# Emergent Collaborative Behavior in a Multiagent Online Learning Simulation


## Student names removed (2 undergraduates)

Department Of Computer Science and Engineering
University of Nebraska - Lincoln
256 Avery Hall, Lincoln, NE 68588-0115 USA



### Abstract

In this paper, we describe the results of simulating collaborative learning amongst students in an online learning environment using supportive agents to make local decisions benefiting their assigned students. Different environments were produced and results were collected to determine whether or not emergent behavior could occur from local decision making by the agents. Here, emergent behavior is defined as converging knowledge levels for all students, as well as the formation of lasting relationships between students. After varying the size of groups, the number of learning activities, and the duration of activities, we discovered that given our agent design, students did learn from one another and their knowledge levels did converge to nearly maximal values. However, due to the complexity of our agent interactions, we were not successful in achieving many lasting relationships amongst students (although most relationships were enduring).


## Introduction

In order to test the effects of local decision making versus global coherence within a multiagent system, we have designed a simulation involving collaborative learning between students in an online learning environment where students are supported by intelligent agents. The simulations were implemented and conducted with the Repast Agent Toolkit [1] in the C# language. We designed our agents to make primarily local decisions but still try to achieve emergent behavior. The decisions made by agents include what actions students should take during learning activities to maximize their knowledge gains, as well as trying to form transitive buddy relationships after activities to continue working with peers beneficial to their students. Our desired emergent behavior includes: 1) converging average knowledge levels for all students, and 2) lasting buddy relationships formed between many students. We created a set of hypotheses positing how different environmental parameters would influence emergent behavior and created a set of experiments to conduct to verify our hypotheses.

The rest of this report is organized as follows: in the next section, we present an overview of our system design, including the design and implementation of both the environment and our agents, as well as a little more detail on the desired emergent behavior from the agents' decisions. Next, we describe our experimental setup, including our hypotheses and experiment design. Then, we present the results of our experiments and try to give justifications and implications for observed trends. Afterwards, we provide a discussion on the results gathered, including evaluating our hypotheses and remarking on problems encountered during the project. Finally, we describe directions and avenues for future work, and we conclude with a brief summary of our work.

## Simulation Design

In this section, we describe the design and implementation of our simulation. We begin by outlining the design of the environment, followed by a discussion on the design of our student agents, concluding with the emergent behavior desired from our simulation.

### Environment Design

The environment is inhabited by one teacher who assigns *A* collaborative learning activities, *S* students who can learn from one another or by themselves during an activity, and *S* student agents, each assigned exclusively to a particular student to make decisions to improve their student's learning. The collaborative activities each last a set duration *D* and occur in a particular knowledge domain -- computer science, history, and music. The teacher is a dummy agent who makes no decisions but only assigns activities and randomly groups students without enough buddies for activities. Every student in the environment is also a dummy agent and each is defined by a set of attributes (randomly generated, ranging from [0, 1]) that govern their learning behavior and skills, including motivation, resourcefulness, learning ability, knowledge (per domain), and communication skill (per mode of interaction – chat, whiteboard, and audio/video). The students begin with an initial knowledge base that increases over time with learning from collaborative activities, while no other attributes change.

The learning experience occurs as follows: to begin the class, a teacher automatically assigns each student to a random group of size *G*. This is necessary because students do not yet know each other in the online learning environment, so they cannot form groups based on local decisions. Next, the first collaborative activity begins.

During the activity, students can either share information with their peers using a specific mode, learn from their peers using the mode chosen by the sharer, or learn independently. Each of these actions will cause the students' knowledge levels to increase by a predictable but varying amount, depending on the students involved and their attributes (as per functions to be defined in the next subsection). Depending on the duration *D* of the activity, sharing and learning will occur multiple times within an activity. After the activity is over, student agents can evaluate the relationships between other students and their assignee and try to form lasting, transitive relationships with other student agents (referred to as "buddy groups") if such a decision benefits the students involved. These buddies remain working together through future activities until one or more agents choose to break the relationship, in which case any remaining buddies still constitute a buddy group. Next, if more activities are left to be performed, the teacher assigns a new activity, filling in groups with randomly assigned members until each group is full.

In this system, the goal of each student agent is to form relationships with other student agents to produce the best learning groups for their student, and to choose the best actions within a learning activity (i.e., sharing, listening, or independent learning) to maximize improvement of their student's knowledge. Such a system is effective if students' knowledge levels increase over time and if lasting group relationships are formed. The system is efficient if students' knowledge increases quickly to a stable level and if full, lasting groups are formed after only a few activities.

## Agent Design

Student agents are rational agents employed to help students create buddy groups and perform activities that maximize the learning benefit for the student. In order to assist their student, the agents control which learning action a student will take during learning activities, find and recruit peers for buddy groups, and perform tracking and learning to improve their performance over time. An agent tracks the knowledge gained by the student when working with different peers in different communication modes, domains, and learning styles. The agent also uses simple machine learning to predict how successful learning will be based on the attributes of other students and how well their attribute combinations match the student.

**Attribute Match Learning:** The machine learning to find good attribute matches for the local student occurs as follows: for each attribute, the range of values is split into four uniform regions (e.g., [0 – 0.25), [0.25 – 0.5), etc.). The agent uses a matrix, where each attribute is a row and the corresponding regions are the columns. Whenever the student learns with another student, the corresponding (attribute, range) entries in the matrix are updated to reflect the average amount of knowledge gained when working with a student with the corresponding attributes. This matrix can then be used later to determine how much knowledge a student can expect to gain, based on the attributes of a peer, where the expected knowledge gain is the average of all the corresponding entries (one for each attribute) for the peer.

**Learning Activities:** During the learning steps of our simulation which begin after receiving the start activity message from the teacher, every student agent first sends updated information about its student's attributes to all other agents in its group in order to help each agent make accurate decisions. Next, each agent determines which of three types of learning is in the best interest of their student to perform. The three types of learning are to self learn in which the student learns alone, sharing in which the student shares knowledge with other students, and listening in which a student listens to a sharing group member. Deciding which action is best is based on taking the maximum expected knowledge gain, as determined by the following formulas parameterized by the attributes of the students involved:

<u>Expected Sharing value:</u>
**Share** (Knowledge Domain a, Communication Mode b) $= \Delta K_T = (K_{aTeach} - K_{aListen}) * M_{Teach} * L_{Teach} * C_{bTeach} * C_{bLearn} * \alpha_T$

<u>Expected Listen value:</u>
**Listen** (Knowledge Domain a, Communication Mode b) $= \Delta K_L = (K_{aTeach} - K_{aListen}) * M_{Listen} * L_{Listen} * C_{bTeach} * C_{bLearn} * \alpha_L$

<u>Expected Self-Learn value:</u>
**Self-learn** (Knowledge Domain a) $= \Delta K_S = K_a * M * L * R * \alpha_S$

Where $K_a$ is the knowledge of the given domain *a*, *M* is motivation, *L* is learning ability, $C_b$ is the communication mode employed, and α values are normalizing constants where:

$$\alpha_T = \alpha_L = \alpha_S / 5.$$

These normalizing constants are useful to prevent any one of the learning actions from being weighted much higher than the others, without respect to the students' attributes.

The learning formulas where created to be simple in nature (not requiring any complex mathematics) while relying on the student attributes involved in the execution of their respective actions. For example, when a student shares information with another student, the success of the activity depends on the knowledge levels of the students, their motivation levels, their learning abilities, and the communication skills of both agents. For self-learning, the communication skills of the agents are not important since they are not working with anyone, but their resourcefulness is important for helping the student find sources of information useful in self-learning.

At the beginning of this action selection process, each agent predicts the expected knowledge gain of their student assuming their student shares with all members of the group. It compares this prediction against the expected self-learn utility, as well as the utility gained from listening to the average group member. If the agent decides it is in the best interest of the student to share, this decision is announced to the rest of the group. If the agent feels listening or self learning would probably best, it remains undecided.

At the next step in the decision process, each undecided agent looks to see who has announced that they will be sharing. These undecided agents will then predict which student is the best to listen to, and calculates if it is better to listen to the optimal sharer or better to self-learn. If the agent decides that listening is in the best interest of the student, the agent then sends a confirm listen message to the sharing agent.

