2003-2004 Faculty Learning Community Portfolio

Course: Parallel Computing II

Nan C. Schaller

Table of Contents

Part 1. Philosophy of Teaching, including metaphor for teaching and assumptions about students

I view myself as a guide participating in an interactive dualism: teaching and learning. I believe that all parties must invest themselves to make it work. As such, I set high standards for both my students and myself. Here are the responsibilities I believe that each of us has in this dualism:

My responsibilities:

- To keep current
- To provide a safe environment
- To come to class prepared
- To invite interaction
- To present the material that the student needs to know (using lecture, active learning, etc.)
- To make the material challenging and yet accessible
- To do whatever is needed to stimulate interest in topic
- To encourage mastery the material and how to learn more
- To answer all questions in a timely fashion
- To be flexible
- To provide timely feedback on homework and examinations
- To be available in my office during office hours
- To go the extra mile

Student responsibilities:

- To want to learn about the topic
- To come to class
- To participate in class
- To be prepared to learn
- To pay attention, take notes, ask questions
- To be willing to put the effort into all assignments: readings, homework, projects and to seek help when anything is not thoroughly understood
- To review and correct graded homework, projects, and examinations
- To go the extra mile
- To be prepared to interact, to take a chance, to trust

Part 2. The syllabus of the class, including course objectives, learning outcomes, and policies regarding grades and attendance.

Parallel Computing II

The catalog description of Parallel Computing II says: *Parallel Computing II is a study of the principal trends in parallel algorithm design, through the analysis of algorithms used in various areas of application. Specific techniques that have gained widespread acceptance will be highlighted. The course will investigate the interplay between architecture and algorithmic structure and will discuss the effect that these issues have on the complexity and efficiency of parallel algorithms.*

However, the course has been evolving and will be a collaborative learning effort that will focus on topics decided by the class as a whole. The programming assignments will focus on applying the techniques discussed in Parallel Computing I to current research efforts, some of these interdisciplinary in nature.

Table of Contents

Syllabus -- 20033

Keeping Up (736 only)

Computing Facilities

Assignments -- 20033

- **General Documentation**
  - Sun Online Documentation (Sun HPC ClusterTools 5 Collection)
  - Sun Perfomance Library User's Guide
  - Sun Online Documentation (Forte Developer 7 Collection)
  - file:///usr/local/Forte7/SUNWspro/prod/lib/locale/C/html/index.html - Forte(TM) Developer 7 Documentation Index
    (NOTE: You will need to be on a CS RIT machine to get to this. You may need to use your browser's *Open file...* menu selection)
  - Sun Perfomance Library User's Guide - Forte Developer 7

- Programming Assignment 1 (FAQ's) (/usr/local/pub/ncs/parallel/Anhinga) (gradesheet)

Research Report Session

Supplementary Materials

In the computer science and information technology recommended databases - just a few selected items, i.e., there's much more....

From IEEE Xplore

- Journals
  - IEEE - Computer (Search on parallel)
  - IEEE - Concurrency
  - IEEE - IEEE Parallel & Distributed Technology: Systems & Applications
  - IEEE - Parallel and Distributed Systems, IEEE Transactions on
- Proceedings
  - Advances in Parallel and Distributed Computing, 1997. Proceedings
  - Advances in Parallel and Distributed Systems, 1993., Proceedings of the IEEE Workshop on
  - Algorithms and Architectures for Parallel Processing, 1995 - 2002.
  - Databases in Parallel and Distributed Systems, 1988, 1990.
  - Databases, Parallel Architectures and Their Applications,. PARBASE-90, International Conference on
  - Design and Application of Parallel Digital Processors, 1988, 1991.
  - Fault-Tolerant Parallel and Distributed Systems, 1994., Proceedings of IEEE Workshop on
  - Frontiers of Massively Parallel Computation, 1988 - 1999.
  - High Performance Applications of Parallel Architectures, IEE Colloquium on
  - High-Level Parallel Programming Models and Supportive Environments, 1998. Proceedings. Third International Workshop on
  - Massively Parallel Computing Systems, 1994., Proceedings of the First International Conference on
  - Massively Parallel Processing Using Optical Interconnections, 1994 - 1997.
  - Massively Parallel Processing, 1992, 1998.
  - Massively Parallel Programming Models, 1997. Proceedings. Third Working Conference on
  - Medical Imaging: Transduction and Parallel Processing, IEE Colloquium on
  - Parallel Algorithms/Architecture Synthesis, 1995, 1997.
  - Parallel and Distributed Information Systems, 1991 - 1996.
  - Parallel and Distributed Processing Symposium, 1990 - 2002.
  - Parallel and Distributed Real-Time Systems, 1994 - 1997.
  - Parallel and Distributed Simulation, 1995 - 2002.
  - Parallel and Distributed Systems, 1994 - 2002.
  - Parallel and Large-Data Visualization and Graphics, 2001. Proceedings. IEEE 2001 Symposium on
  - Parallel Architectures and Compilation Techniques, 1996 - 2002.
  - Parallel Architectures for Image Processing Applications, IEE Colloquium on
  - Parallel Architectures for Image Processing, IEE Colloquium on
  - Parallel Architectures, Algorithms and Networks, 1994 - 2002.
  - Parallel Computing in Electrical Engineering, 2000, 2002.
  - Parallel Interconnects, 1999. (PI '99) Proceedings. The 6th International Conference on
  - Parallel Operation of Generating Plant within a Public Electricity Supply Network, IEE Colloquium on

unethical behavior in the execution of assigned work in a computer science course will be treated as follows:

1. For a first offense the student involved will receive a grade of zero on the assignment. [A stronger penalty may be exacted if, in the judgement of the instructor, the offense involves a flagrant violation of basic ethical standards.]
2. For a second offense, in the same or a different course, the student will receive a failing grade for that course.
3. A third offense will be referred to judicial affairs.

Furthermore, the following action will be taken for each person involved in the incident, whether currently enrolled in the course or not:

If the student is a computer science major, a letter recording the incident will be placed in the student's departmental file; otherwise, the letter will be forwarded to the student's department chair or program coordinator.

Nan Schaller - Courses Page

Nan C. Schaller
Rochester Institute of Technology
Computer Science Department
102 Lomb Memorial Dr.
Rochester, NY 14623-5608
telephone: +1.585.475.2139
fax: +1.585.475.7100
e-mail: ncs@cs.rit.edu
20 April 2004
http://www.cs.rit.edu/~ncs/Courses/532.shtml

20033 Syllabus for 532/736

Instructor: **Nan Schaller**
Office: **70-3637**
Office Hours: **4-6 PM on Monday and Wednesday or make an appointment, 475-2139**
E-mail: **ncs@cs.rit.edu**
URL: **http://www.cs.rit.edu/~ncs/**
Course URL: **http://www.cs.rit.edu/~ncs/Courses/532.shtml**

TEXTBOOK: None required.

COURSE OBJECTIVES

The objectives of this course are to:

1. Extend your knowledge of parallel computing.
2. Work with a team to apply your knowledge of parallel computing to real-world interdisciplinary problems.

These objectives support the first four Computer Science Program Objectives:

1. Demonstrate a strong theoretical and practical background across the computer science discipline, with an emphasis on software development.
2. Possess the required practical computing experience to become valueable contributors to the field immediately upon graduation.
3. Demonstrate effective communication and teamwork skills in order to collaborate with colleagues and clients successfully.
4. Exhibit the skills necessary for continuing education and long-term professional growth.

LEARNING OUTCOMES

At the end of this course, students should be able to

1. Research and present a current topic of interest in the field of parallel computing
2. Lead a discussion in a current topic of interest in the field of parallel computing
3. Identify, analyze, evaluate, and implement alternative approaches to parallelizing code for real-world, interdisciplinary problems.

HOW WE WILL DO THIS:

This course will have a computational component. As such, we will be working on real-world, interdisciplinary problems. In addition, given the number of students this quarter, this course will be run as a seminar, i.e., we will take a collaborative learning approach. As such, we will be looking into areas of interest to us as a whole, using research into texts, journal articles and the web as a basis. The topics below will be filled in as they are determined. Each student will be responsible for leading one class and working with one other to generate discussion questions in a timely fashion. (See the Research Report Session below.)

SCHEDULE:

NOTES

a. Maximum grades for work turned in late are 89% for 1 day late, 79% for 2 days late, 69% for 3 days late, etc. Friday and Sunday are considered valid due dates.

b. Please note that if you have a complaint about the grading of any assignment or exam, you must bring it to my attention within one (1) week after I have handed back the graded material in class. After that your coupon will have expired and no grade adjustments will be considered!

c. I use a standard grading scheme, i.e., an average of greater than or equal to 89.5 is required to receive an A for the course.

d. Class participation grading will be somewhat subjective. I will be looking at class attendance, team work, timeliness, and participation in classroom discussions.

C.S. DEPT. POLICY ON ACADEMIC DISHONESTY

Students are expected to maintain the highest standards of ethical behavior.

Those who behave in a dishonest or unethical manner in computer science courses, or in their dealings with the Computer Science Department, are subject to disciplinary action. In particular, dishonest or unethical behavior in the execution of assigned work in a computer science course will be treated as follows:

1. For a first offense the student involved will receive a grade of zero on the assignment. [A stronger penalty may be exacted if, in the judgement of the instructor, the offense involves a flagrant violation of basic ethical standards.]
2. For a second offense, in the same or a different course, the student will receive a failing grade for that course.
3. A third offense will be referred to judicial affairs.

Furthermore, the following action will be taken for each person involved in the incident, whether currently enrolled in the course or not:

If the student is a computer science major, a letter recording the incident will be placed in the student's departmental file; otherwise, the letter will be forwarded to the student's department chair or program coordinator.

Nan C. Schaller
Rochester Institute of Technology
Computer Science Department
102 Lomb Memorial Dr.
Rochester, NY 14623-5608
telephone: +1.585.475.2139
fax: +1.716.475.7100
e-mail: ncs@cs.rit.edu
May 13, 2004

Part 3. The problem to be addressed, including analysis of how it currently interferes with instruction and how it would improve your class if successfully addressed.

Last Spring quarter (20023), I taught Parallel Computing II for the first time in many years. This course was taught most of the time that it has been offered by Andrew Kitchen who recently retired. When Andrew taught the course, it was a very interesting (to me) lecture course. However, the textbook that he used is no longer in print and, indeed, the face of parallel computing has changed sufficiently in the past few years that the course needed to be revamped.

I had two goals in mind when setting up the course last spring: 1. To introduce an element of computational science to the course by having the course projects involve parallelizing "real-world" programs in use in other disciplines. 2. To run the course in a collaborative learning manner. I was able to meet the first goal without any difficulties. However, my success with the second goal was much less.

To meet the second goal, I had the students decide on the topics that would be discussed in the course, i.e., we decided together which topics in parallel computing were interesting to them, individually, and to the class as a whole. Each student prepared a lecture, including a reading list and discussion questions. I had requested that the reading lists and discussion questions be made available one week before the class in which the presentation was to be made. While the students selected a relatively diverse set of topics and, for the most part, had the reading lists available in a timely fashion, the discussion questions were NOT made available in a timely fashion and often were not such as to engender discussion. Indeed, they were instead often made available only a few minutes before the class met. As a result, the discussion aspect of the course was sorely lacking. This is the problem I wanted to address.

Part 4. Proposed solution

While I had no concrete evidence as to why there were problems with the discussion portion of the course, I did have the feeling that the students prepared their presentations at the last minute and were so focused on getting that done that the discussion question aspect fell by the wayside. It was not a specifically graded component. In addition, I was not sure if they knew what constituted a good discussion question.

