Active and Group Learning in the Computer Architecture Classroom

Mark Fienup
Computer Science Dept.
University of Northern Iowa
Cedar Falls, IA 50614-0507
e-mail: fienup@cs.uni.edu

Abstract

Lecture is the traditional instructional format in Computer Science (and in most disciplines) at the college level. However, research studies have shown what we all suspect -- lecture is not very effective. Students retain only a small fraction of the material covered, attendance only has a marginal effect on performance, and learning via lecture is independent of the lecturer's quality.

This paper presents an alternative to the lecture format, i.e., active and group learning. Application of active and group learning to teaching Computer Architecture within a Computer Science program is described.

Introduction

Learning is not a spectator sport -- so why is the majority of classroom time spent lecturing? Students need to be more engaged than just listening and note taking! Studies have shown that after 10-15 minutes of lecturing students essentially stop learning [6], but their "attention-span" clock is reset by interjecting activities to break up the lecture. Additionally, we (as teachers) have all discovered that we learn a lot about a topic when trying to teaching it, so why can't our students learn a concept at a deeper level by trying to explain it to other students? Issues such as these motivated my interest in active and group learning techniques.

In the past my own teaching experiences tend to support that non-lecture activities in the classroom are very helpful. For example, when teaching about binary and hexadecimal arithmetic in a Computer Organization course, a useful exercise is to show several examples and then immediately in the classroom let students work several problems. Students immediately see if they really understand the concept. However, these active learning activities were few and far between with lecture being my predominate teaching technique. After a workshop by Jeffery J. McConnell (Canisius College) on active and group learning at the SIGCSE 1999 conference, I decided to take the plunge and apply active and group learning concepts the following semester in my Computer Architecture class on a daily basis. This paper discusses my experience of applying active and group learning concepts in this course.

Preparation - Group Selection

I wanted to apply active and group learning techniques from the first day of the course, so students knew what to expect. This necessitated me to decide on groups before the first class. It had been suggested in a workshop I attended on active and group learning that groups of size 5 seemed to work best. Since 37 students were enrolled in the course, five groups of 5 and two group of 6 seemed reasonable.

I wanted to select groups with roughly equal "potential," but I knew very few of the students before the semester started. The method that I used was to sort students by cumulative GPA and then assign the "top" student to group 1, the next "best" student to group 2, etc. until 7, then I assigned the next student to group 7 down to group 1, etc. In

other words I assigned the sorted list of students to groups 1, 2, 3, ..., 7, 7, 6, ..., 1, 1, 2, ..., 7, etc. For this class, this technique proved successful for all but one group. To stimulate discussion is each group it is useful to have at least one leader or extroverted individual per group. Since high GPA does not necessarily correspond to these traits, you run the risk of forming a group of all introverted students. This is what appeared to be the problem with one of the initial groups.

Ideally, you would know your students well enough when forming groups to ensure at least one "leader" or extroverted personalities per group. As a result I reorganized the groups after the first test which occurred about a third of the way through the semester. By then, I knew the students well enough to make sure that I had several extroverted personalities per group while still maintaining roughly equal ability based on their current standing in the course.

General Instructional Format

The format chosen for the class involved assigning students to groups of approximately 5. Before each 75 minute class period, students arranged the tables so each group could sit together. Class time alternated between me performing mini-lectures (predominately using overhead transparencies of figures from the text or of my own design) on a topic, each group discussing and answering one or more exercises related to the topic just covered, and then a class discussion of group answers. While answering the exercise questions, each group recorded their answer before the class discussion of group answers. At the end of class each group handed in a single copy of their answers. During the class discussion of the group answers, groups were not allowed to modify their individual group's answers that were to be turned in. In grading the course ten percent of a student's grade was based on the quality of their group's answers to these in-class exercises.

Students were expected before each class to read assigned sections of the text in preparation for the in-class exercises. All of the overheads for the course and the in-class questions were made available via the web after each class period.