At the third step in the decision process, each agent that announced plans to share will check how many confirm listen messages it has received. If it has received no listen confirmations, it will change its choice to undecided. If it has received listen confirms it will send out a share confirm to the rest of the group.

At the final step in the decision process if the agent is undecided (because it decided to share initially but gathered no listeners), it then compares the optimal confirmed sharers with the option of self learning. If it is best to listen to the confirmed sharers, it will send those agents a confirm listen message. However if it is best to self learn it sets its decision to self learn.

Finally, learning is executed based upon the decisions the agents have made. The actual knowledge gained by the student is calculated using the formulas for expected value gained multiplied with a uniform random variable ranging between 0 and 1. After learning, the agents update their local models of their own student, as well as update their own machine learning and tracking. If the learning activity is not yet over, the agents repeat the learning process.

**Group Formation:** After an entire learning activity (not just the learning process steps within a learning activity) has completed, the teacher sends out a stop activity message and student agents enter the buddy group formation stage. During buddy group formation, each agent tries to form transitive "buddy" relationships with other agents in an effort to maximize the potential group benefit for the upcoming activity. They determine a list of acceptable peers based on the previous history working with a peer (as tracked during learning), as well as through the attribute match learning.

Buddy group formation is accomplished by individual agent decisions whenever possible. When several agents are buddies and another individual agent or set of buddies try to merge buddy groups, the agents will decide the next course of action by voting. Decisions must be made upon a majority vote with an unbiased coin toss breaking ties. All steps of buddy formation are done through message passing in order to reduce the need for a centralized stepping process, increasing the amount of local decision making and eliminating control of a group by one agent alone.

The messages involved in buddy group formation are as follows:

- Express Interest
- Vote Request
- Vote Response
- Request Merger
- Confirm Merger
- Inform Merger
- Back Out
- Repeal Merger

Buddy group formation is a multistep process of message passing. Express Interest messages initiate a buddy group formation negotiation. Vote Request messages query the buddies for their opinion on potential mergers. Voting replies are sent in Vote Response messages. When an agent approaches another agent with the prospect of a merger, it sends a Request Merger message. If a request has been accepted by another agent, that agent sends back a Confirm Merger message. Inform Merger messages are sent to inform current buddies of a confirmed merger. If a confirmed merger must be aborted, a Back Out message is sent to cancel the merger. When an agent is to be removed from the current list of buddies a Repeal Merger message is sent. The intricate decisions required to process these messages are described below.

During the first step of buddy group formation, every agent evaluates every peer their student has interacted with previously. From this they develop the average expected utility of every other student in order to arrive at an estimate of what benefit a randomly assigned student would provide. During this step, the agent also creates a list of students the agent is interested in pursuing as buddies based upon the students expected utility (where utility is measured as expected knowledge gain). Each agent then sends Express Interest messages to students on this list. In order to minimize message passing, we designed each agent to only send Express Interest messages to one interesting agent in each group formation cycle.

When an agent receives an Express Interest message, it computes three values: 1) utility of the sender joining the receiver's group, 2) utility of the receiver joining the sender's group, and 3) the utility of the two groups merging. It then finds the highest utility of these three options. If the receiver agent wishes to invite the other agent to the current group or merge groups, it requests its current group to vote to decide if such a merger should be performed. If the agent would rather leave its current group and join the sender's group, the agent sends a Request Merger message to the sender.

When an agent receives a Request Merger message, it first determines if the merger benefits its student. The agent determines this by comparing its current group utility filled to the maximum with average agents to the utility of the merged group filled similarly. If the receiver agent benefits from the merger and it has no buddies, it would begin the confirmation process. If, however, it is already buddies with at least one other agent, it would request its current buddies to vote on the merger.

When an agent receives a Vote Request message, it evaluates whether the merged group would benefit the agent's student. It then sends its response in a Vote Response message. When an agent who initiated a vote receives the Vote Response messages, it sums all the replies. If the vote is successful, the agent then proceeds with the intended course of action. If it had requested the vote to decide whether or not to make a merger request, it would then make the request. If it had requested the vote to decide whether to accept a received request it would begin the confirmation process.

The confirmation process is a very complicated procedure. To maximize autonomy and allow each agent to pursue mergers that most benefit their student, we allow every agent to negotiate these mergers. There arises the possibility that several agents within a single buddy group guarantee mergers with several other groups. Because voting is based upon the specific merger in question, an agent may vote to merge with group one, and group two, but to merge with both might not be in that agent's best interest. Thus only one merger should be allowed at a time.

When an agent decides to confirm a merger, it is committing its buddy group to that merger. The agent sends a Confirm Merger message to the other group and an Inform Merger message to all of its buddies (including itself). On the next tick, each agent may receive several Inform Merger messages indicating confirmed mergers. The first Inform Merger message is the only allowed merger, and all others must be cancelled. Thus every agent will assume the merger related to the first Inform Merger message has occurred, and add these students to its buddies list. All other confirmed mergers must be cancelled, so each agent must search through Inform Merger messages for messages it has authored. The agent must then send Back Out messages to groups it previously confirmed a merger with. If in this step it has received a Confirm Merger message from another group, it will also send Back Out messages cancelling these mergers

When an agent receives a Confirm Merger message, it sends an Inform Merger message to its group alerting them of the merger. If another group has sent a Back Out message, the agent will send its buddies a Repeal Merger message indicating to them that certain agents have backed out of a merger and must be removed from their buddy lists.

Repeal Merger messages are of a higher priority than Inform Merger messages because if an agent is being removed from the buddy group, then any previously confirmed mergers will be flawed. Thus, if a repeal message is received, all Inform Merger and Confirm Merger messages must be replied to with Back Out messages alerting the other buddies that the mergers must be tried again later. The priority structure of processing messages has this structure:

1. Repeal Merger, Back Out Merger
2. Inform Merger
3. Confirm Merger
4. Vote Request, Vote Response, Request Merger
5. Express Interest

If a message high on the priority list is sent or received, then messages and corresponding actions lower on the list are canceled to avoid concurrency problems (remember that this is a highly distributed problem). Repeal Merger and Back Out messages are placed on the same priority because they both must be processed if received and they also are independent as multiple agents may decide to leave a group without affecting the autonomy of the group. Vote Request, Vote Response and Request Merger are also of the same priority because they can all be sent and received in parallel when a confirmation has not been received. In our current implementation, this process is only run once, thus Express Interest messages are only sent once every buddy group formation period, so its placement on the priority list is currently arbitrary as it is always received independent of the others.

Finally during each step of buddy group formation, each agent is given the option of leaving its current buddy group. To make this decision the agent calculates the utility of a group composed entirely of the average student. If this value is higher than its current group utility value, this indicates the student is better with the teacher's randomly assigned group. If the agent decides to leave its current buddy group, it sends a repeal message to everyone in its buddy group indicating its decision to leave.

## Emergent Behavior

From our environment and agent design, we hope to see two emergent behaviors arise from local decisions: 1) the average knowledge level amongst all students in each discipline will rise and globally converge to relatively high levels (i.e., close to 1) based on the groups formed by agents and the actions chosen within collaborative activities, and 2) lasting buddy relationships between many students would be established that produce successful collaborative learning. Closely tied to behavior 2 is the notion that over time, the number of students not in buddy groups and randomly assigned to groups will diminish, indicating that coherent behavior is achieved through local decisions and not global imperatives.

## Experimental Setup

In this section, we begin with the set of questions and hypotheses we wished to evaluate with our simulations, followed by the design of our experiments created to test those hypotheses.

### Hypotheses

Based on the emergent behavior desired from our system design, we have determined several questions we wish to investigate and have developed a hypothesis for each. These questions and hypotheses include:

***Question 1***: How does group size affect the convergence of knowledge levels?

***Hypothesis 1***: We hypothesize that time to convergence (as measured in ticks) will decrease with group size up until a point, then larger groups will actually cause slower convergence.