As I had identified three possible causes of the problem, I proposed to address each of them, in particular, by

1. Making the lecture preparation a shared experience: I would pair students to work together on each lecture topic. One would be responsible for preparing the lecture and the other for preparing the discussion questions, having them available in a timely fashion, and leading the discussion portion of the lecture.

2. Using an active learning exercised used by Laurie Richlin at the Lilly Conference in her session *Designing Questions that Get Students to Speak* to help attendees discover how to create questions that are more likely to engender discussion. This technique was my first exposure to a totally active learning situation and was extremely convincing to me. It was based on an article "The Verbal Structure of Teacher Questions: Its Impact on Class Discussion" written by John D. W. Andrews and appearing in the *POD Quarterly*, Vol. 2, Nos. 3 & 4 (Fall/Winter 1980). This session is described below:

    a. When we walked into the room, we were given a *Question Matching Exercise*, which asked us to pair similar questions. Similarity was defined to be similarity in question types rather than content.
    b. After we had finished the matching, we were asked to find other people in the room who had the same answers as we did.
    c. The questions were typed in large font on separate sheets of paper. A group of agreeing people hung their pairing of these questions on the wall.
    d. Laurie led a discussion that corrected the mismatches.
    e. Next she handed out a page entitled *Types of Questions* by John D.W. Andrews. We then typed the pairs of questions.
    f. Laurie reported that there had been a study done in which a classroom was taped for some period of time. During this time period, questions of these types were asked and the number of responses for each type of question was recorded. We were then asked to order the pairs of questions with respect to the highest to lowest number of responses.
    g. The results were discussed and this ordering was corrected.
    h. The final exercise was to write questions of the type that could be expected to produce the largest response. We were asked to share our questions and an analysis ensued as to whether our questions met the stated criteria.

While I am not on top of the questions types as a result, I did feel like I had learned something! I hoped to give my students a similar experience. And, I hoped this would aid them in designing more effective discussion question, which in turn would make the class a more lively and interesting place to be.

I plan to this exercise by assigning a related technical article to be read and asking the students to bring in one discussion question, i.e., to identify what they would lie the class to discuss about the article. This would be done in the classes directly preceding and following the question matching exercise. My faculty and student partners would attend both classes, where they would record student response in both classes. I hoped that there would be a qualitative difference as a result of the active learning exercise regarding types of discussion questions. Students would be asked about their response to this exercise in an assessment session later in the term.

(Complete information about the exercise is available in Appendix A.)

3. Adding a component to the assessment portion of the course that would include the effectiveness of the discussion questions and another that would include participating in the discussion.

4. Factoring in student characteristics by having the students complete the Myers-Briggs Personality (http://similarminds.com/myers-briggs.html).

Part 5. A Teaching Goals Inventory for the class (http://www.uiowa.edu/~centeach/tgi/)

Teaching Goals Inventory... Online!

Teaching Goals Inventory Results

Parallel Computing II
12/10/2003

This table contains your results. The third column contains the percentage of items within each cluster that you rated "essential." The fourth column contains the average rating you assigned to items within each cluster.

Cluster	Goals Included in Cluster	Percent Rated "Essential"	Mean Rating
I. Higher Order Thinking Skills	1-8	50%	4.13
II. Basic Academic Success Skills	9-17	-- none --	1.67
III. Discipline-Specific Knowledge and Skills	18-25	25%	3.50
IV. Liberal Arts and Academic Values	26-35	-- none --	1.40
V. Work and Career Preparation	36-43	25%	4.25
VI. Personal Development	44-52	22%	2.11

You identified your primary role as a teacher as "Helping students develop higher-order thinking skills."

It may be useful to compare your results to those of a large sample of teachers. The following table provides mean cluster ratings and the average percentage of items in each

Math skills (17)						61		60	84
Terms and facts (18)						61		60	
Wise decisions (52)							70		
Analytic Skills (2)			66						73
Self-esteem (45)				63					
Think for self (51)	66	59	75	65	50				
Responsible for self (44)							68		
Value of subject (21)		56			52				
Concepts and theories (19)								71	
Creativity (7)	69								
Writing skills (15)			84						
Aesthetic appreciation (31)	78								
Openness to ideas (27)		56							
Problem solving (3)						57			84

The rest of this report lists the goals you rated sorted into groups according to the rating you assigned.

subject
25. Learn to appreciate important contributions to this subject
27. Develop an openness to new ideas
30. Develop a lifelong love of learning
44. Cultivate a sense of responsibility for one's own behavior

Goals You Rated "Unimportant"

Goals You Rated "Not Applicable"

8. Develop ability to distinguish between fact and opinion
9. Improve skill at paying attention
10. Develop ability to concentrate
11. Improve memory skills
15. Improve writing skills
16. Develop appropriate study skills, strategies, and habits
17. Improve mathematical skills
18. Learn terms and facts of this subject
26. Develop an appreciation of the liberal arts and sciences
28. Develop an informed concern about contemporary social issues
29. Develop a commitment to exercise the rights and responsibilities of citizenship
31. Develop aesthetic appreciation
32. Develop an informed historical perspective
33. Develop an informed understanding of the role of science and technology
34. Develop an informed appreciation of other cultures
35. Develop capacity to make informed ethical choices
45. Improve self-esteem/self-confidence
46. Develop a commitment to one's own values
47. Develop respect for one's own values
48. Cultivate emotional health and well-being
49. Cultivate physical health and well being
50. Cultivate an active commitment to honesty

Center for Teaching

Part 6. Faculty and student partners – who and roles

Faculty Partner: Larry Quinsland, NTID Science and Mathematics, 2285 Johnson, lkq9999@rit.edu, 475-6237, cell: 230-3369

Role:

- Help design experiment
- Help design assessment
- Participate in experiment (March 15, March 22)
- Participate in assessment (April 28)
- Will debrief faculty member

Student Partner: Robert Whitcomb, rjw2183@cs.rit.edu, Computer Science graduate student who participated in the first offering of the course.

Role:

- Will attend classes (March 15, March 17, March 22) when the three experimental phases will take place and will record student interaction
- Will record assessment session (April 28)
- Will debrief faculty member

Part 7. Plan for assessing projects effectiveness, including feedback from students and classroom assessment techniques.

Besides using the normal student evaluation forms, the success of the project will be assessed using a technique suggested by my teacher partner and found in *Cooperative Learning For Higher Education Faculty* by Barbara J. Millis and Philip G. Cottell, Jr., pages 226-228. (See Appendix B.) The technique is called Small-Group Instructional Diagnosis and is designed to gather data that may be useful in strengthening classes predicated on cooperative learning. It consists of four parts:

1. **Pre-interview meeting:** My faculty partner and I will come up with the set of interview questions.

2. **Thirty-minute in-class interview:** Week 8, Day 2
   My faculty partner will meet with the class and runs the interview.
   My student partner will record the results.
   My faculty partner will analyze and organize the materials to make it meaningful/useful to me

   Potential Questions:
   Why do teachers want you to discuss?
   Being a participant in a discussion is a skill.
   Perception of discussion skills
   Attitude toward Lilly Conf. Experiment
   Did your ability to create questions that generated discussion improve?
   Did you want it to?

3. **Debriefing:** My faculty and student partners will meet with me to discuss the interview process, analyze the data, and plan further action

4. **Class debriefing:** Thank students for honest comments and discuss the results and what I will do to address the issues that have been identified.

d. **Week 3 or 4**: I will meet with student partner (and teacher?) to discuss findings.

OUTCOME:
Done with the results as indicated above.

e. **Weeks 5-7:** Students will prepare lectures and discussion questions

OUTCOME:
In general, the lectures were very well done. The discussion questions were in on time. The student partners led the discussion sessions. For the most part, there was much more discussion than last year! Unfortunately, the only measure of this is qualitative on my part.

Also, I was fortunate that the very first set of discussion questions was very good and may have set the tone for the rest. For example, all students prepared exactly three questions as did the first student. There was no set number specified.

f. **Week 8:** I will be out of town, but student and teacher partners will do an in class assessment of student reaction to the experiment.
Student and teacher partners will debrief me.
Results will be added to poster session for April 30.

OUTCOME:
Both my faculty and student partners attended the assessment session. My faculty partner led the discussion. He said the students warmed up after about five minutes. He found them to be very open, although he had to specifically solicit input from some of them. My faculty partner provided a written report, shown in Figure 6 with which my student partner concurred.

g. **Week 9:** Debrief class on results of the experiment.
Answer concerns that came up as a result of the assessment session.

OUTCOME:
Although I did do this, there was very little discussion as a result.

Part 9. Outcomes

Spring Quarter:

a. **Week 2, Day 1:** Pre-active learning exercise assignment to bring discussion question relevant to reading assignment to class
Both teacher and student partners will attend.
Student partner will (learn to) record class discussion interaction.

OUTCOME:
Both my faculty and student partners attended the class and recorded the interaction. Figure 1 shows the key to the other figures. Figure 2 shows the recording of the interaction that was done by my faculty partner. Figure 3 shows the recording of the interaction that was done by my student partner. Notice that my student partner tended to get involved with the discussion rather than attending to the recording. Figure 4 shows the summary of the interaction as provided by my faculty partner.