It should be pointed out that this format represents only my implementation of active and group learning. Active and group learning techniques must be adapted for each instructor's style, i.e., what works for me might not work for you. Frederick [2] and Silbererman [4] are two good sources of ideas.

Examples of Typical Class Periods

The in-class questions are designed to stimulate student discussion, and improve learning about the current topic being discussed. Figure 1 is a sample set of questions from a class period when the class was just finishing talking about system buses and starting to discuss main memory. In the previous class we had discussed the PCI bus, including a brief looking at a timing diagram for a read operation from Stallings [5]. At the end of the previous class, I asked the students to develop the write-operation timing diagram as homework. Additionally, they knew that they should read the chapter on main memory

On this particular day, group-work on question 1 (Figure 1) started the day. From e-mail and questions received before class, I knew that many students were having difficulty understanding the PCI timing diagrams. The purpose of this question was for students to compare and discuss their answers to the PCI write operation assigned at the end of the previous class. After group discussion, most groups had a correct timing diagram, and after class discussion of group answers, I'm certain (although I have no proof) students understood the PCI timing diagrams much better than if I would have spent a significant part of a class period lecturing to them.

Next, I gave a quick mini-lecture review of how we implemented a register file when we were talking

1. Combine your group's answers for the PCI write operation, i.e., draw and explain a timing diagram for a PCI write operation (similar to Figure 3.22). Clearly indicate: a) the address phase, data phase(s) and any wait states, and b) when the "target" reads the data off the bus

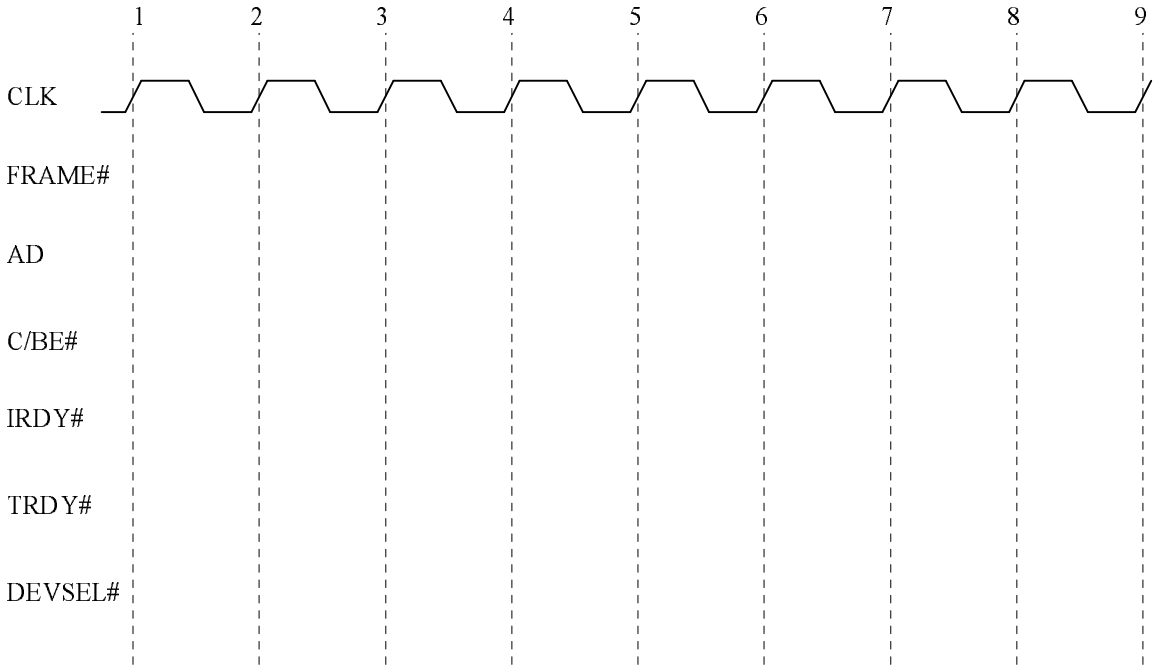

2. Suppose that we are implementing a 4 M x 4 memory similar to the register file organization with one write-port and one read-port.
a) How many and what type of decoder(s) would be needed? *<space for answers removed>*

b) How many total gates (assume only 2-input gates are used) would be needed to implement this (these) decoder(s)?

c) How many and what type of MUX(s) would be needed?

d) How many total gates (assume only 2-input gates are used) would be needed to implement this (these) MUX(s)?