***Question 2***: How does group size affect buddy relationship size and duration?

***Hypothesis 2***: Similar to our first hypothesis, we predict that increasing group size will provide agents with more students to consider for forming buddy groups, creating longer lasting relationships, but at some point, increasing group size will make it difficult to fill groups with close-knit buddies, resulting in less coherent buddy groups and more student agents breaking relationships.

***Question 3:*** How do the number of activities and duration of activities affect knowledge convergence?

***Hypothesis 3:*** We predict that the number of activities will not affect the rate of convergence of knowledge levels amongst all students, but convergence will not occur if the number of activities is too low. We also predict that activity duration will be directly proportional to convergence rate because a larger proportion of time will be spent on learning actions in an activity and less on group formation and activity assignment.

***Question 4:*** How do the number of activities and duration of activities affect the length of relationships?

***Hypothesis 4:*** We hypothesize that the number of activities will be directly proportional to the length of relationships because as students participate in more activities, they will encounter more students who could be potential buddies, so agents will have a wider selection of students to choose from when selecting buddies. Similarly, we believe that longer activities will also promote longer relationships because students will work together for longer periods of time.

In our original project proposal, we had more questions and hypotheses, but due to time constraints and the amount of results required to evaluate the first four hypotheses, we have moved the other hypotheses to Future Work, as described in that section.

### Experiment Design

In order to test our hypotheses and collect information about how local decisions lead to global, emergent behavior, we have designed our experiments to vary different environmental parameters (shown in Table 1) and collect data about both relationships and student knowledge levels.

**Table 1: Environmental Parameters with Ranges**

| Parameter | Range of Values |
|---|---|
| Number of students $S$ | 150, 300, 450, 600 |
| Size of Groups $G$ | 3, 5, 10, 15 |
| Number of activities $A$ | 10, 25, 50, 100 |
| Activity Duration $D$ | 50, 150, 300, 500 |

For each possible parameter combination (of which there are $4^4 = 256$), we ran each simulation five times to reduce the variance within each combination. We also selected a different set of random seeds for each run. Since there are five runs per parameter combination, each run receives a different seed. To reduce the number of seeds necessary, we just used four sets of five random seeds, with each simulation with the same number of activities using the same set of seeds, as shown in Table 2. To generate our random seeds, we used the same uniform random generator used by our simulations, created several random numbers, and picked 1 out of every 10 to serve as our seeds.

**Table 2: Random Seeds for Simulations**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **Seed 1** | 7065570 | 6908840 | 5660494 | 11583068 |
| **Seed 2** | 2706659 | 5272011 | 11025165 | 4122670 |
| **Seed 3** | 3108799 | 12012915 | 10424710 | 8080341 |
| **Seed 4** | 2423705 | 3960409 | 2204078 | 9922688 |
| **Seed 5** | 1189151 | 8656125 | 11576508 | 9354767 |

During each simulation, we tracked several key features in order to evaluate both our hypotheses and our system.

These features include the attributes of all students at the beginning of our experiments, their knowledge levels at the end of the simulations, the average knowledge level of all students after each learning action within a learning activity, the constitution of every group and buddy group formed before every activity, the average knowledge of each group and buddy group before and after each learning activity, the average size and age of each buddy group after learning activities, the total number of each type of message passed between agents, and the final tick counts for each simulation. We would like to note that not all of this data was required to evaluate our hypotheses, and as such only those necessary will be used in the remainder of this report.

## Results

In this section, we present the results of our experiments to measure emergent behavior. We begin by evaluating the hypotheses involving converging knowledge levels, followed by the hypotheses dealing with buddy group formation. We conclude this section with some observations about local decisions vs. global control based on the number of messages sent by teachers. We would like to note here that for the tables presented in this section, unless otherwise specified, the numbers of activities are the columns, the rows are the duration of activities, and each size of buddy group is given a separate table. Also worth noting is the fact that for all of the results presented in this section, the number of students was held fixed to 150. This was done for several reasons. First of all, the number of students did not belong as part of any of our hypotheses, so there was no need to vary it. Leaving it constant also reduced one dimension from our analysis, which would have been even more complicated and harder to present if we had included the other buddy group sizes. We chose to use 150 arbitrarily, but partly because it was the fastest to run in case we needed to rerun any experiments (which we did not).

### Knowledge Convergence

In order to test our hypotheses and collect information about the convergence of knowledge levels amongst all students (one form of emergent behavior), we recorded several key measurements. These include the average knowledge amongst all students for all knowledge domains after 500 learning ticks (the duration of the shortest experiment), the average final knowledge levels for all students across all domains, and the percentage of the final knowledge achieved after only 500 learning ticks. These measurements reflect several important factors: 1) how knowledge grows after only a few learning ticks (500 might seem like many, but some of our experiments were 50,000 learning ticks which is much, much larger), 2) what knowledge levels the students approach after many activities, and 3) how close they are to their final values after only a short period of time.

**Knowledge after 500 learning ticks:** Our collected measurements of the average knowledge across all domains for all students after only 500 learning ticks are given in Tables 3 – 6.

**Table 3: Average Knowledge @ 500 Ticks for Group Size of 3**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.53877 | 0.54029 | 0.52235 | 0.53770 |
| **150 Ticks** | 0.53455 | 0.53816 | 0.52565 | 0.53527 |
| **300 Ticks** | 0.58788 | 0.50128 | 0.44769 | 0.49833 |
| **500 Ticks** | 0.58848 | 0.49967 | 0.43338 | 0.50465 |

**Table 4: Average Knowledge @ 500 Ticks for Group Size of 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.54172 | 0.51383 | 0.55288 | 0.54178 |
| **150 Ticks** | 0.57249 | 0.51892 | 0.54045 | 0.53627 |
| **300 Ticks** | 0.59496 | 0.50043 | 0.48616 | 0.53323 |
| **500 Ticks** | 0.59121 | 0.50212 | 0.43613 | 0.50801 |

**Table 5: Average Knowledge @ 500 Ticks for Group Size of 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.54617 | 0.53544 | 0.55601 | 0.54664 |
| **150 Ticks** | 0.57059 | 0.54213 | 0.51099 | 0.55193 |
| **300 Ticks** | 0.57019 | 0.52072 | 0.48075 | 0.55311 |
| **500 Ticks** | 0.59624 | 0.50960 | 0.44157 | 0.51454 |

**Table 6: Average Knowledge @ 500 Ticks for Group Size of 15**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.54908 | 0.55079 | 0.55975 | 0.54764 |
| **150 Ticks** | 0.61367 | 0.56365 | 0.55580 | 0.53433 |
| **300 Ticks** | 0.61508 | 0.49982 | 0.46697 | 0.50094 |
| **500 Ticks** | 0.60010 | 0.51451 | 0.44534 | 0.51842 |

From these tables, we can observe several key points:

- The total number of activities in the simulation seems to affect the average knowledge levels of students after only 500 learning ticks.
- Increasing the duration of activities decreases the early knowledge gains in almost all cases (except for simulations with only 10 activities, which experienced an opposite trend).
- Increasing the group size provides a small but mostly consistent gain in average knowledge after only 500 learning ticks.
- Some average knowledge levels are below the expected starting point of 0.5.

The observation that stands out greatest to us is the first one. Because every simulation runs for at least 500 learning ticks, the number of activities should play absolutely no role in the knowledge convergence of students after only 500 ticks – it should only impact the final knowledge levels. Upon further inspection, we discovered the source of this anomaly. As mentioned previously in our experiment design, we selected a fixed set of random seeds, with one set of seeds for all experiments sharing the same number of activities. Thus, the discrepancy observed here is most likely caused by an odd set of random seeds chosen for the experiments consisting of only 10 activities. Because the knowledge levels of students are randomly generated for each experiment, it is possible that the knowledge levels for these students were already larger than the expected 0.5 level when the experiment began, causing the students to reach higher levels faster (given their higher base knowledge levels), or the discrepancies between their starting knowledge levels could have been higher, emphasizing faster learning through sharing and listening (where gains are proportional to the difference between the students' knowledge levels).

To remove the erroneous affect of random seeding for different numbers of activities from our consideration, we have consolidated Tables 3 - 6 into one table (Table 7), where the columns are now the different group sizes. The values of this table were simply averaged across all activities. This table reinforces our second and third observations previously noted.