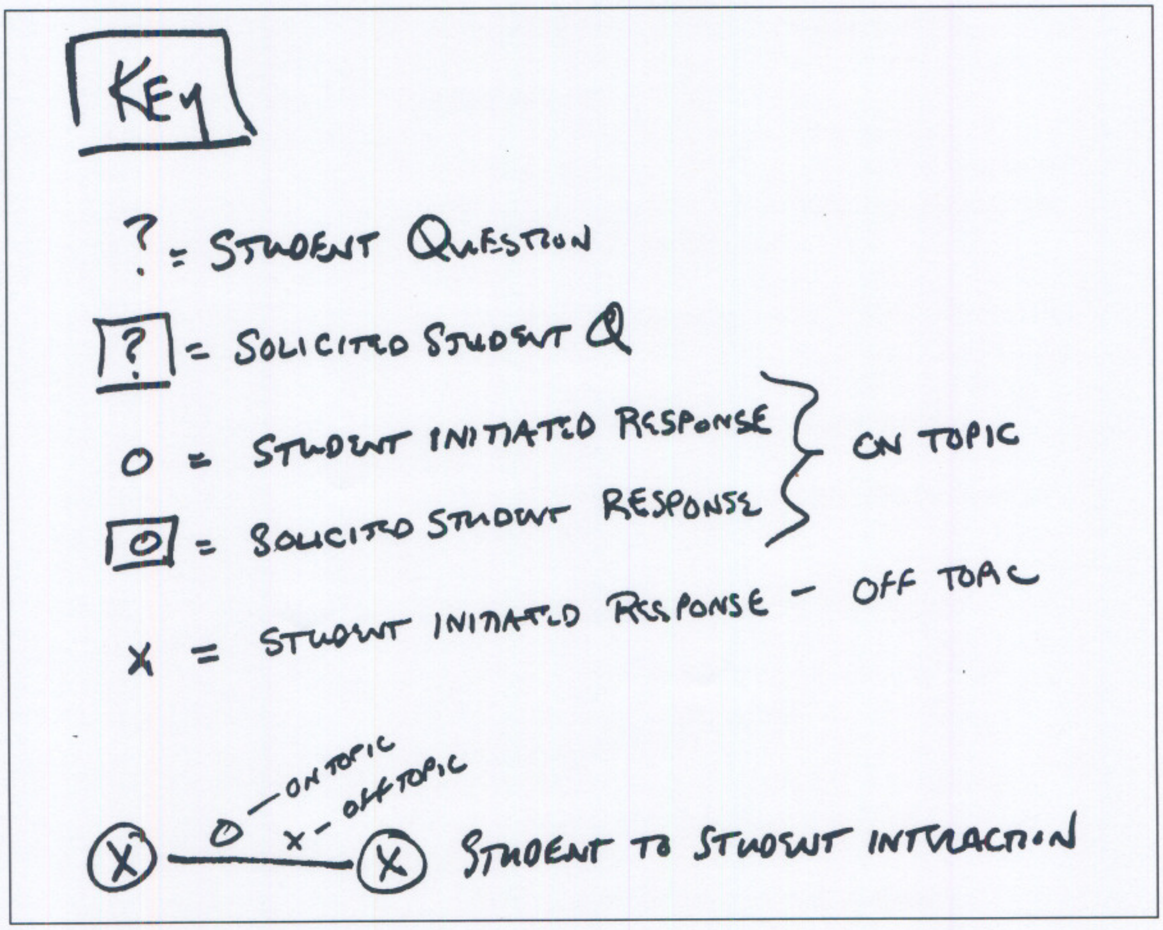

Figure 1. Key for Figures 2, 3, and 5

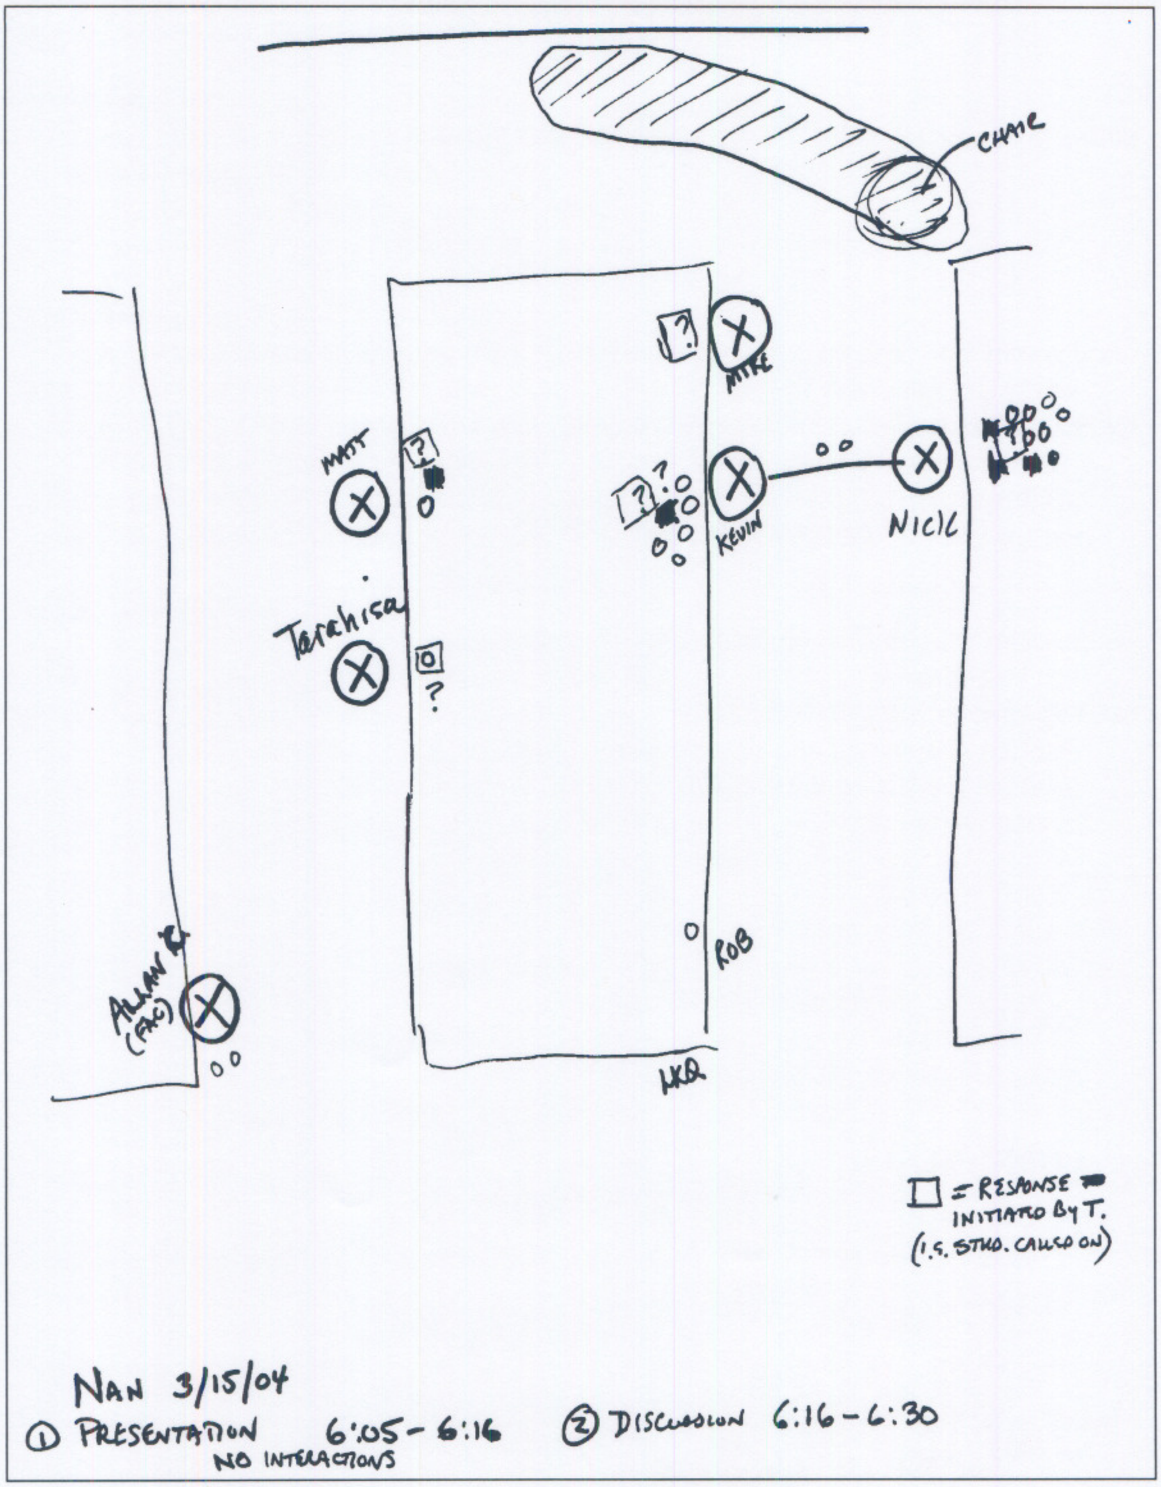

Figure 2. Interaction as recorded by faculty partner

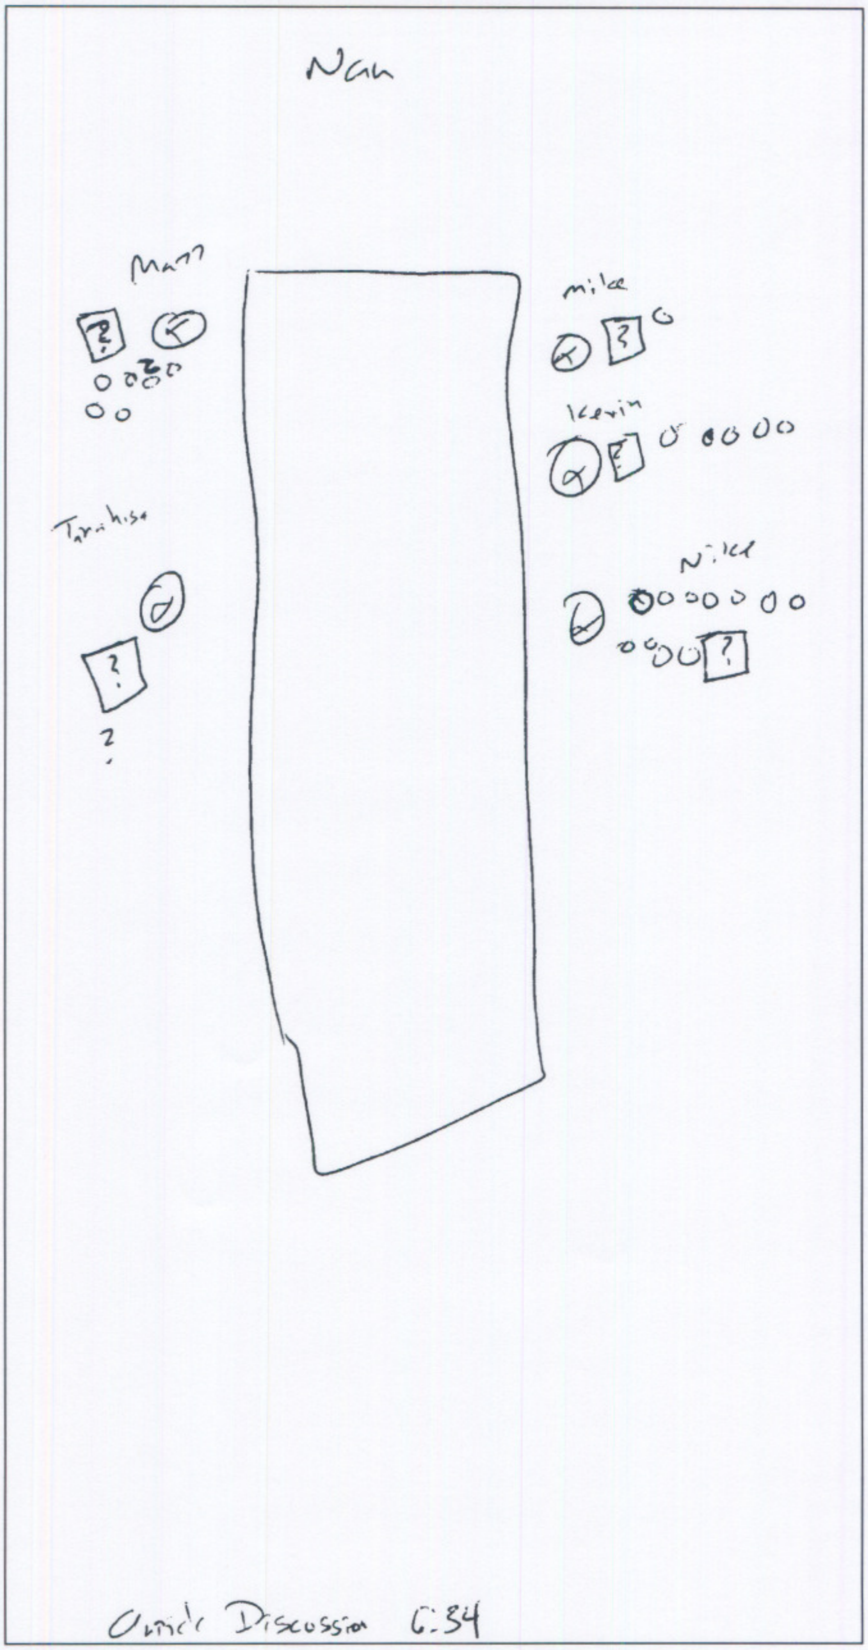

Figure 3. Interaction as recorded by student partner.