3. Consider the 4 M x 4 memory implementation shown in Figure 4.4 *<2D array of bits>*.
a) How many total gates (assume only 2-input gates are used) would be needed to implement the decoders?

b) How many total gates (assuming only 2-input gates are used) would be needed to implement the MUX?

c) How do these values compare to the answers from question 2 above?

Figure 1: Sample in-class questions for a lecture

about flip-flops previously in the semester. In the past, students always had difficulties understanding the subtleties of implementing main memory chips as a two-dimensional array of bits. Question 2 was designed for students to discover the unfeasibility of implementing main memory chips similar to a register file. After group discussion of question 2, a mini-lecture about implementing memory chips as a two-dimensional array of bits made a lot more sense to the students. Finally, the 75-minute class period ended with students answering question 3 which was intended to "drive home" the need to implement memory chips as a two-dimensional array of bits.

The second example of a typical class period occurred at the start of discussion about hardware support for the operating system. Previously in the semester we had discussed what an interrupt was in the context of external devices, and students were to have read the chapter on operating system support in the text. Figure 2 is the set of in-class discussion questions from that period.

The class period started with a mini-lecture overviewing the goals of an operating system followed by group discussion of question 1. This question and follow-up discussion was designed to make the students realize that if a user program is running, then the operating system cannot also be running to monitor it. Question 2 was immediately discussed by the groups without any lecture since all of the students had a prerequisite Computer Organization course that focused on assembly language programming. I was surprised by their lack of understanding of the concepts of privileged instructions and dual-mode execution of the processor. One of the advantages of this classroom format is that you can determine where students are having conceptual problems.

Next, I did a mini-lecture that addressed the concept of a process, and process scheduling in a multiprogrammed system. Question 3 was designed to allow students to see the advantages of multiprogramming. After group discussion, most groups were able to correctly explain the improvement of CPU utilization.

1) How does OS/hardware protect against a user program that going into an infinite loop?

2) How does the OS/hardware prevent a user program from accessing files of other user?

3) Explain the near linear relationship between CPU utilization and degree of multiprogramming observed in the below graph.

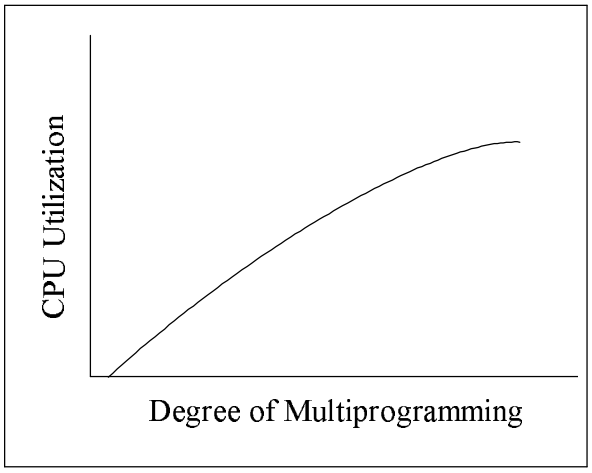

4) How does the OS/hardware prevent a user program from accessing (RAM) memory of other user programs or the OS?

Figure 2. Second example of in-class questions.

The groups immediately trying to answer question 4. I used the follow-up discussion to motivate the topic of virtual memory for the next class period.