**Table 7: Average Knowledge @ 500 Ticks Averaged Across Number of Activities**

| | 3 GS | 5 GS | 10 GS | 15 GS |
|---|---|---|---|---|
| **50 Ticks** | 0.534778 | 0.537553 | 0.546065 | 0.551815 |
| **150 Ticks** | 0.533408 | 0.542033 | 0.54391 | 0.566863 |
| **300 Ticks** | 0.508795 | 0.528695 | 0.531193 | 0.520703 |
| **500 Ticks** | 0.506545 | 0.509368 | 0.515488 | 0.519593 |

Given that we are comparing knowledge levels after only 500 learning ticks, it makes sense that increasing the duration of activities decreases the average knowledge level after 500 ticks. This is because for longer activities, each domain is emphasized less (since only one knowledge domain is used in each learning activity), so some domains do not experience much gain, which is especially true when the duration of activities is 500 ticks where only one learning activity occurs, leaving two domains unused. Also, because more learning activities occur when durations are shorter, students are exposed to more peers, so they are more apt to find students who work well with them to increase their gains. This observation also implies that improvements in student learning occur at a decreasing rate over time, even after only 500 learning ticks. This is evident because otherwise it would not matter how many domains were emphasized in learning. If students learned at a constant rate, we would expect their average knowledge levels to be consistent after 500 ticks regardless of how many activities were performed in those ticks.

Similarly, it makes sense that increasing group size causes increased gains in knowledge levels. While in larger groups, students are exposed to more peers who could work well with them, leading to increased gains in knowledge. However, early on agents have not yet learned enough about other students to make optimal decisions in buddy pairing, so this explains the relatively small size of the increases caused by increased group size.

Finally, although some average knowledge levels are below the expected starting point of 0.5, we do not believe that there is a significant finding here. Due to the random generation of knowledge levels, the students probably began with a low average knowledge level in one of the domains. All of the below 0.5 findings were in the 300 and 500 duration simulations where only 1 or 2 activities are performed within the first 500 learning ticks. If the lowest domain was not selected for learning, it would not improve before 500 learning ticks had occurred. Once its value was averaged with the other domains, the overall knowledge level could be drawn below 0.5.

**Percent of Final Knowledge after 500 Ticks:** In order to assess not only the magnitude of knowledge but also its convergence, we also recorded the measurements detailing how close to our final knowledge values each simulation was after only 500 learning ticks. These measurements are presented in Tables 8 – 11.

**Table 8: % Final Knowledge @ 500 Ticks for Group Size of 3**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 100% | 90.877% | 76.954% | 70.288% |
| **150 Ticks** | 87.425% | 74.695% | 63.772% | 60.274% |
| **300 Ticks** | 84.986% | 60.868% | 49.858% | 53.280% |
| **500 Ticks** | 79.973% | 57.493% | 46.357% | 52.625% |

**Table 9: % Final Knowledge @ 500 Ticks for Group Size of 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 100% | 84.908% | 79.571% | 69.224% |
| **150 Ticks** | 91.972% | 69.665% | 64.178% | 59.793% |
| **300 Ticks** | 84.331% | 60.249% | 53.436% | 56.602% |
| **500 Ticks** | 76.025% | 56.179% | 46.264% | 52.785% |

**Table 10: % Final Knowledge @ 500 Ticks for Group Size of 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 100% | 86.054% | 77.687% | 68.193% |
| **150 Ticks** | 88.806% | 70.236% | 59.639% | 60.754% |
| **300 Ticks** | 80.128% | 60.985% | 52.311% | 58.358% |
| **500 Ticks** | 74.064% | 56.392% | 46.562% | 53.252% |

**Table 11: % Final Knowledge @ 500 Ticks for Group Size of 15**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 100% | 87.273% | 77.345% | 67.764% |
| **150 Ticks** | 94.629% | 72.205% | 64.265% | 58.610% |
| **300 Ticks** | 84.281% | 58.127% | 50.618% | 52.804% |
| **500 Ticks** | 74.987% | 56.702% | 46.914% | 53.574% |

From these measurements, we can observe one major trend:

- After only 500 learning ticks, students have not yet come close to converging on maximal knowledge levels.

These tables clearly indicate that students still have improvements to be made to their knowledge levels, due to the low percentages of final knowledge given for experiments with more time spent on learning activities. This is consistent with our observed knowledge levels presented in Tables 3 – 6 (which were not much higher than the expected 0.5), and indicates that more time spent on learning activities will increase student knowledge much further. This can be accomplished through either increasing the duration of activities, increasing the number of activities, or both.

**Final Knowledge Levels:** Finally, to assess how high knowledge levels converge to in our simulations, we also recorded the final knowledge levels of all students. The averages across all domains for all students are given in Tables 12 – 15.

**Table 12: Final Knowledge for Group Size of 3**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.53877 | 0.594533 | 0.67878 | 0.76499 |
| **150 Ticks** | 0.61144 | 0.72048 | 0.82425 | 0.8880 |
| **300 Ticks** | 0.69173 | 0.82355 | 0.89793 | 0.93530 |
| **500 Ticks** | 0.73584 | 0.86911 | 0.93488 | 0.95895 |

**Table 13: Final Knowledge for Group Size of 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.54172 | 0.60516 | 0.69483 | 0.78264 |
| **150 Ticks** | 0.62246 | 0.74488 | 0.84212 | 0.89688 |
| **300 Ticks** | 0.70550 | 0.83060 | 0.90980 | 0.94208 |
| **500 Ticks** | 0.77766 | 0.89379 | 0.94269 | 0.96240 |

**Table 14: Final Knowledge for Group Size of 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.54617 | 0.62222 | 0.71571 | 0.80162 |
| **150 Ticks** | 0.64251 | 0.77187 | 0.85680 | 0.90846 |
| **300 Ticks** | 0.71160 | 0.85384 | 0.91902 | 0.94779 |
| **500 Ticks** | 0.80503 | 0.90367 | 0.94836 | 0.96624 |

**Table 15: Final Knowledge for Group Size of 15**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.54908 | 0.63111 | 0.72371 | 0.80815 |
| **150 Ticks** | 0.64850 | 0.78062 | 0.86486 | 0.91167 |
| **300 Ticks** | 0.72980 | 0.85987 | 0.92252 | 0.94867 |
| **500 Ticks** | 0.80028 | 0.90739 | 0.94927 | 0.96767 |

These data indicate several important trends, including:

- Final knowledge levels increase laterally with both the number of activities and duration of activities.
- Increasing group size improves knowledge levels up until a certain point, but after awhile, the final knowledge levels between group sizes of 10 and 15 were pretty consistent.
- Knowledge levels experience diminishing returns, increasing at a decreasing rate, indicating convergence without reaching a static plateau.

First of all, as expected, increasing the amount of time spent on learning activities by increasing the number of activities and the duration of activities increases the final knowledge levels. In fact, in some simulations, the average knowledge for all students across all domains reaches some very large values near a perfect 1.0. If we assume that the knowledge level of students is their average grade in different domains, these students are very intelligent! This indicates that students continue to learn the more opportunities they have to do so.

Second, increasing the group size improves student learning performance, but it does so at a greater rate when less time is spent on learning. As mentioned previously, the gains come from working with more peers who are better potential learning partners. However, for longer learning times, the gains slow as everyone spends enough time learning that they converge to their nearly maximal values. At this point, further increasing the number of students they work with does not improve performance.

Finally, although the average knowledge levels continue to climb as more time is spent on learning, the rate of gains decreases drastically. For example, looking at the largest cases, increasing from 50 activities to 100 activities for 500 tick durations, or from 300 tick durations to 500 tick durations for 100 activities only provide average gains of 0.02 in knowledge. Considering that both of these scenarios result in around twice as much learning time (after an already very long time learning), the gains in learning are very small, especially when compared against the gains caused by increasing learning time when learning time was already small. However, it is important to note that there was always at least a small increase in knowledge when increasing the total amount of time spent learning, so although the knowledge levels have converged to *near* maximum levels, they have not yet hit a pinnacle or plateau. This implies that further gains could be had if even more time were spent on learning, but given the sharp diminishing returns, the gains would probably not outweigh the costs.