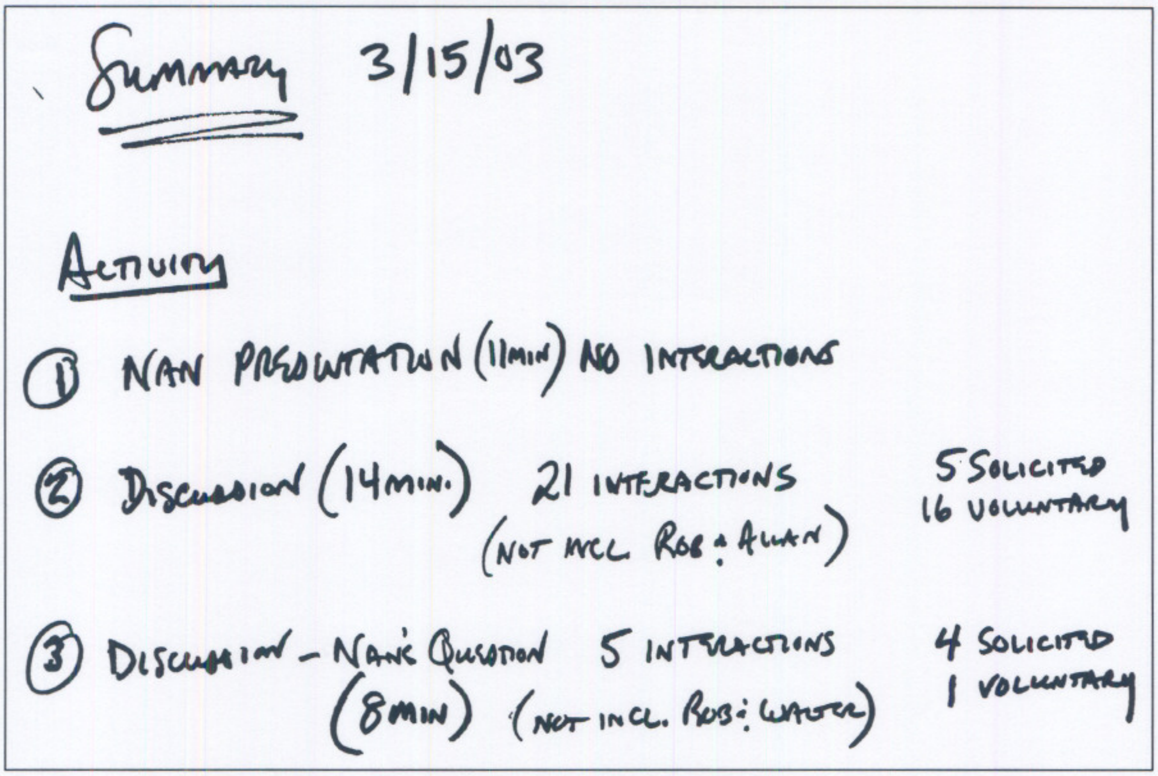
Summary 3/15/03

Activity

① NAN PRESENTATION (11MIN) NO INTERACTIONS

② DISCUSSION (14MIN.) 21 INTERACTIONS (NOT INCL. ROB & ALLAN) 5 SOLICITED 16 VOLUNTARY

③ DISCUSSION – NAN'S QUESTION (8MIN) 5 INTERACTIONS (NOT INCL. ROB & WALTER) 4 SOLICITED 1 VOLUNTARY

Figure 4. Faculty partner summary

b. **Week 2, Day 2:** Run the active learning exercise.
Student partner will attend (participate in?) active learning exercise.
Students will be asked to take Myers-Briggs personality type quiz and to discuss the distribution of personality types in the class and how that might affect the amount of discussion we are able to generate.

OUTCOME:
The students did amazingly well with this exercise, much better than the group at the Lilly Conference! They correctly matched the types of questions and had just one of the levels of response out of order!!!! I was amazed and all five of them worked to come to a consensus. (Unfortunately, my student partner forgot to come to class.)

Myers-Briggs Personality Types of class; I was amazed that there were no two with the same results, but was equally interested to see that there were a couple of E's amongst the students:

Instructor: ESTJ
Nick: ENTP
Kevin: ISTP
Mike: ISTJ
Matt: ENFP
Takahisa: INFP

c. **Week 3, Day 1:** Post-active learning exercise assignment to bring discussion question relevant to reading assignment to class
Student partner will attend and record class discussion interaction.

OUTCOME:
While my student partner did attend the class and agreed that there was much more discussion, his recording of the interactions was incomplete and did not show this. Figures 5 shows what he recorded. Notice even the solicited questions were not shown. Unfortunately as well, my faculty partner was out of town at a conference and could not attend

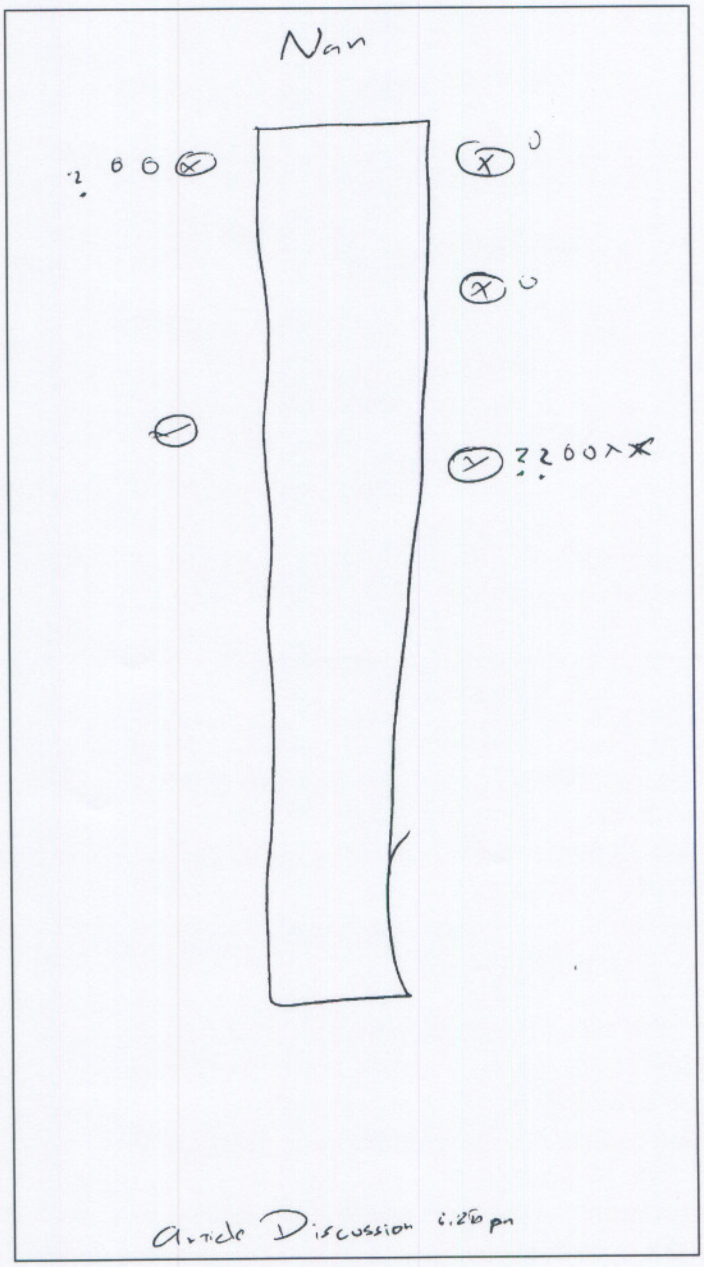

Figure 5. Student recording of interaction for after experiment

d. **Week 3 or 4**: I will meet with student partner (and teacher?) to discuss findings.

OUTCOME:
Done with the results as indicated above.

e. **Weeks 5-7:** Students will prepare lectures and discussion questions

OUTCOME:
In general, the lectures were very well done. The discussion questions were in on time. The student partners led the discussion sessions. For the most part, there was much more discussion than last year! Unfortunately, the only measure of this is qualitative on my part.

Also, I was fortunate that the very first set of discussion questions was very good and may have set the tone for the rest. For example, all students prepared exactly three questions as did the first student. There was no set number specified.

f. **Week 8:** I will be out of town, but student and teacher partners will do an in class assessment of student reaction to the experiment.
Student and teacher partners will debrief me.
Results will be added to poster session for April 30.

OUTCOME:
Both my faculty and student partners attended the assessment session. My faculty partner led the discussion. He said the students warmed up after about five minutes. He found them to be very open, although he had to specifically solicit input from some of them. My faculty partner provided a written report, shown in Figure 6 with which my student partner concurred.

g. **Week 9:** Debrief class on results of the experiment.
Answer concerns that came up as a result of the assessment session.

OUTCOME:

NAN - Classroom Assesment - Student Interview
Dr. L.K. Quinsland 28 April 2004

(NOTE = Anything that is in parentheses, is Dr. Quinsland's comment)

1. **Did Nan describe this project to you? What is your perception?**

NO CLUE

2. **Describe the discussion component of this class. (What did Nan do? What did you do?)**

Nan…
First we read articles and students wrote questions
Discussions were stimulated by student questions
She introduced the question matching experiment (they did not know the name of this)
Asked Kevin to read recent news articles (because he is the only grad student)… did not lead to much discussion (Kevin)
Students did presentation …. other students developed questions groups of two.
 person presenting was paired with other student
 presenter had week to prep answer… (Students agreed that this led to most discussion of ALL activities)

Students… (Only 1 stud. responded to this…)
Student said he sometimes asked questions to get back into mode because was drifting (falling asleep)

3. **Why do you think professors want you to discuss?**

It gets student thinking about materials… rather than passive
Sparks interest
Makes student look deeper into topic
Helps professor READ what student know
So professor can talk less (cynical guess?)
Assess what student interest is

4. **Do students need discussion skills?** (They said they had never thought about this)

Forces student to think more… not passive
Might learn something vicariously (from other students… something did not know)
Sometimes student can explain more clearly than professor

5. **Have your discussion skills changed this quarter?**

n/a

6. **About the Question Matching Exercise**

Thought about it a little bit
Led to variation in questions development when I had to make up questions
Our questions were better for presentation than for article (students agreed that they developed better questions for presentations than for initial article review discussions… probably due to Lilly activity)
More aware of questions types and intent (due to Lilly activity)

Did you ability to create questions that generated discussion improve?

See above

Did you want it to improve? Do you care?
(long pause… they clearly had not considered that improving their respective discussion skills was part of the deal of being a student… However, all 5 students suddenly realized how they benefited from the process!)

Yes
Won't hurt
If it helps learn yes
I probably will think about questions more
Yes.. I have a presentation at my business… I am considering questions much more due to this experience…

Is there anything Nan could do that would encourage you to discuss more?

Do questions match earlier (group expressed pros and cons…… maybe would not have appreciated activity if did not have article activity first… experience with developing inappropriate questions … the key here was that the initial experience of developing questions had no guidelines… after the question matching activity… they had a reference upon which to judge the questions they generated for the presentations… they clearly did not realize this until we had the discussion tonight… but seemed to be proud of the ability to CONSIDER/THINK ABOUT and EVALUATE the questions they would ask of the presenter)
Maybe would not work with other classes… e.g. = hard time in CS1
Would like more specific feedback on our questions… (e.g. appropriateness of the questions we generated… would have benefited from feedback re: the questions they generated…

Figure 6. Classroom Assessment Report

Part 10. Reflections that process the discussion in each meeting of the Faculty Learning Community and apply topics to your own experience.

Fall Quarter:
The fall quarter suffered somewhat because of an inability to find a common meeting time. As a result, I was able to attend only one hour of the two-hour meetings for two out of three meetings. This limited community bonding, my interaction with the rest of the community and limited the effectiveness of the discussions for me.

