choice." [1] Through this act, parents and children are given "Public School Choice", which allows them to transfer to another public school if their school is listed as "In Need of Improvement." The performance level of each school is determined primarily by the average grade on state standardized tests; schools which fall far below this average are considered to be under-performing.

To consider the effects of No Child Left Behind, we want to show what happens when many students transfer from one school to another. We use the STAR scenario (described in section 1) to vary student-teacher ratios and demonstrate what happens when students are given a choice of changing

5 Although we have not discussed the feature value of a teacher agent in this paper

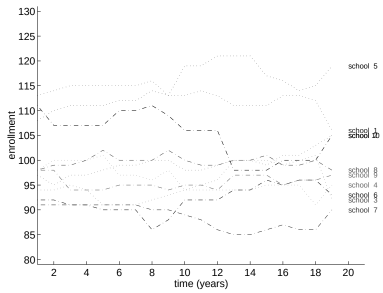

(a) change in enrollments, base scenario

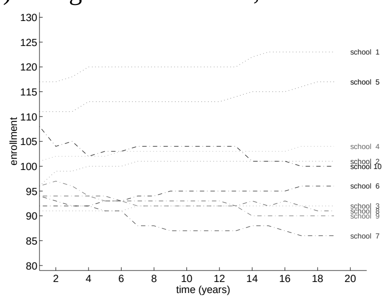

(b) change in enrollments, STAR scenario

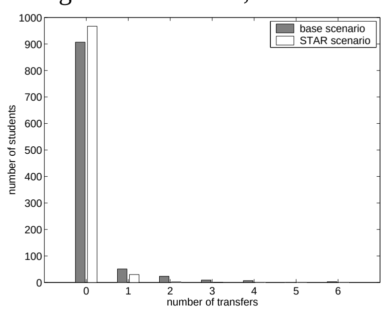

(c) difference in transfer rates

Figure 6. STAR and No Child Left Behind.

schools, based on school performance, following the reasoning behind the No Child Left Behind legislation. Results of a simple experiment investigating the effects of the STAR policy are shown in figure 6. The experiment was run over a 20-year period (i.e., across 20 scenes), using 10 schools, 100 teachers and 1000 students; a comparison of two scenarios was made. In the first scenario, students and teachers were assigned randomly to schools in the district (using a uniform distribution), roughly 10 teachers and 100 students to each school. Students want to transfer when their satisfaction level (which takes into account school performance) drops below a certain threshold; they do transfer if their mobility level is high enough that they are able to transfer. The second scenario imitates the STAR policy and initializes the student-teacher ratio to be lower in some of the schools and higher in others[6].

The top two line graphs in figure 6 contain one line per school and show how enrollments fluctuate at each school over the 20 year period. The bottom bar graph contains a histogram, grouping students according to how frequently they transfer in total, over the 20 year period. The first scenario is illustrated in figure 6a, where students are given the opportunity to transfer to another school, if the school they are attending is under-performing. It can be seen that transfers occur with some regularity, although, as indicated by the leftmost (grey) bar in figure 6c, approximately 90% of students never transfer. The STAR-like scenario (figure 6b) shows that students transfer with much less frequency. This is also evident in the unshaded bars of figure 6c, where nearly 97% of students never transfer.

5. Summary and Future Work

We have presented the basic framework for our multi agent simulation model of the education system which we call "SimEd." We have demonstrated our classroom and school district tools, applying them to two current, real-world issues in education today.

SimEd is structured as an electronic institution, with well-defined agents, roles and interaction models. We discussed the notion of a "sub-institution" as an element within a larger electronic institution framework that could also conceivably stand on its own. In our case, we define three sub-institutions that are related to each other hierarchically ($\Theta \succ \Lambda \succ \Omega$). Our dialogic framework is highly simplistic, but as explained in the text, we are building a formal model of dialogues for education in related work. We plan to incorporate these dialogues into SimEd in the near future. Further experimentation will lead to the addition of performative structures, as the system matures and allows modeling of a wide variety of situations and issues.

Development of SimEd has included the use of packages such as Matlab [2] and RePast [3], but we are driving toward a fully agent-oriented version combining the RePast simulation engine with the ISLANDER electronic institution specification tool [15], working in collaboration with researchers at IIIA who developed ISLANDER. We have consulted, and will continue to consult, experts in education research and policy, psychology, sociology, anthropology and related fields in order to inform our models and advise on the social science aspects of the project. We believe that SimEd, in turn, will be able to help these researchers explore and test their hypotheses using the techniques of multi agent simulation. The challenges we face in the immediate future involve linking together each of the toolkit sub-institutions so that we can demonstrate cause-and-effect across levels of the hierarchy, and the design and

6 The exact allocation was as follows: 4 schools got 15 teachers, 4 schools got 5 teachers and 2 schools got 10 teachers.

implementation of a user-interface that is powerful, yet understandable by the typical social scientist who represents the end-user of SimEd.