## Buddy Group Behavior

In order to measure our other desired form of emergent behavior (lasting buddy relationships amongst many students), we recorded several key measurements, including the average buddy group size after 500 learning ticks, the final average buddy group size and age (measured in the number of activities each pair of buddies sticks together), and the maximum buddy group size and age during the experiments.

**Buddy Group Size after 500 Ticks:** Similar to our analysis of knowledge convergence, we wanted to identify how quickly buddy groups were formed amongst students, so we recorded the average buddy group sizes after 500 learning ticks (again the shortest total duration of a simulation). However, unlike our other measurements for buddy group behavior, we did not record the average age at 500 ticks since this a short amount of time and all ages will be low. The results are displayed in Tables 16 – 19.

**Table 16: Average Buddy Group Size @ 500 Ticks For Group Size = 3**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.0053 | 1.00533 | 1 | 1.00267 |
| **150 Ticks** | 1 | 1 | 1 | 1 |
| **300 Ticks** | 1 | 1 | 1 | 1 |
| **500 Ticks** | 1 | 1 | 1 | 1 |

**Table 17: Average Buddy Group Size @ 500 Ticks For Group Size = 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.024 | 1.03733 | 1.03467 | 1.03733 |
| **150 Ticks** | 1 | 1 | 1 | 1 |
| **300 Ticks** | 1 | 1 | 1 | 1 |
| **500 Ticks** | 1 | 1 | 1 | 1 |

**Table 18: Average Buddy Group Size @ 500 Ticks For Group Size = 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.15467 | 1.16533 | 1.12 | 1.112 |
| **150 Ticks** | 1.008 | 1.00267 | 1.00267 | 1.01067 |
| **300 Ticks** | 1 | 1 | 1 | 1 |
| **500 Ticks** | 1 | 1 | 1 | 1 |

**Table 19: Average Buddy Group Size @ 500 Ticks For Group Size = 15**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.18133 | 1.22933 | 1.10933 | 1.15067 |
| **150 Ticks** | 1.04 | 1.02667 | 1.02933 | 1.02133 |
| **300 Ticks** | 1 | 1 | 1 | 1 |
| **500 Ticks** | 1 | 1 | 1 | 1 |

These results show several trends:

- 500 learning ticks is not enough time to form buddy relationships with peers.
- Increasing the duration of activities decreases the average size of buddy relationships.
- Increasing group size increases the average size of buddy groups.

First of all, nearly all of the values recorded for this measurement are near 1, so it is pretty obvious that group formations were not occurring in such a short span of time. This is detrimental to shorter simulations because it means that they will not experience much in the way of emergent group formation behavior. The cause for this is probably due to the lack of confident knowledge agents have about one another. Tracking information and learning about attribute combinations takes time to produce reasonable results, so after only 500 learning ticks, agents themselves have not yet learned enough about one another to adequately assess whether or not others would make good buddies, resulting in a lack of group formation.

Second, as noted in the previous results analysis, for longer activity durations, very few learning activities have occurred so agents have worked with less agents (further reducing available, accurate information about peers), and there have been less opportunities for buddy group formation since this stage only occurs after each learning activity. Because it takes time for agents to learn about one another and form groups, several activities must occur before groups are formed. Since very few activities of long duration can occur within the first 500 learning ticks, this short amount of time is simply not enough to encourage much group formation. However, for short activity durations, agents are exposed to more peers and have more opportunities for buddy group formation, resulting in at least some bonding between agents.

Finally, increasing group size provides at least a small increase in the average size of buddy groups because agents are again exposed to more peers, countering previously mentioned problems. When exposed to a larger variety of students, agents become more aware of whom to work with and are more likely to attempt group formation early on.

**Final Buddy Group Size:** After observing the lack of buddy group formation early in simulations, we wished to observe the final group sizes after all learning and group formation activities. The average sizes of buddy groups at the end of each simulation are presented in Tables 20 – 23.

**Table 20: Final Buddy Group Size for Group Size = 3**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.00533 | 1.06133 | 1.10667 | 1.24267 |
| **150 Ticks** | 1 | 1.04 | 1.1 | 1.14933 |
| **300 Ticks** | 1.016 | 1.05067 | 1.12933 | 1.16533 |
| **500 Ticks** | 1.0053 | 1.05067 | 1.088 | 1.23467 |

**Table 21: Final Buddy Group Size for Group Size = 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.024 | 1.10133 | 1.13867 | 1.328 |
| **150 Ticks** | 1.056 | 1.05867 | 1.12267 | 1.24267 |
| **300 Ticks** | 1.02667 | 1.09067 | 1.21067 | 1.24 |
| **500 Ticks** | 1.03467 | 1.10133 | 1.22933 | 1.272 |

**Table 22: Final Buddy Group Size for Group Size = 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.15467 | 1.12533 | 1.25067 | 1.416 |
| **150 Ticks** | 1.04 | 1.184 | 1.344 | 1.436 |
| **300 Ticks** | 1.05067 | 1.2 | 1.256 | 1.32533 |
| **500 Ticks** | 1.05067 | 1.19733 | 1.28 | 1.40667 |

**Table 23: Final Buddy Group Size for Group Size = 15**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.18133 | 1.17867 | 1.26933 | 1.49067 |
| **150 Ticks** | 1.056 | 1.144 | 1.33067 | 1.59467 |
| **300 Ticks** | 1.088 | 1.19467 | 1.28533 | 1.57467 |
| **500 Ticks** | 1.08267 | 1.24667 | 1.26933 | 1.62133 |

From these results, several key observations jump out at us:

- Large buddy group relations are not being formed, regardless of environmental parameters
- Increasing the number of activities generally increases the size of buddy groups
- Increasing the duration of activities results in inconsistent buddy group size changes
- Increasing group sizes provides slight but disproportionate increases in buddy group size

The most glaring observation from these results is the fact that buddy groups do not appear to be forming. Average buddy group sizes never even get very close to 2.0, which would be the case if at least everyone were paired up with another student. This indicates that there are many students without any buddies (except themselves). Combing through the data logs proves that this is the case. Interestingly enough, we did observe some groups of size larger than 2 (up to 5), but these anomalies were very rare. This lack of group formation indicates a serious problem with our group formation design since we are not achieving emergent behavior from local decisions. A detailed description of this problem is provided in the Discussion section of this report.

From the results we did observe, we can note that as the number of activities increases, so to does the average size of buddy groups (to at least a little extent), as should be expected. Like we indicated in the previous results, as students work together in more activities, they are both exposed to more peers and have more opportunities for relationship development, resulting in (at least slightly) larger buddy groups.

Unexpected was the trend that increasing the duration of activities did not result in larger groups. We believed that when agents work together for longer, they would learn more about one another and be more willing to form buddy relationships, but this was not the case. However, this surprising result could be caused by the general lack of group formation, or even with more knowledge, agents were not willing to branch out and form relationships.

Finally, it appears that increasing group sizes (i.e., the number of students work with which consist first of buddies then randomly assigned students) increases the average buddy size at the end of simulations, as we expected. Once again, this is because agents sample more students and find better matches for their students. Comparing to our last observation, this implies that knowing about many students is more important to group formation than knowing a lot about students. Thus, quantity of information is more important for building larger relationships than quality.

**Final Buddy Group Age:** Similar to observing the average size of buddy groups at the end of simulations, we also wanted to record and analyze the average length of buddy relationships (as measured in the number of activities spent in a relationship). These data are presented in Tables 24 – 27.