Meeting 1: (9/8/03)
Group introduction

Meeting 2: (9/15/03)
Preparing to offer a course
Reading: Part I of McKeachie's Teaching Tips
Bring: Your teaching philosophy
Your metaphor for teaching
Given: Teaching Goals Inventory

Meeting 3: (9/29/03)
Meeting the class for the first time
Reading: McKeachie Chapter 3

Meeting 4: (10/13/03)
Classroom Assessment
Readings: McKeachie Chapter 6
Angelo: pp 3-32
Blooms Taxonomy on the web

Meeting 5: (10/27/03)
Facilitating Active Learning
Reading: McKeachie Chapters 13-15

Meeting 6: (11/10/03)

Winter Quarter:
Most of the meetings in the winter quarter involved talking about and trouble shooting our projects. This was very helpful in several ways: helping to narrow the project focus, helping to identify and fill holes, and providing necessary support.

Spring Quarter:
Continuation of winter quarter activities, plus discussed continuation and the community with Lynn Wild and David Neumann, as well as getting ready for our presentation for the RIT community and FITL. While I enjoyed these meetings, this quarter did not really add much to the process for me.

Appendix A. Question Matching Exercise Materials, including answers and article

Question Matching Exercise

Instructions:

The 14 questions below consist of seven pairs of similar question types in mixed order. The purpose of this exercise is to heighten awareness of the characteristics of various sorts of questions through noting similarities and differences. Please compare the questions in the two columns below, in terms of their wording and their likeliness to engender discussion. When you identify two that seem similar, indicate this by putting the number from the question in the left hand column in the blank next to the right hand matching question.

1. So where is this wild boy better off? In the forest where he started, or in civilization being socialized?
2. So in this story, when's the point of truth for Kurtz?
3. What about the lecture?
4. What possibilities are there for refuge in *A Farewell to Arms*?
5. So we're talking about the fact that everybody's roles are changing, how—we've mentioned religion and education, how did religion and education during this period affect these changes, or how did the changes affect the kind of religion and education people had? … Let's start with religion … Have women always had a divine place in religion?
6. Let's see if we can make any generalizations about the play as a whole, from the nature of the opening lines.
7. What was the name of that institution?

___ A. (Instructor reads a sentence from the novel under study): Well that's a very rich sentence … there's a lot there… OK, what's there?

___ B. Any comments on Pluto?

___ C. He talks about envying one character. Who was it?

___ D. How do you interpret what the narrator tells you about the hero? What do you make of his return from law school? Why did he decide he didn't really expect too much?

___ E. Is Ivan Illyich a victim of society, or did he create his problems by his own choices?

___ F. What was the most important reason for the revolution's failure?

___ G. What kinds of things is Hamlet questioning? Not just in his soliloquy, but broadly, throughout the whole play?

Question Matching Exercise
ANSWERS

Instructions:

The 14 questions below consist of seven pairs of similar question types in mixed order. The purpose of this exercise is to heighten awareness of the characteristics of various sorts of questions through noting similarities and differences. Please compare the questions in the two columns below, in terms of their wording and their likeliness to engender discussion. When you identify two that seem similar, indicate this by putting the number from the question in the left hand column in the blank next to the right hand matching question.

1. So where is this wild boy better off? In the forest where he started, or in civilization being socialized?
2. So in this story, when's the point of truth for Kurtz?
3. What about the lecture?
4. What possibilities are there for refuge in *A Farewell to Arms*?
5. So we're talking about the fact that everybody's roles are changing, how—we've mentioned religion and education, how did religion and education during this period affect these changes, or how did the changes affect the kind of religion and education people had? … Let's start with religion … Have women always had a divine place in religion?
6. Let's see if we can make any generalizations about the play as a whole, from the nature of the opening lines.
7. What was the name of that institution?

_6_ A. (Instructor reads a sentence from the novel under study): Well that's a very rich sentence … there's a lot there… OK, what's there?

_3_ B. Any comments on Pluto?

_7_ C. He talks about envying one character. Who was it?

_5_ D. How do you interpret what the narrator tells you about the hero? What do you make of his return from law school? Why did he decide he didn't really expect too much?

_1_ E. Is Ivan Illyich a victim of society, or did he create his problems by his own choices?

_2_ F. What was the most important reason for the revolution's failure?

_4_ G. What kinds of things is Hamlet questioning? Not just in his soliloquy, but broadly, throughout the whole play?

Types of Questions

John D.W. Andrews

The Analytic Convergent Question: These questions imply that there is a single correct answer. The asker has a single endpoint in mind that the responder can be expected to identify. They usually involve comparisons, cause and effect statements, drawing of straightforward inferences from limited material, but for a desired end.

The Brainstorm Question: Any and all ideas or solutions are sought in response to a specific question or problem. The primary subject is thematic, with the subject matter slightly unfocused.

The Quiz Show Question: These questions call for facts, definitions, and descriptions of events. They call for memory or comprehension level responses. The responses are generally brief.

The Focal Question: Is ordered around an issue that calls for decisions. The asker poses a fixed number of alternatives, usually two or three, and asks the responders to take stands that they are to justify during discussion. These stands call for higher order thinking and require the responder to marshal various kinds of information to support their views.

The Playground Question: The hallmark of the playground question is the invitation to explore. The asker designates a specific intellectual sphere (the "playground") for discussion and then gives the responder the widest possible latitude in approaching it. Playground questions frequently begin with the phrases like "How do you interpret", "What can you draw from", or "What are the possible meanings of".

Shotgun Question: Many questions all at once that may not even relate to each other.

The General Invitation: Is a playground question that has lost its boundaries. A general invitation question is broad, vague, and often defined by a handy or conventionally obvious unit of the course, lecture, or book. It is a license to wander all over the place in discussion.

Types of Questions (ANSWERS)

John D.W. Andrews

__2F The Analytic Convergent Question:__ These questions imply that there is a single correct answer. The asker has a single endpoint in mind that the responder can be expected to identify. They usually involve comparisons, cause and effect statements, drawing of straightforward inferences from limited material, but for a desired end. **(#6– 1.95)**

__4G The Brainstorm Question:__ Any and all ideas or solutions are sought in response to a specific question or problem. The primary subject is thematic, with the subject matter slightly unfocused. **(#2 4.88)**

__7C The Quiz Show Question:__ These questions call for facts, definitions, and descriptions of events. They call for memory or comprehension level responses. The responses are generally brief. **(#7 1.45)**

__1E The Focal Question:__ Is ordered around an issue that calls for decisions. The asker poses a fixed number of alternatives, usually two or three, and asks the responders to take stands that they are to justify during discussion. These stands call for higher order thinking and require the responder to marshal various kinds of information to support their views. **(#3 4.29)**

__6A The Playground Question:__ The hallmark of the playground question is the invitation to explore. The asker designates a specific intellectual sphere (the "playground") for discussion and then gives the responder the widest possible latitude in approaching it. Playground questions frequently begin with the phrases like "How do you interpret", "What can you draw from", or "What are the possible meanings of". **(#1 5.08)**

__5D Shotgun Question:__ Many questions all at once that may not even relate to each other. **(#5 2.5)**

__3B The General Invitation:__ Is a playground question that has lost its boundaries. A general invitation question is broad, vague, and often defined by a handy or conventionally obvious unit of the course, lecture, or book. It is a license to wander all over the place in discussion. **(#4 2.6)**

Reordering of Types of Questions Producing Most Discussion
John D.W. Andrews

6A The Playground Question: The hallmark of the playground question is the invitation to explore. The asker designates a specific intellectual sphere (the "playground") for discussion and then gives the responder the widest possible latitude in approaching it. Playground questions frequently begin with the phrases like "How do you interpret", "What can you draw from", or "What are the possible meanings of". **(#1 5.08)**

6. Let's see if we can make any generalizations about the play as a whole, from the nature of the opening lines.

A. (Instructor reads a sentence from the novel under study): Well that's a very rich sentence … there's a lot there… OK, what's there?

4G The Brainstorm Question: Any and all ideas or solutions are sought in response to a specific question or problem. The primary subject is thematic, with the subject matter slightly unfocused. **(#2 4.88)**

4. What possibilities are there for refuge in *A Farewell to Arms*?

G. What kinds of things is Hamlet questioning? Not just in his soliloquy, but broadly, throughout the whole play.

1E The Focal Question: Is ordered around an issue that calls for decisions. The asker poses a fixed number of alternatives, usually two or three, and asks the responders to take stands that they are to justify during discussion. These stands call for higher order thinking and require the responder to marshal various kinds of information to support their views. **(#3 4.29)**

1. So where is this wild boy better off? In the forest where he started, or in civilization being socialized?

E. Is Ivan Illyich a victim of society, or did he create his problems by his own choices?

3B The General Invitation: Is a playground question that has lost its boundaries. A general invitation question is broad, vague, and often defined by a handy or conventionally obvious unit of the course, lecture, or book. It is a license to wander all over the place in discussion. **(#4 2.6)**

3. What about the lecture?

B. Any comments on Pluto?

5D Shotgun Question: Many questions all at once that may not even relate to each other. **(#5 2.5)**

5. So we're talking about the fact that everybody's roles are changing, how—we've mentioned religion and education, how did religion and education during this period affect these changes, or how did the changes affect the kind of religion and education people had? … Let's start with religion … Have women always had a divine place in religion?

D. How do you interpret what the narrator tells you about the hero? What do you make of his return from law school? Why did he decide he didn't really expect too much?

2F The Analytic Convergent Question: These questions imply that there is a single correct answer. The asker has a single endpoint in mind that the responder can be expected to identify. They usually involve comparisons, cause and effect statements, drawing of straightforward inferences from limited material, but for a desired end. **(#6– 1.95)**

6. So in this story, when's the point of truth for Kurtz?

F. What was the most important reason for the revolution's failure?

7C The Quiz Show Question: These questions call for facts, definitions, and descriptions of events. They call for memory or comprehension level responses. The responses are generally brief. **(#7 1.45)**

7. What was the name of that institution?

C. He talks about envying one character. Who was it?

the result is positive, people accept their good fortune happily enough but with little sense that they know how to make it recur. The thoughtful use of our more powerful question types should give teachers another tool for constructing successful discussions.

I find little support in our findings for the fear that rigor will be sacrificed in producing open, involved discussions (or vice versa). I suspect that this fear may stem from some half-truths about both qualities. Specifically, if one associates rigor with gaining a lot of knowledge, and thus with Quiz Show Questions or Analytic Convergent Questions, the best discussion *will* be lost. And conversely, if one associates free discussion with the shelving of standards or the substitution of opinion for the thoughtful use of information, rigor *will* diminish. The latter is most apt to happen through overuse of General Invitations with the vagueness and lack of direction which results.

The Structured Divergent Questions help us to have our cake and eat it too. They provide space for students to bring in new ideas, yet each type has its own built-in directionality. While they encourage students to think for themselves, there is also encouragement to do that thinking *about* something: to incorporate the facts of the discipline into the discussion.

Use of these questioning strategies should help students develop intellectually and become more nearly full partners in the learning process. Because they can be explicitly defined and explained to students, the Structured Divergent Questions are a good basis for giving roles in the class. They can organize their own thinking through constructing questions, leading discussions, and so on. The overall result should be a more sophisticated approach to learning in general. In this connection, it is interesting in that the three Structured Divergent Questions call for thinking characteristic of the higher intellectual and ethical stages delineated by Perry (1970) in his developmental study of college students. It may turn out that these teaching tools can help us foster such maturational progress as well as helping students acquire cognitive skills *per se*.

Finally, the conclusions reached here suggest many issues for further exploration. Are these question categories exhaustive, or will other "species" turn up? What differences would we see in studying other types of classes, for example in the natural sciences? Is the questioning pattern of the instructor related to the way stu-

dents rate him or her on evaluation questionnaires? Is questioning pattern related to instructor personality? Is it related to outcome: that is, to what students learn as measured by various testing criteria? Do students with different personality patterns, learning styles, or levels of development respond differently to our various questioning patterns? Answering these and other questions should help us construct a more complete anatomy of the effective discussion.