6. Acknowledgments

We would like to thank Simon Parsons for his collaboration on the SimEd project. This work was made possible by funding from NSF #REC-01-13317 and NSF #REC-02-19347.

References

[1] http://www.nycenet.edu/nclb/.

[2] http://www.mathworks.com.

[3] http://repast.sourceforge.net.

[4] The Capstone Report: What We Have Learned About Class Size Reduction in California. Technical report, California, Department of Education, 2002.

[5] Class Size: Counting Students Can Count. *Research Points: Essential Information for Education Policy*, 1(2), 2003.

[6] S. F. Abernathy and C. J. Mackie. Skimming, Social Welfare, and Competition Among Public and Private Schools: An Agent-based Perspective. Boston, MA, 2002.

[7] J. R. Anderson. Acquisition of cognitive skill. *Psychology Review*, 89, 1982.

[8] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[9] R. Axelrod. Advancing the art of simulation in the social sciences. *Complexity*, 3(2):16–22, 1997.

[10] K. Binmore. *Fun and games: A text on game theory*. D. C. Heath and Company, Lexington, MA, USA, 1992.

[11] J. S. Brown and R. B. Burton. Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science*, 2(2), 1978.

[12] W. J. Clancey. Intelligent tutoring systems: A tutorial survey. Technical Report STAN-CS-87-1174, Stanford University, 1986.

[13] D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England, 1997.

[14] M. Davies and E. Sklar. Modeling human learning as a cooperative multi agent interaction. In *AAMAS Workshop on Humans and Multi-Agent Systems, at the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2003)*, 2003.

[15] M. Esteva, D. de la Cruz, and C. Sierra. ISLANDER: an electronic institutions editor. In *First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)*, pages 1045–1052, Bologna, Italy, 2002.

[16] M. Esteva, J. A. Rodriguez, C. Sierra, P. Garcia, and J. L. Arcos. On the formal specifications of electronic institutions. *Agent-mediated Electronic commerce (The European AgentLink Perspective, LNAI 1991*, pages 126–147, 2001.

[17] J. Forrester. System Dynamics and the Lessons of 35 Years. In K. B. D. Greene, editor, *Systems-Based Approach to Policymaking*. Kluwer Academic Publishers, 1993.

[18] R. Ganeshan, W. L. Johnson, and E. Shaw. Pedagogical Agents on the Web. In *Proceedings of the Third Int'l Conf on Autonomous Agents*, pages 283–290, 1999.

[19] J. E. Greer and G. I. M. (eds.). *Student Models: The Key to Individualized Educational Systems*. Springer Verlag, New York, 1994.

[20] S. Jackson, J. Krajcik, and E. Soloway. The design of guided learner-adaptable scaffolding in interactive learning environments. In *Proceedings of the Human Factors in Computing Systems Conference (CHI98)*, 1998.

[21] H. Kitano, S. Tadokor, H. Noda, I. Matsubara, T. Takhasi, A. Shinjou, and S. Shimada. Robocup-rescue: Search and rescue for large scale disasters as a domain for multi-agent research, 1999.

[22] H. Kitano, M. Tambe, P. Stone, M. Veloso, S. Coradeschi, E. Osawa, H. Matsubara, I. Noda, and M. Asada. The RoboCup Synthetic Agent Challenge, 97. In *International Joint Conference on Artificial Intelligence (IJCAI97)*, 1997.

[23] J. Lester, B. Callaway, and B. Stone. Mixed initiative problem solving with animated pedagogical agents, 1997.

[24] F. Mosteller. The Tennessee Study of Class Size in the Early School Grades. *Future of Children*, 5(2):113–27, 1995.

[25] J. B. Pollack and A. D. Blair. Co-evolution in the successful learning of backgammon strategy. *Machine Learning*, 32:225–240, 1998.

[26] J. Rickel and W. L. Johnson. Integrating pedagogical capabilities in a virtual environment agent. In W. L. Johnson and B. Hayes-Roth, editors, *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, pages 30–38, New York, 5–8, 1997. ACM Press.

[27] R. Schank and C. Cleary. *Engines for Education*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1995.

[28] W. Schwartz. Class Size Reduction and Urban Students. *ERIC Digest*, (ED472486), 2003.

[29] E. Sklar, A. D.Blair, and J. B. Pollack. Co-evolutionary learning: Machines and humans schooling together. In *Workshop on Current Trends and Applications of Artificial Intelligence in Education: 4th World Congress on Expert Systems*, 1998.

[30] E. Sklar and S. Parsons. Towards the Application of Argumentation-based Dialogues for Education. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-04)*, 2004.

[31] K. VanLehn, Z. Niu, S. Slier, and A. Gertner. Student modeling from conventional test data: A bayesian approach without priors. In *Proceedings of the 4th Intelligent Tutoring Systems Conference (ITS'98)*, pages 434–443, 1998.