**Table 24: Final Buddy Group Age for Group Size = 3**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.6 | 3.87724 | 8.15860 | 18.79594 |
| **150 Ticks** | 0.2 | 3.63968 | 10.85108 | 27.86915 |
| **300 Ticks** | 0.46667 | 5.11667 | 15.13333 | 24.71179 |
| **500 Ticks** | 0.8 | 5.86667 | 16.81887 | 28.44215 |

**Table 25: Final Buddy Group Age for Group Size = 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 2.1 | 3.83828 | 7.10031 | 15.96682 |
| **150 Ticks** | 1.99 | 3.76799 | 11.41115 | 21.20492 |
| **300 Ticks** | 2.26667 | 6.57667 | 15.35176 | 25.20002 |
| **500 Ticks** | 1.83333 | 7.05333 | 9.74530 | 31.98362 |

**Table 26: Final Buddy Group Age for Group Size = 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|

| | | | | |
|---|---|---|---|---|
| **50 Ticks** | 1.88519 | 3.65911 | 7.14123 | 14.44613 |
| **150 Ticks** | 1.69319 | 4.85973 | 13.63242 | 21.78567 |
| **300 Ticks** | 2.32333 | 5.21047 | 12.76406 | 23.73653 |
| **500 Ticks** | 2.73333 | 7.40010 | 13.97395 | 23.31784 |

**Table 27: Final Buddy Group Age for Group Size = 15**

| | **10 Act.** | **25 Act.** | **50 Act.** | **100 Act.** |
|---|---|---|---|---|
| **50 Ticks** | 1.91490 | 4.17541 | 6.80061 | 16.44740 |
| **150 Ticks** | 1.66655 | 3.85950 | 10.57924 | 29.30073 |
| **300 Ticks** | 2.03105 | 5.43281 | 11.88034 | 35.96596 |
| **500 Ticks** | 2.47818 | 6.00278 | 13.59128 | 32.22764 |

These results indicate several important points:

- Although buddy group sizes don't show emergent behavior, lasting relationships are formed.
- Increases in the number of activities and duration of activities both increase the average age of buddy relationships.
- Increasing the size of groups does not result in a general increase in relationship age.

After viewing these results, it was refreshing to observe that at least some emergent behavior from buddy group formation occurred. Even if relationships were not being formed often, at least they lasted when they existed.

Since consistent relationships were formed, their age increased when the number of activities increased, as should be expected. This makes sense because relationship age is measured in terms of the number of activities, and any increase in this number should increase relationship age, if relationships are lasting. However, we did not expect the average age of relationships to increase with the duration of activities. Given the lack of additional group formation caused by increasing activities and the evident lack of impact of additional knowledge about peers, we thought that, despite our hypothesis to the contrary, duration of activities would have either no effect or a negative effect on relationship age. However, the observed tread implies that as agents learn more about one another, they are not more confident in forming new groups, but the extra knowledge does increase their confidence in the relationships they have already have formed. Thus, acquiring more accurate information about others was beneficial to the agents.

Continuing the inverse relationship with buddy group size, increasing the group size did not increase buddy group ages. Since relationships were few and far between, but lasting when they occurred, agents with buddies were not willing to form new relationships no matter how many other agents they worked with, but instead stuck it out with their current buddies. Thus, for buddy group duration, quality of information was more important than quantity.

**Maximum Buddy Group Size:** After observing the lack of large buddy groups at the beginning and end of experiments, we wanted to verify that large groups were not formed and then quickly dropped a few activities later (even though evidence of lasting buddy groups would argue otherwise). Thus, we also recorded the maximum buddy group size observed for each simulation. These results are presented in Tables 28 – 31.

**Table 28: Max Buddy Group Size for Group Size = 3**

| | **10 Act.** | **25 Act.** | **50 Act.** | **100 Act.** |
|---|---|---|---|---|
| **50 Ticks** | 1.00533 | 1.096 | 1.15467 | 1.25333 |
| **150 Ticks** | 1.00267 | 1.05067 | 1.12667 | 1.168 |
| **300 Ticks** | 1.016 | 1.064 | 1.136 | 1.17067 |
| **500 Ticks** | 1.008 | 1.056 | 1.09333 | 1.25333 |

**Table 29: Max Buddy Group Size for Group Size = 5**

| | **10 Act.** | **25 Act.** | **50 Act.** | **100 Act.** |
|---|---|---|---|---|
| **50 Ticks** | 1.024 | 1.144 | 1.18667 | 1.37067 |
| **150 Ticks** | 1.06133 | 1.06933 | 1.16 | 1.26933 |
| **300 Ticks** | 1.02933 | 1.10133 | 1.21867 | 1.26933 |
| **500 Ticks** | 1.03467 | 1.10933 | 1.23467 | 1.33333 |

**Table 30: Max Buddy Group Size for Group Size = 10**

| | **10 Act.** | **25 Act.** | **50 Act.** | **100 Act.** |
|---|---|---|---|---|
| **50 Ticks** | 1.176 | 1.25067 | 1.27733 | 1.46667 |
| **150 Ticks** | 1.06667 | 1.19467 | 1.34667 | 1.45733 |
| **300 Ticks** | 1.06667 | 1.2 | 1.272 | 1.36533 |
| **500 Ticks** | 1.05667 | 1.19733 | 1.296 | 1.45467 |

**Table 31: Max Buddy Group Size for Group Size = 15**

| | **10 Act.** | **25 Act.** | **50 Act.** | **100 Act.** |
|---|---|---|---|---|
| **50 Ticks** | 1.22133 | 1.30933 | 1.30667 | 1.496 |
| **150 Ticks** | 1.088 | 1.17067 | 1.344 | 1.65067 |
| **300 Ticks** | 1.096 | 1.208 | 1.304 | 1.59067 |
| **500 Ticks** | 1.09333 | 1.252 | 1.29333 | 1.66133 |

These results do not offer us any key trends not already observed by in final buddy group sizes, except that buddy groups truly never do grow much in size. In fact, it appears that size tends to peak sometime before the end of the simulation (given the differences between the final and maximum values), but these differences are small.

**Maximum Buddy Group Age:** To complete our analysis of buddy group relationships, we decided to also record the maximum age of buddy relationships to see if these numbers differ much from the final values. These results are presented in Tables 32 – 35.

**Table 32: Max Buddy Group Age for Group Size = 3**

| | **10 Act.** | **25 Act.** | **50 Act.** | **100 Act.** |
|---|---|---|---|---|
| **50 Ticks** | 0.6 | 4.55860 | 8.37144 | 18.87427 |
| **150 Ticks** | 0.2 | 3.80833 | 14.11506 | 27.86915 |
| **300 Ticks** | 0.46667 | 5.11667 | 16.78154 | 24.91374 |
| **500 Ticks** | 0.8 | 6.15 | 17.10221 | 28.50981 |

**Table 33: Max Buddy Group Age for Group Size = 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 2.1 | 4.02817 | 7.10031 | 15.96682 |
| **150 Ticks** | 2.36 | 4.00410 | 11.44146 | 21.22468 |
| **300 Ticks** | 2.53333 | 6.65333 | 15.42235 | 26.17531 |
| **500 Ticks** | 1.83333 | 7.45333 | 11.18364 | 32.09088 |

**Table 34: Max Buddy Group Age for Group Size = 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.92963 | 3.66936 | 7.25330 | 14.44613 |
| **150 Ticks** | 1.71302 | 4.86180 | 13.78417 | 21.78566 |
| **300 Ticks** | 2.45 | 5.26714 | 12.76557 | 23.75364 |
| **500 Ticks** | 2.73333 | 7.53083 | 13.97829 | 23.46713 |

**Table 35: Max Buddy Group Age for Group Size = 15**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.98037 | 4.17857 | 6.80061 | 16.44740 |
| **150 Ticks** | 1.77424 | 3.86860 | 10.74251 | 29.30073 |
| **300 Ticks** | 2.09811 | 5.43281 | 11.89149 | 35.96596 |
| **500 Ticks** | 2.47818 | 6.03981 | 13.88003 | 32.47518 |

As with maximum buddy group size, these results do not provide any observable trends not already discussed. Once again, the maximum values are very close to the final values, but are slightly higher in many cases, indicating that groups peak slightly before the end of simulations, then are broken. Since this has nothing to do at all with the fact that simulations are ending (since agents are not aware of this fact), this implies that some variations do occur over time in buddy group relationships, but not to a large extent. Given longer simulations, more groups might be formed or broken, but not many changes will be made.

## Local Decisions vs. Global Control

The final set of results we collected data for were to observe the levels of local decisions versus global control in our simulations. In our setup, global control is exercised by a teacher when she assigns random students to groups (which happens often due to our lack of buddy relationship formation) or starts and stops activities. In order to determine the amount of global control occurring in our simulations, we recorded the proportion of messages sent out by teachers against the total number of messages sent during a simulation. The results are displayed in Tables 36 – 39.