BIBLIOGRAPHY

Bloom, B. S., *Taxonomy of Educational Objectives.* New York: Longmans Green, 1956.

Erikson, E. H., "The Problem of Ego Identity." In Stein, M. R., Vidich, A. J., & White, D. M., *Identity and Anxiety.* Glencoe, Ill.: The Free Press, 1969. Pp. 37–87.

Guilford, J. P., *Intelligence, Creativity, and Their Educational Implications.* San Diego: R. R. Knapp, 1968.

Napell, S., "Six Common Non-Facilitating Teaching Behaviors." *Contemporary Education,* Vol. XLVII, No. 2, Winter, 1976. School of Education, Indiana State University.

Osborne, A. F., *Applied Imagination.* New York: Scribner, 1963.

Perry, W. G., Jr., *Forms of Intellectual and Ethical Development in the College Years: A Scheme.* New York: Holt, Rinehart & Winston, 1970.

improbable-looking connections between the theme and other aspects of the material. One of the virtues of this thematic approach is that it helps students unify a topic, book, or other unit of material, as the theme suggests interrelationships between apparently disparate subparts. Once having selected the theme, introduce it to students with some variant of the following: "This theme (issue) arises in lots of ways when we examine ________________. What are some ways you can think of?"

c) Focal Questions involve either-or decisions or controversies, and so the first choice of structure is to set these terms. The alternatives should all be tenable or defensible. They should be pivotal in that one's choice among them has significant ramifications for one's view of a larger topic or sphere of discussion, and they should enable the student to draw on a sizable portion of the material in this topic area in support of one or another viewpoint. The question may involve a complex problem with several possible solutions or interpretations; a controversy (concerning values or analytic conclusions) with several possible stands to be taken; a phenomenon which can be interpreted according to two or more theories; or a complex, provocative statement which can be agreed or disagreed with.

A good process for generating all three types of questions is brainstorming, either for yourself or with the class. When the latter is done, students become acquainted with the question form, which helps them participate effectively. They can use their own interests as criteria of selection, and everyone will feel more investment and responsibility for the ensuing discussion. For example, one class concerned with ecology and environment chose as a Focal Question to discuss whether Eskimos should be restricted in their killing of seals. They recognized this as internally dynamic because it pitted against each other two values which this group of students supported strongly: the preservation of the Eskimos' independence and traditional way of life on the one hand, and the protection of an endangered animal species on the other. And they could draw widely in search of support for one argument or the other: on ethics, political decision-making and the role of government, ecological biology, anthropology, and so on.

Structured Divergent Questions can be used in many formats. They can be introduced *ad hoc* in discussion; assigned as study questions or as pre-planned discussion questions; and used as essay exam questions. Learning is strengthened when the type of thought encouraged in discussions is mirrored in the evaluation procedure as

well. Such questions can also be a powerful stimulus to learning factual information, because the latter is seen as useful for dealing with meaty problems and issues. One approach is to pose a question near the end of one class session, elicit enough discussion to catch students' interest, and then ask them to think about it during their studying in preparation for the next meeting.

The three types of Structured Divergent Questions can be expanded into classroom formats useful for structuring one or more whole sessions.[16] Projecting them on such a giant screen also helps clarify their characteristics and the differences among them.

a) The format corresponding to Brainstorm Questions is, not suprisingly, the brainstorming session. This process was developed some time ago (Osborne, 1963) in order to stimulate creativity by reducing the inhibiting effects of criticism and argumentation. The basic ground rule divides the session into two parts: the first is devoted to idea production; no holds are barred, and apparently zany ideas are actively encouraged while evaluative commentary is out of bounds. Generally, the leader (teacher) records all contributions visibly, on the blackboard for example. The underlying attitude is, "Why come up with two ideas and then spend the next 30 minutes arguing over them, when in the same time period we can come up with 20 ideas and then select the best among them?" The second phase of the brainstorming session is concerned with combining ideas, applying criteria, and deciding which are truly valid and useful.

b) The Focal Question can be used in any format which enables students to declare themselves on an issue and argue back and forth. One approach is to ask for a hand vote periodically. Focal Questions also lend themselves well to a debate format, and I have found a debate variant which is extremely effective in college classes. This is called by its inventor[17] the "change-your-mind-debate" because, unlike traditional debates, it encourages the participants to let the discussion influence them toward new views. The room is set up with three banks of chairs (when there are two alternatives posed); one bank is for those who espouse each of the two main positions, and the third is for those who have mixed reactions or are undecided. Students are instructed that whenever they change their minds during the debate, they are to get up and move to the area which reflects their current views. Thus,

[16] While we do not have enough examples of these formats to conduct a systematic study, the three which are available (one of each) are extremely effective by the discussion mileage criteria used in this resarch. They all have NSS scores well beyond the range of the questions analyzed in Table 7.

[17] My thanks to Zachary Seech, a graduate student in the Philosophy department, University of California, San Diego, for this idea.

gent questions are more effective, it follows mathematically that one who uses a greater proportion of them will have a higher overall NSS score.

This distortion can be controlled by examining the relationship between %D and NSS *within* the divergent and convergent subcategories taken separately. Here, the results are quite clear. Among our seven instructors, the higher the %D score, the greater the question mileage (mean NSS score). This is most striking when we consider only divergent questions: the rank-order correlation between %D and NSS is $+.82, P < .03$. Among the convergent questions, the relationship is still present, but weaker: $rho = +.60$, $P < .15$. Moreover, these mileage differences are far from trivial, especially among the divergent questions; the highest scoring instructor had a mean NSS score of 5.71, while the score of the lowest was 1.43. The range for convergent questions was 3.33 to 0.94.[14]

The most plausible interpretation of this finding is that when the overall style is divergent, students develop expectations which make them more ready to respond; probably they gain practice in dealing with such questions and feel more expansive and free from judgments. It is interesting that this climate factor apparently carries over even when convergent questions are asked; perhaps in divergent-oriented classes students give ambiguous questions the benefit of the doubt, interpret even a convergent structure in a divergent way, or feel freer to raise divergent issues around overtly convergent questions.

To summarize the findings: the verbal form of questions does influence the extent of discussion response, and it is possible to identify several types of questions which have the most productive results. These, collectively called Structured Divergent Questions, have the common characteristic of providing students considerable freedom of expression within the frame of a definite focus which holds the class together and provides a sense of continuity and direction. This impact is predictable and consistent cross instructors. However, it occurs within a larger context of instructor style: although all seven instructors got better mileage when using such questions, the most fruitful results seem to come from building a

[14] It should be borne in mind that these instructor differences do not contradict the finding that for *all* instructors divergent questions were superior to convergent questions in NSS score; the evidence presented here concerns the *magnitude* of that superiority.

consistent style around the structured divergent approach. Finally, the wait-time findings show the importance of allowing students adequate time for information-processing. And they provide a further reason for avoiding certain types of questions: those which generate too much ambiguity and confusion, and thereby produce long, discussion-chilling, post-question silences.

IMPLICATIONS FOR BETTER TEACHING

These data begin to establish a firmer footing for teaching improvement recommendations. Some key points are:

1) It is probably better to avoid ambiguity about convergence or divergence. Ask questions in such a way that you communicate *clearly* that multiple answers are acceptable or discussible, and follow this up with prompting and encouragement which invites second and third responses to the same question.

2) Familiarize yourself with Bloom's (1956) taxonomy and sprinkle your discussions liberally with questions requiring higher levels of thought: analysis, synthesis and evaluation.

3) Don't let complexity run rampant. The data show us two forms of cognitive overload, the Shotgun and Funnel Questions, which seem to produce confusion and withdrawal. In addition to the "brain cell torque" (to borrow a student phrase) needed to grasp the question, search one's storage-retrieval system and verbalize an answer, students must sort through mixed messages and overcome the fear of seeming foolish by answering the wrong question. This leads to poorer mileage and longer wait-times; so state your question crisply and take your chances with it.

4) Bear in mind that students may read your question in terms of the kind of thinking expected, and that it is confusing and inhibiting to switch levels drastically. One reason for this may be motivational as well as cognitive: people who are at one moment confronted with a challenging, adult issue, and at the next with a Micky Mouse, "When was the war of 1812?", type of question may oscillate psychologically, advancing a few years at one moment and regressing to high school the next. For students at a late-adolescent transition point, whose adult identities are still very fluid (Erikson, 1960), this process can be unsettling.

5) Don't be seduced by "pure" freedom. Many instructors, when they first decide to break away from the rigidity of low-level convergent questions, swing to the other extreme and give students a gigantic, formless playground in which to get lost. General Invitations do produce a bit more interaction than Quiz Show Questions, but they are not consistently effective. Probably they are most likely to succeed

must absorb two or more versions of the question before responding. The data indicate that these questions function somewhaat better than Single Questions in terms of student response. For example:

> Well, does Kafka like religion? . . . Is our impression that Kafka's favorable to the development of Christianity? Are we meant to clap our hands and praise Christianity after reading this story?

10. The Shotgun Question (NSS = 2.50)

5. So, we're talking about the fact that everybody's roles are changing, how—we've mentioned religion and education, how did religion and education during this period affect these changes, or how did the changes affect the kind of religion and education people had? . . . Let's start with religion . . . have women always had a sort of divine place in religion?

D. How do you interpret what the narrator tells you about the hero? What do you make of his return from Law school? Why did he decide he didn't really expect too much?

The Shotgun Question is multiple, in the sense of containing more than one question-sentence, and inconsistent. It may contain two or more separate content areas; it may embody multiple levels of thinking, as when an instructor simultaneously asks for complex analysis and factual information; or it may involve both of these inconsistencies. We call these Shotgun Questions because they often seem to be fired off in the hope that at least one fragment will hit the target. The instructor may be hoping to provide "something for everyone," be insecure about the effectiveness of any given question, or simply be thinking aloud between the beginning and the end of the question. The result is often self-defeating, because students are puzzled about what is expected; they must sort through the various questions being asked, select among them, and/or reconcile incompatibilities before responding. This is apt to heighten anxiety and inhibit participation.

11. The Funnel Question (NSS = 2.18)

(Note: Question #5 under "The Shotgun Question" is also a Funnel Question.)

The Funnel Question is a Shotgun Question with some particular characteristics. It is consistent as far as topic is concerned, but inconsistent as to level of thinking; and the level of thinking moves

down the hierarchy of complexity. Thus it "funnels" students from a general playground into the narrow chute of a convergent question. Sometimes these questions tumble out all at once, and at other times there is a notable pause between each subquestion. In either case, it is as though the instructor began with a broad question—a General Invitation or a Structured Divergent Question—and, fearing that it would fall flat, rushed in with progressively more and more structure in hopes of pulling a response from students. At times one has the impression that the opening, general phase of a Funnel Question is really "fishing": the instructor has a quite specific answer in mind but hopes that students will find it on their own. Then, when this fails to occur, he or she keeps giving hints and eventually arrives at a convergent reformulation which conveys quite clearly the sort of response originally desired. Funnel Questions are about as effective as Shotgun Questions in eliciting student response: about two-thirds as efficient as Single or Multiple Consistent Questions.