Student Comments about the Course Format

Toward the end of the semester, students were asked to provide anonymous, written comments about the course format. The vast majority of comments were positive with a few negative comments and some suggestions for improvements.

One of the goals of using active and group learning is to try to keep the students alert and learning throughout the whole class period. Some supporting students comments were:

- "I liked the group activities. … I pay closer attention when there is questions than when there is just lecture."
- "Not lecturing the whole time with the problems are really nice to have in that we can apply things right away and get an answer swiftly when the lecture resumes. Plus this keeps the student interested and also awake during class."

Another goal of active and group learning was to improve the amount and quality of learning that occurred in the classroom. Supporting students comments were:

- "I really appreciate the effort being made to help us understand the material, rather than just memorize the material. I know that I have learned and understand, and will remember things because I relate problems to the things we do in class. I learn more in class than I do from the book. In almost all other classes you learn more from the book, and class reinforces what the book says. In this class the materials reinforce what the book says."
- "I think the group structure worked well. The questions were geared along with what you were talking about so that helped. I think that when you explain the topic and then provide an exercise for us to do relating to it, it reinforces the idea in our head."
- "I think group work and discussion questions after short periods of lecture worked very well. By analyzing the concepts after they were presented to the class, the concepts became more clear and easier to remember. Studying for tests was easier because of this."
- "The groups worked really well. I seemed to remember more examples done in group then if it had just been lecture. Sometimes in lecture I think I understand something until I go to do it then figure out I really didn't understand. In the groups we get instant feedback."

Finally, peer instruction was also achieved using this technique as noted in the following student comments:

- "First, I really enjoy having the group sessions. First, it breaks up the class a little bit. Secondly, employers always ask what type of work we've done in teams, this is one of the few examples we can give. Third, sometimes a team member can explain something in a different way that is more understandable."
- "The group idea does work well, it allows us to bounce ideas off of each other and combine our thoughts."

Only three negative comments were received, but they were:

- "Sometimes some of our group members generally don't read the book and they talked something that are not related to the question."
- "This seems good to have groups discussion the topic in class. Because it creates more active environment....Personally I like formal lecture classes, because it better for me."
- "I think that the group idea is good. But I don't learn much in the groups. I think that hands on classes would help students out as well."

Interesting enough the first two comments were from foreign students from the Far East (I could tell by the handwriting and sentence structure). So this instructional format was not uniformly liked, especially by the Oriental students.

Conclusions

Originally, I was reluctant to adopt active and group learning techniques because of fears that not enough content could be covered and preparation time would

be too high. Surprisingly, I was able to cover about the same amount of material as previous semester with the same textbook and course outline. This was partly due to the use of overhead transparencies and on-line notes.

Unfortunately, it did require a substantial amount of time to develop in-class activities for students. It typically required one to two hours to develop the questions for each lecture. I tried to develop questions that helped students understanding of concepts that they traditionally have problems with.

Student reaction to the course format was extremely positive. Informally, student performance on homeworks and tests seem to be slightly better, but no quantitative study was performed.

On a personal note, I really enjoyed teaching using the active and group learning format. If you are feeling a little burned out on teaching, I would encourage you to try this approach.

Reference

1. Felder, R. & Brent, R. (1996). Navigating the Bumpy Road to Student-Centered Instruction. *College Teaching*, 44, 43-47.

2. Frederick, P. (1986). The lively lecture - 8 variations. *College Teaching*, 34, 43-50.

3. McConnell, J. (1996). Active learning and its use in Computer Science. *SIGCSE Bulletin*, 28, 52-54.

4. Silberman, M. (1996). Active learning: 101 strategies to teach any subject. Boston: Allyn & Bacon.

5. Stallings, W. (1999). *Computer Organization and Architecture* (5th ed.). Upper Saddle River, NJ: Prentice Hall.

6. Stuart, J. & Rutherford, R. J. (1978, September 2). Medical student concentration during lectures. *The Lancet*, 514-516.