**Table 36: Proportion of Teacher Messages for Group Size = 3**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 2.677% | 2.196% | 1.961% | 1.690% |
| **150 Ticks** | 0.850% | 0.706% | 0.645% | 0.593% |
| **300 Ticks** | 0.391% | 0.335% | 0.311% | 0.296% |
| **500 Ticks** | 0.214% | 0.193% | 0.183% | 0.181% |

**Table 37: Proportion of Teacher Messages for Group Size = 5**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 1.573% | 1.228% | 1.121% | 1.010% |
| **150 Ticks** | 0.485% | 0.404% | 0.380% | 0.368% |
| **300 Ticks** | 0.222% | 0.195% | 0.186% | 0.179% |
| **500 Ticks** | 0.126% | 0.114% | 0.110% | 0.107% |

**Table 38: Proportion of Teacher Messages for Group Size = 10**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.759% | 0.602% | 0.557% | 0.585% |
| **150 Ticks** | 0.236% | 0.204% | 0.198% | 0.208% |
| **300 Ticks** | 0.111% | 0.101% | 0.099% | 0.0947% |
| **500 Ticks** | 0.064% | 0.060% | 0.058% | 0.059% |

**Table 39: Proportion of Teacher Messages for Group Size = 15**

| | 10 Act. | 25 Act. | 50 Act. | 100 Act. |
|---|---|---|---|---|
| **50 Ticks** | 0.495% | 0.404% | 0.386% | 0.387% |
| **150 Ticks** | 0.159% | 0.145% | 0.139% | 0.132% |
| **300 Ticks** | 0.077% | 0.071% | 0.069% | 0.077% |
| **500 Ticks** | 0.045% | 0.0044% | 0.042% | 0.044% |

These results indicate several important trends:

- Increasing every parameter generally decreases the proportion of messages sent by teachers.
- Global control by teachers is a very limited portion of our simulations.

First of all, as each of the parameters increases, the proportion of messages sent by teachers decreases. This makes sense because as the number of activities increase, both the number of messages sent while learning and group formation increase. Similarly, as activity duration increases, so does the number of messages within groups sent to support student learning. Increasing group sizes will also increase the amount of traffic between agents within groups. Each of these increases in other types of traffic are greater than any increases in the number of teacher messages (which are only affected by the number of activities), so the total proportion of messages sent by teachers should decrease.

Because the percentage of messages sent by teachers decrease, the amount of global control decreases as well, increasing the local control by each agent. However, some of this control is sacrificed to groups, where decisions like group formation are a little less local (i.e., they depend on more than one agent for completion), but even these intragroup decisions are still based primarily on local decisions made by each agent (e.g., what type of learning action to perform, voting yes or not to a proposition, etc.).

## Discussion

In this section we provide some discussion based on the results presented in the previous section. Of special highlight is a discussion on the emergent behavior observed within our system, as well as a discussion describing the problems encountered in buddy group formation. We also provide a short discussion on applications from our simulations to the actual classroom domain.

### Emergent Behavior

Based on the results we have collected and the important trends we have observed, we are now able to evaluate our hypotheses in terms of emergent behavior.

**Knowledge Convergence:** First, we evaluate affects of different environmental parameters on knowledge convergence. It is important to note that we did see emergent behavior with respect to knowledge convergence – average knowledge levels did increase faster at first, the converge to large values and the rate of increases in knowledge slowed down.

Our first hypothesis stated that knowledge convergence would speed up as group size increased, but after a certain size, the speed to convergence would actually slow down. From our results, especially focusing on the average knowledge after 500 learning ticks and at the end of simulations, we can see that as group size increased, knowledge levels were higher, holding all other parameters constant. Assuming that for every simulation the average knowledge level started around the expected 0.5 value, students were converging faster to higher levels in the same number of learning ticks. However, when comparing the results of our simulations with group sizes of 10 and 15, we can see that there was not much of a difference in the knowledge levels (holding other parameters constant), indicating that the advantage of larger groups was pretty constant for these two sizes. Thus, convergence did not occur slower for very large groups, but it wasn't much faster. It is still possible that even larger groups would have experienced a slowdown (given the decreasing rate of increase in knowledge while increasing groups), but our data does not support this claim.

Our third hypothesis claimed that the number of activities would not affect the rate of convergence, but convergence would not occur if there were not enough learning activities to promote growth. We also posited that activity duration would be proportional to knowledge convergence because more time would be spent proportionally on learning activities. Our findings verify some of these claims. First of all, judging from the average knowledge levels after 500 ticks, the number of activities appeared to play a role in convergence rate, but our analysis has lead us to believe that this was due to random seeds instead of the activities themselves. We also observe that convergence does not occur if not enough activities are performed, as evidenced by the jump in final knowledge levels between 10 and 25 activities. With respect to duration of activities, we observed that the overall convergence rate was actually faster for smaller durations than longer durations because all domains are more likely to be selected at least once in a fixed amount of time when activities are shorter. The first few activities in a domain produce the fastest convergence, so balancing the amount of time between domains results in faster convergence, as observed based on our average knowledge levels after 500 learning ticks. However, the ultimate convergence height was proportional to activity duration, as we expected.

**Buddy Group Formation:** The second type of emergent behavior we wished to achieve was large, lasting buddy relationships amongst students. However, we were not as successful in achieving this behavior.

Our second hypothesis posited that as group size increased, the duration of buddy relationships would increase, but after a certain size, it would become harder to form coherent groups. Based on our data, this hypothesis was incorrect. First of all, we found that increasing group size didn't have much of an effect on buddy relationship age, but it did result in slightly larger group sizes. This is due to the fact that increasing group sizes exposed agents to more students, which made it easier to form a relationship, but it played no role in the duration of buddy groups since relationships were very likely to last a long time if they existed. Also, we did not find that as group sizes increased it became harder to form coherent groups. Our buddy group formation was already very incoherent (as described in the next subsection), so we cannot prove this portion of the hypothesis to be true.

Our final hypothesis stated that as the number of activities and duration of activities increased, so too would the length of buddy relationships. Both of these claims were verified by our data. The increase in relationship age due to the number of activities is pretty straightforward since age is measured in terms of the number of activities students are together. Because relationships lasted long when they occurred, more activities meant students could be buddies for longer. Also, we believe the duration of activities increased the average age of relationships due to increased confidence in buddy selection caused by more accurate knowledge gained about peers over time.

### Buddy Group Formation Problems

As stated previously, buddy group formation did not show a high level of convergence or emergent behavior. We believe that this is likely caused by three factors: 1) a lack of robustness in the event of group formation failure, 2) synchronization issues, and 3) static ordering of agents.

First of all, in our current student agent design, each agent sends Express Interest messages only once each buddy group formation period. For each buddy group formation period, if an agent's current prospect collapses, it doesn't have an opportunity to pursue another. If forming groups is more likely to fail than succeed, it will not be very frequent that groups are formed without needing repeated attempts per formation period.

The lack of an iterative buddy formation period leads to the second factor disallowing buddy group formation. Because buddy group formation is run only once, the steps in the process are essentially synchronized. Nearly all agents reach the confirmation step in mergers at the same time. This causes conflicts as only one confirmed buddy group can be handled at a time, causing most mergers to be cancelled. Currently, agents might agree to multiple mergers at once and only the first one accepted is formed. When multiple agents in a group are trying to merge with other groups, only a lucky few buddy groups will be formed. In most cases, one or both of the parties involved will have to bail out due to actions taken by buddies. This problem arises from the lack of central control within a group. Because we wanted to stress the individual decisions of each agent, each agent in a group can do whatever it thinks is in its own best interests, leading to a lack of coherent behavior by the group. If instead the group could elect a single leader to handle all activities with other groups, it is likely that less formation attempts would fail because less would be accepted then canceled due to other deals. In fact, with central control, we could limit each group to only one merger attempt at a time, preventing the need for back outs.