Table 7 displays the question typology data in a form amenable to cross-validation. The pattern of mileage scores is quite similar when Samples 1 and 2 are compared, both in terms of absolute means and in terms of the ordering of question types. The latter point can be stated as a hypothesis: if there are consistent differences among question types in the degree of student response, the rank-

TABLE 7

MEAN NSS SCORES FOR ALL QUESTION TYPES

	Samples 1 & 2			Sample 2			Combined Sample		
Question Type	Mean	S.D.	N	Mean	S.D.	N	Mean	S.D.	N
Major Question Categories									
Playground Question	5.00	2.73	7	5.17	4.37	6	5.08	3.59	13
Brainstorm Question	5.57	7.41	7	4.33	3.13	9	4.88	5.47	16
Focal Question	4.43	3.94	14	4.00	3.59	7	4.29	3.83	21
General Invitation Question	2.75	3.83	4	2.00	0.00	1	2.60	3.90	5
Analytic Convergent Question	1.38	0.70	8	2.33	1.31	12	1.95	1.20	20
Low Level Divergent Question	1.67	1.08	6	2.14	1.61	7	1.92	1.41	13
Quiz Show Question	1.50	1.51	18	1.41	0.78	22	1.45	1.17	40
Supplementary Question Categories									
Multiple Consistent Question	3.30	3.60	23	4.42	3.58	12	3.69	2.67	35
Single Question	3.30	4.54	26	2.53	2.59	30	2.89	2.52	56
Shotgun Question	2.57	1.55	7	2.43	1.68	7	2.50	1.14	14
Funnel Question	2.22	1.83	9	2.13	2.03	8	2.18	1.36	17

1. *The Playground Question* (NSS = 5.08)[10]

6. Let's see if we can make any generalizations about the play as a whole, from the nature of the opening lines.
A. (Instructor reads a sentence from the novel under study): Well, that's a very rich sentence . . . there's a lot there . . . OK, what's there?

This type is so named because the instructor designates a specific intellectual sphere (the "playground") for discussion, and then gives students the widest possible latitude in approaching it. In effect, he or she is saying, "Here's a playground that should be interesting; let's all stay here so we can do something in common, but you go ahead and choose the games." This "playground" may be a poem, a character in a play, a philosophical concept, "the opening lines," a sentence, etc., and is circumscribed quite carefully in light of the instructor's teaching goals. The hallmark of the Playground Question is the invitation to explore, which is typically offered with such phrases as "How do you interpret . . . ?", "What can you draw from . . . ?", "What are the possible meanings of . . . ?", and the like. This sort of invitation leaves open the question of what concept, category, or theme the student will use to make sense of the raw material offered.

2. *The Brainstorm Question* (NSS = 4.88)

4. What possibilities are there for refuge in *A Farewell to Arms*?
G. What kinds of things is Hamlet questioning? Not just in his soliloquy, but broadly, throughout the whole play?

This name is borrowed from the brainstorming technique, in which any and all ideas or solutions are sought in response to a specific question or problem. In contrast to the Playground Question, the primary structure is *thematic*, with the subject matter being less tightly focused. Thus, in the two examples from the exercise, the entire novel or play is within the scope of the question, but the issue to be addressed is quite specific ("refuge" in one instance, and "questioning" in the other).

[10] As indicated earlier, the NSS score refers to the average number of separate student comments following questions falling into a given category.

3. *The Focal Question* (NSS = 4.29)

1. So where is this wild boy better off? In the forest where he started, or in civilization being socialized?
E. Is Ivan Illych a victim of society, or did he create his problems by his own choices?

The Focal Question is oriented around an issue which calls for decisions. The instructor poses a limited number of alternatives (usually two or three, and rarely more than five) and asks students to take stands which they are to justify during the discussion. These stands call for higher order thinking and require students to marshal various kinds of information to support their views. There is a tone of debate or persuasion generated by Focal Questions; while the groups as a whole may or or may not reach a single conclusion, the effort to exert influence toward this end affects the shape of the discussion. Thus, the question's structure comes from the alternatives posed, while openness stems from the fact that students may draw on a wide range of information to support their views. For example, in the "Wild Child" film discussion (Exercise Question 1), students could (and did) bring in everything from ethics to developmental psychology to back up their positions.

The three questions discussed so far, the Playground, Focal and Brainstorm Questions, are clearly the most productive quantitatively and share characteristics which suggest the label "Structured Divergent Questions." Each in its own way provides a clear focus which prevents confusion and wandering (a delimited chunk of the material, a thematic thread, or a set of debatable alternatives); at the same time they encourage students to express a variety of ideas and to marshal many kinds of information in order to pursue the discussion. By contrast, the question types to be described below seem to miss this synthesis, by being either too open and vague or too narrow and mechanical.

4. *The General Invitation* (NSS = 2.60)

3. What about the lecture?
B. Any comments on Plato?

This question is a Playground Question that has lost its boundwait-time than low level questions, but the difference does not even approach statistical significance. Morever, this finding is mirrored when the data are broken down into subcategories defined by the other variables (convergent versus divergent and consistent versus inconsistent). All differences are small and inconsistent in direction.

Hypothesis IIIb compares structured and unfocused questions, predicting longer wait-times for the latter. The approach taken in testing this hypothesis follows the same line of reasoning described in connection with the test of Hypothesis IIb: that since unfocused questions are divergent and ambiguous with respect to intellectual level, they should be compared with structured divergent questions at both levels of complexity; and that since all the unfocused questions in the sample were consistent in content, the comparison group should be made up of consistent questions only. The results of both these comparisons support Hypothesis IIIb. As Table 5 shows, the shortest mean wait-time occurs with the structured, high-level questions, and the structured low-level questions are not too dissimlar. Furthermore, both comparisons are statistically significant, although the small number of unfocused questions must make this conclusion a cautious one.

Finally, the data in Table 5 test Hypothesis IIIc, which states that inconsistent questions will be associated with longer wait-times. Like the NSS scores, the wait-time data show marked differences among single, multiple-consistent, and multiple-inconsistent questions, and so these three groupings are analyzed separately. From this it is clear that the inconsistent group has the longest mean wait-time (significantly different from the other two), which supports Hypothesis IIIc.[8] Moreover, the ordering of groups is constant across Samples 1 and 2, and when examined separately within the convergent, divergent and high level subcategories. There is a reversal of order between the single and inconsistent groups among low level questions, but by and large it can be said that the pattern is a general one.

There is still, however, the puzzling finding that single-consistent questions have noticeably longer wait-times than multiple-consistent questions, to a statistically significant degree. This mirrors the pattern found with NSS scores (Hypothesis IIc), and such a similarity across two quite different indices of student response is worth some thought. It has led me to an ambiguity in the concept of wait-time, one which leads in turn to a broader and more productive view of question response.

If we treat "wait-time" as literally that, an empty period in which the teacher awaits a reply, then it is reasonable to simply count the number of seconds between the end of the question and the beginning of the first response. But, as already mentioned, this empty space must be filled with considerable information processing before a reply can be made; "wait-time" is really "think-time" from the student perspective. This point bears on the single versus multiple distinction in an important way: for while wait-time begins when the teacher stops talking, *think-time begins after the first complete statement of the question*. During the repetitions or paraphrases used in multiple-consistent questions, the student can be checking his or her understanding of the question, enriching his or her appreciation of its meaning, and formulating a response. In other words, the answer is being prepared while the instructor is still speaking, and it is really not surprising that students leap in to talk so quickly when the opportunity arises. By contrast, of course, single questions do not embody any think-time (other than wait-time); and multiple-inconsistent questions complicate its function by adding new questions which the student must digest and try to reconcile with the initial question.

Looked at in this way, the wait-time ranking of the three groups makes a good deal of sense. This further reinforces the assumption underlying Hypothesis III: that building in appropriate time for information processing is an important ingredient in creating good discussions. And it seems likely that the same or related factors could influence the extent of response as well—not as markedly as with wait-time, but enough to account for the marginal NSS score superiority of multiple-consistent as opposed to single questions. For example, having two or three alternate versions of the question to respond to may draw in more students; the opportunity for unpressured think-time during the question may stimulate greater response; and the positive climate created by the lack of a post-question pause may encourage hesitant individuals to speak up.

The data pertinent to Hypothesis III make the assumption of a

[8] This analysis includes only question-types which are represented in all of the three categories of Table 5. Unfocused questions (N = 4) and lower-level divergent questions (N = 6) are not included.

TABLE 2

MEAN NSS SCORES FOR THREE SETS OF VARIABLES: CONVERGENT VS. DIVERGENT, HIGH VS. LOW LEVEL, AND CONSISTENT VS. INCONSISTENT

Variable	Mean	S.D.	N
1. Divergent	4.13	3.23	68
2. Convergent	1.67	1.16	60
Difference: 1–2	+2.46**		
3. High Level	3.90	3.95	70
4. Low Level	1.57	1.32	53
Difference: 3–4	+2.33**		
5. Consistent	3.16	3.63	94
6. Inconsistent	2.24	1.77	34
Difference: 5–6	+0.92*		

NOTE: The N for categories 3 and 4 combined is smaller than the other groupings because five questions had so little content that they were unclassifiable as to level. They are discussed later as "unfocused" questions.

* $P < .10$ (t-test, one-tailed)

** $P < .001$ (t-test, one tailed)

than convergent questions, we ask: is the direction of difference the same within the multiple question subgroup as it is within the single question subgroup? And is it the same within the high level subgroup as it is within the low level subgroup? In other words, do the results run consistently through the data, or are there reversals?

The general answer is that the patterns are quite consistent throughout the data. There are 12 direction-of-difference comparisons to be made using the cross-classification system just described, and if we do this separately for Samples 1 and 2, there are a total of 24. Of these, 22 are in the directions predicted by Hypotheses Ia-c. While the individual cells are not statistically significant due to the small N's involved, this shows that all three of the variables do have consistent effects. There are two reversals, however, such that in each case a direction in Sample 1 is reversed by that in Sample 2. Both of these are in the convergent subcategory, and the result is to weaken the differences involved in testing Hypotheses IIb and IIc. Among the divergent questions, however, both hypotheses are well supported, as Table 3 shows.

While the reasons for this variation in pattern are not wholly clear, it is possible to speculate that the convergent versus divergent variable sets a *context* for the operation of the other qualities. Perhaps asking a high level question in a "one answer" frame inhibits the richness of response that its intellectual complexity implies, and thus overshadows the differential impact of high versus low level

TABLE 3

COMPARISN OF MEAN NSS SCORES FOR TWO PAIRS OF VARIABLES (HIGH VS. LOW INTELLECTUAL LEVEL, AND CONSISTENT VS. INCONSISTENT) WITHIN THE DIVERGENT SUBCATEGORY

Variable	Mean	S.D.	N
1. High Level	4.68	4.38	50
2. Low Level	1.92	1.73	13
Difference: 1–2	+2.95**		
3. Consistent	4.70	4.38	47
4. Inconsistent	2.74	1.35	19
Difference 3–4	+1.96*		

* $P < .05$ (t-test, one-tailed)

** $P < .025$ (t-test, one-tailed)

structuring. And it may be that the confusion engendered by an inconsistent question is kept in check by the specificity provided by a convergent focus. Conversely, the openness of the convergent questions would allow more free play for the other factors to exert influence on discussion response.

Another type of analysis helps clarify the picture with respect to Hypothesis IIc. Although the consistent questions are presented as one group in Tables 2 and 3, this group is actually composed of two subtypes: single questions, consisting of only one sentence; and multiple questions, which contain two or more repetitions or paraphrases of the same question. Thus there are two variables, not one, involved: the consistent versus inconsistent distinction referred to in Hypothesis IIc, and the distinction between single-sentence and multiple-sentence questions. We have, as a result, three subtypes of questions; there are, of course, no single-inconsistent questions since there must be more than one question-sentence for inconsistency to arise.