Finally, the order in which agents are run in each step is critical to formation as well. Within Repast [1], agents execute their actions sequentially, so agents at the head of the agent list always perform their actions before later agents. Because only the first agent to send an Inform Message has its merger accepted by the group, there exists in the current system a few privileged agents that receive their preferred mergers at a higher rate than others. This may reduce the number of buddy groups formed as only the preferred agents' confirmation steps are processed. Once again, if these preferred agents' actions across groups do not match up, it is very likely that merger attempts will fail due to other deals being wrongly accepted first and a lack of coherence within and between groups.

## Real World Application

Assuming that our experiments truly simulation classroom behavior (which is certainly unproven but worth considering), we can draw several conclusions from our simulations that could aid in designing real online collaboration environments.

First of all, it appears that increasing group sizes for collaborative activities provides some benefits to students, but in a real world scenario, finding roles for all students in large groups could be difficult and cause some students to simply "piggyback" off of the work of others. Additionally, increasing the number of activities and duration of activities does allow students to learn more, but the added time cost will no longer be worth it after some point. Since we have diminishing returns on knowledge gains, the students' time could be better spent on other non-learning activities, or learning in other domains. If our knowledge levels approximate the average grade of students, we should probably stop after the average is near 75% (a C grade). Thus, stopping could happen pretty early, but giving more/longer activities would cause a continued increase in learning. However, since many students max out their knowledge, they have nothing to gain from sharing and essentially become forced tutors for less intelligent students. This results in the small knowledge gains after convergence that we observed. This is very important to understand when designing an online collaboration environment for students and is an issue that must be addressed. If such a scenario does occur, some reward aside from knowledge gains should be provided to intelligent students to enlist their help. An alternative would be to allow these students to work on something else and group the slower students together, but we do not believe that this will result in the same gains for the remaining students.

Second, based on our difficulties with buddy group formation between agents, we believe that similar difficulties would arise in a real classroom environment without some sort of group control exerted by only a few students. If all members were allowed to go and try to increase their central set of core buddies, conflicts would arise in their choices, leading to difficulties in a group, and possibly even relationships being broken. However, if only a few students in each group tried to merge with others based on the inputs of all group members, more coherent behavior would probably occur, as we expect would occur if we changed our agent design.

# Future Work

In this section, we focus on some ideas we currently have for future work in order to further evaluate the outcome of simulations involving collaborative learning amongst students in online learning environments supported by intelligent agents. These include comparing emergent behavior across different numbers of students, testing additional hypotheses evaluating the impact of student attributes on learning, performing more statistical analyses, and finally improving our buddy group formation procedure.

## Number of Students

Within our simulation setup, one of the adjustable environmental parameters includes the number of students in the classroom. However, we did not create any questions or hypotheses useful for determining how the number of students in a learning environment affects the emergent behavior of the system. While running our experiments, we did collect data for various numbers of students (as outlined in the Experimental Setup section of this report), so we have all of the data necessary for determining the affect of this parameter on emergent behavior. Before writing this report, we thought about adding a couple hypotheses relating the number of students to learning and group formation, but we realized this would add a fourth dimension to our results, which would be very difficult to present and analyze for this report. Instead, we fixed the number of students to 150 and evaluated our results on this subset of our data. We would like to analyze the importance of the number of students, as well as analyze the impact of the other parameters for a different fixed number of students, but we save such analysis for future work.

## Student Attributes

When we originally designed our simulations and experiments, we also had four other hypotheses we wished to evaluate, each of which was created to determine the importance of student attributes on student learning and group formation. These hypotheses and related questions include:

***Question 5*****:** How does group similarity (i.e., variance of student attributes in groups) affect convergence?

***Hypothesis 5*****:** We hypothesize that groups that contain very similar students (i.e., low variance) will exhibit higher levels of learning and thus achieve convergence faster. This should be true of both working groups (i.e., all students in a group) and buddy groups

***Question 6*****:** How does group similarity affect relationship duration?

***Hypothesis 6*****:** We hypothesize students with similar attributes will learn more from each other, thus encouraging future interactions. Thus, similar students will be part of lasting groups and have longer relationships, while students with widely different attributes will not remain in groups for long.

***Question 7*****:** What attribute combinations within buddy groups will provide the fastest convergence of knowledge levels?

***Hypothesis 7*****:** We hypothesize that students with high communication abilities in the same modes will result in the fastest convergence on higher knowledge levels, regardless of disparities in initial knowledge. If we assume collaborating causes greater knowledge gains and that resourceful students can work well alone, individuals with only a moderate resourcefulness may show the highest knowledge gain as their agents will push them to collaborate rather than work alone. Similarly, more motivated students will be more likely to learn quickly than unmotivated students.

***Question 8*****:** What attribute combinations cause the longest and shortest relationships?

***Hypothesis 8*****:** Student agents will retain relationships that are beneficial to the student. We believe that communication abilities are the most important attribute to forming beneficial relationships as they allow knowledge to be shared with other students. We hypothesize that the longest relationships will be between agents with high communication skills, and conversely the shortest relationships will result from the lack of communication between students due to poor communication skills.

In the future, we wish to evaluate each of these hypotheses given our current simulation setup. We already track all of the necessary information to perform the required evaluations, and in fact we have already collected the necessary data. However, due to both time constraints and the complexity of presenting information about attribute combinations while varying multiple environmental parameters, which would have resulted in many more tables even more complex than those already presented in our Results section, we have left out these hypotheses for this report.

## Statistical Analyses

Many of the analyses performed in our results section were rather ad-hoc and uninspired by actual statistical tests. In the future, we will work on using more statistical tools to produce better, more accurate results. These tools include correlation studies to determine exactly how different effects are related to the different parameters. We would also increase the number of runs of each simulation from 5 to 30 to reduce variance even further (which seems necessary due to the odd effects of random seeds on initial knowledge levels in the 25 activity experiments) and take advantage of the benefits of the Central Limit Theorem.

## Buddy Group Formation

As mentioned in the Results and Discussion sections, the small size and rarity of buddy group formation indicates a low level of emergent behavior. The three causes for this were identified as only allowing interest messages to be sent once, problems with synchronization and coherence across all agents within a group, and the preferential treatment of agents early on the processing queue. In future implementations of this system, these factors could be overcome fairly easily. The message driven buddy management system was designed for aynchronous ad-hoc buddy formation, and because of the reasons mentioned above, it is expected that the buddy group emergent behavior hoped for would be prevalent in a buddy formation step of this type. If, instead of sending Express Interest messages only once at the beginning of group formation, they were sent at multiple times corresponding to when an agent does not have a current buddy merger being negotiated, it would allow every agent to attempt several joins if initial buddy mergers are unsuccessful. This could also presumably solve the synchronization issue as after a few attempts at merging every agent would be at a different stage in group formation, allowing more mergers to complete successfully instead of everyone trying to perform the same step at the same time, essentially blocking one another. Finally, if the execution order of the agents was selected at random for each tick, it would guarantee that no agent would receive preferential treatment, allowing many agents to get their Confirm Merger messages accepted by their buddy group.

If we were able to relax the constraint that agents must make local decisions, we could also try to add some global group control to the formation procedure, adding coherence and synchronization to a highly distributed problem.

## Conclusions

In conclusion, we have created a simulation for testing collaborative learning amongst students in an online educational environment supported by intelligent agents using the Repast software environment [1]. During learning activities, each agent attempts to maximize its student's learning by selecting an appropriate learning action, and after activities, agents attempt to form buddy relationships to continue working with peers who best fit their students. Based on our agent design, we attempt to achieve two forms of emergent behavior from local decisions made by the agents: 1) converging knowledge levels, and 2) full, lasting buddy relationships. The first behavior was achieved successfully, but we were unable to form large buddy relationships, although those relationships that formed were long lasting. We believe that our failure was due to problems inherent in our buddy formation system which allowed agents too much freedom to make their own decisions and interact with other groups, leading to synchronization problems within groups. We also limited the number of invitations an agent could extend, which limited the effectiveness of buddy group formation due to the high failure rate of attempts to create relationships.

## References

[1] Argonne National Laboratory, "Repast: Recursive Porous Agent Simulation Toolkit", available at <http://repast.sourceforge.net/repast_3/index.html>