By comparing the multiple-consistent and multiple-inconsistent questions, we can control for the single versus multiple variable, and thus provide a clearer test of Hypothesis IIc. In contrast to the data in Table 2, the difference here is quite sharp. The mean for multiple-consistent questions is 3.69 (S.D. = 2.67, N = 35), which is significantly higher (P = .01 by t-test) than the mean of 2.24 presented in Table 2 for (multiple) inconsistent questions. The mean NSS score for single-consistent questions is between the other two groups at 2.89 (S.D. = 2.52, N = 56). It is higher than that of the multiple-inconsistent group, and the statistical significance level is about .10. Thus, by breaking the consistent group into its two comber of instances, following a given question, in which one student verbalization is followed immediately by another, without an intervening verbalization by the instructor.

d) Student Talk Time (STT): The duration, in number of seconds, of all student talk following a given question.

e) Percent Student Talk (Versus Teacher Talk) (%S): The number of separate student verbalizations divided by the total number of verbalizations (by instructor and students), which follows a given question.

It is predicted that these five indices will be positively intercorrelated, thus permitting the formulation of a single unified index of discussion participation.

Hypothesis II:

That certain question characteristics will be associated with higher or lower levels of discussion participation.

a) That divergent questions will yield greater participation than convergent questions.

b) That questions calling for higher levels of thinking will produce more participation than lower-level questions.

c) That clear, consistent questions which raise a single point and call for one type of thinking from students will lead to greater amounts of discussion than questions which have inconsistent subparts and thus set contradictory expectations.

d) That questions which provide an orientation or focus for student responses will be more effective in generating discussion than questions that are lacking in guidelines or direction.

Hypothesis III:

That questions requiring considerable interpretation and/or complex thought will be followed by longer "wait-time" pauses than simpler, more obvious questions.

a) That questions involving higher levels of thinking will be followed by longer pauses than those involving lower levels of thinking.

b) That vague, unstructured questions will be followed by longer pauses than those which provide structure and orientation.

c) That questions which have inconsistent subparts will be followed by longer pauses than clear, straightforward questions.

Hypotheses IIIa-c duplicate the variables thought likely to affect discussion mileage (as presented in Hypothesis II), with one exception: the convergent versus divergent dimension. The omission is deliberate, because the basis of prediction concerning this variable is not clear. One might expect divergent questions to lead to longer wait-times because there are more possible answers to consider. On the other hand, convergent questions (at an equivalent level of complexity) require more careful evaluation of the *correctness* of the response; and there may also be more inhibiting anxiety about being wrong. These reciprocally operating factors seem most likely to cancel each other out.

Hypothesis IV:

This is more a heuristic than a hypothesis in the strict sense. It arises from the impression that beyond the single variables enumerated above, questions can be classified by a more encompassing, holistic taxonomy. That is, certain recurrent themes may yield distinct question "species," through combining single characteristics into more complex patterns. I pursued Hypothesis IV, not by a direct quantitative test, but by an impressionistic scanning of the data followed by the gradual crystallization of impressions; these were then formalized, cross-validated, and checked against discussion quality criteria.

PROCEDURES

The research was based on seven 45-minute videotapes of T.A.-led discussion sections in the humanities disciplines (Literature, History and Humanities) at the University of California, San Diego. Most of the students in these classes were freshmen, and the T.A.s had varying degrees of experience in teaching (ranging from 2 to 6 academic quarters).

On each tape, two successive 10-minute segments were marked off and designated as Samples 1 and 2. Sample 1 was divided into subsamples 1a and 1b, and Sample 1a was used to conduct exploratory, inductive analysis. Sample 1b served to further refine and test these procedures, and Sample 2 served as a final cross-validation of the coding system. In the data tables below, Samples 1 and 2 combined; any discrepancies between the two samples are noted in the text. All categorizations, other than purely mechanical counting (which involved primarily the discussion quality and wait-time measurements), were written up in a formal coding manual which was in turn applied by a second rater. After preliminary discussions

It is useful at this point to reformulate our main objective as a pair of hypotheses or questions: *Does* the verbal form given to instructor-initiated questions influence student response, independent of other variables and the context of the class as a whole? And if it does, *what* question forms are the most fruitful? To assemble evidence on these points, two things are required: first, a clear, measurable definition of "good discussions"; and second, a way of classifying questions in order to distinguish what kinds have what effects. The procedure I used was to first assess the quality of discussions, recorded on videotape, from a number of classes in the humanities disciplines; and then to examine the questions which preceded the "good" and "poor" discussions in search of patterns and common themes.

But how can one measure discussion quality? The approach I took was to concentrate on the nature of interaction among instructor and students, an emphasis which corresponds well with what many instructors tell me they want in their discussions. A good discussion is one in which each point raised elicits a variety of student responses. In such discussions, a large portion of the class is active, rather than only a minority. And good discussions have momentum —students will continue interacting for some time without the need of further questions or prompting. Finally, in good discussions students engage each other in conversation; they have abandoned the pattern in which each student contribution must be followed by a comment of some kind from the instructor. These various characteristics, all reflecting the extent of student response, can be encompassed by the inelegant but serviceable term, "mileage."[3]

THEORETICAL CONSIDERATIONS

In searching for the ingredients of good discussions, I was guided by several hunches and hypotheses. The distinction drawn by Guilford (1958) between convergent and divergent thinking has seemed to me a useful guideline for forming discussion questions. Divergent-

[3] Of course, the *caveat* should be added here that because of the many differences among instructors' goals, values, and personal styles, this is not everyone's idea of a good discussion. Some may not feel the interactive pattern decribed here is a fruitful one. And this method says nothing about the content of questions or the quality of thought expressed by the participating students. The purpose of this research is to draw, as clearly as possible, a set of means-end relationships which can be used by instructors to reach certain goals when these are considered desirable.

thinking questions, for which there can be a number of "correct" or discussible answers, are more likely to generate continuing discussion. With convergent questions, by contrast, participation may be inhibited because of fears about being wrong, and the dialogue will probably grind to a halt when "the" answer sought by the instructor is produced. For example, a question like:

> How else would you end this story, if you couldn't have him reach some sort of point of divine revelation . . . all we're allowed to do is rewrite the ending. Tolstoy won't let us rewrite the rest.

seems more likely to generate extended discussion than:

> What's the main message here? What's Hemingway trying to tell us in this story?

Higher level questions,[4] characterized by the intellectual operations of analysis, synthesis, or evaluation, seem more likely to produce extended discussion as students become engaged with the complex issues involved. By contrast, lower level questions, drawing on memorization, comprehension, or application, can generally be dealt with more briefly. For example:

> What about the function of Laertes . . . what is his role . . . for example, could Laertes be left out of the work and would it still be *Hamlet*?

will probably be explored at some length, while this question:

> Okay, men move out to work . . . where did they used to work?

requires only a short answer.[5]

Another lead emerged from the frustrations of trying to decipher the impact of certain questions, in both consulting and research contexts. When one listens closely to classroom questions, or even better sees them transcribed (and thus compares them with written expression), it is immediately, staggeringly obvious that many of them are very poorly constructed and phrased. Often the instructor

[4] This conceptualization is drawn from the taxonomy developed by Bloom *et al* (1956).

[5] The two distinctions drawn so far (convergent versus divergent and higher versus lower level) may seem to overlap considerably. Yet I believe that they are conceptually distinct; the comprehensive typology to be presented later in this article will provide examples of all four combinations of the two dimensions.

The pages that follow are templates for use with the exercise in class. See discussion in Part 4.

1. So where is this wild boy better off? In the forest where he started, or in civilization being socialized?

2. So in this story, when's the point of truth for Kurtz?

3. What about the lecture?

4. What possibilities are there for refuge in *A Farewell to Arms*?

5. So we’re talking about the fact that everybody’s roles are changing, how—we’ve mentioned religion and education, how did religion and education during this period affect these changes, or how did the changes affect the kind of religion and education people had? … Let’s start with religion … Have women always had a divine place in religion?

6. Let's see if we can make any generalizations about the play as a whole, from the nature of the opening lines.

7. What was the name of that institution?

A. (Instructor reads a sentence from the novel under study): Well that's a very rich sentence … there's a lot there… OK, what's there?

B. Any comments on Pluto?

C. He talks about envying one character. Who was it?

D. How do you interpret what the narrator tells you about the hero? What do you make of his return from law school? Why did he decide he didn't really expect too much?

E. Is Ivan Illyich a victim of society, or did he create his problems by his own choices?

F. What was the most important reason for the revolution's failure?

G. What kinds of things is Hamlet questioning? Not just in his soliloquy, but broadly, throughout the whole play?

Appendix B. Assessment Article

COOPERATIVE LEARNING FOR HIGHER EDUCATION FACULTY

Barbara J. Millis
Philip G. Cottell, Jr.

AMERICAN COUNCIL ON EDUCATION
Series on Higher Education
ORYX PRESS

shoulders." The discussion should also be action-oriented with an emphasis on: "Here are the data. Here is how we can interpret them. Now, what can we do to strengthen the learning in this class?" The observer might say, for instance, "I could not read your diagrams on the board from my seat at the back of the class. Students seemed to have the same problem because they continually squinted, peered at each others' notes, and whispered among themselves. Next time, could you draw larger figures, distribute the critical information in a handout, or use an overhead projector?"

Successful classroom observations should accomplish at least two goals: (1) they should reinforce positive behaviors (things that the instructor is doing right) and (2) they should lead to changes in behaviors to strengthen teaching (things the instructor could improve). Thus, a skilled observer will both offer information and provide inspiration. Based on the observation, the instructor should know what to change and be motivated to extend the effort needed to change. It is helpful to record the teaching improvement activities that emerge from the post-visit discussion. This activity might also be considered a "teaching action plan." Followup visits to determine the effectiveness of the changes are extremely important because they place the visit within the context of a continuous improvement process.

SMALL-GROUP INSTRUCTIONAL DIAGNOSIS

A whole-class interviewing technique called Small-Group Instructional Diagnosis can gather consensus-based student data that enable instructors to make informed decisions about teaching/learning changes. This technique is based on research conducted by Joseph Clark (1982) when he served as a project director for FIPSE (Fund for the Improvement of Postsecondary Education) at the University of Washington, Seattle. Because the in-class data collection is based on information gleaned from structured small-group work, the technique uses cooperative-learning methods to gather data that may be useful in strengthening classes predicated on cooperative learning. Practitioners such as Diamond (1988) and Nyquist and Wulff (1988) agree on the basic steps involved.

During the pre-interview meeting, the SGID facilitator collects course-specific information from the instructor. A checklist similar to one designed by Wulff (1996), helps ensure systematic data collection. The two parties discuss important course information and the SGID data-collection process; the instructor selects his or her preferred SGID Feedback Form (generic form appears as Figure 13.1).

During the thirty-minute in-class interview, usually conducted at the mid-term point, the facilitator introduces herself, explains the SGID process, and asks students to form groups of six to eight and to select a recorder. The students

(Cou

No. of

1. List the major strengths in the course. (\
explain briefly or give an example for eacl

Strengths

2. List changes that you would make in th
explain how these suggested changes coul

Changes

3. Optional Question

SGID Feedback Form

FIGURE 13.1

students then discuss the questions c
recorder writes down the points on w
process, which takes only 8–12 minute
comments of each group on a centra
everything from the board for later an
tasks. She asks students for clarificatio
and she seeks to determine whether the
raised. This can be accomplished by
agreement or disagreement with partic

In the next phase of SGID, the
material to make it meaningful to t
arranged, for example, in order of fr
"Things to Continue;" "Things to (
Suggestions." The facilitator should
themes to be shared with the instru
verbatim, particularly ones with pote
tive ones can give the flavor of the in

(Course)

No. of Participants in Group________________

1. List the major strengths in the course. (What is helping you learn?) Please explain briefly or give an example for each strength.

Strengths	Explanation/Example

2. List changes that you would make in the course to assist you in learning. Please explain how these suggested changes could be made.

Changes	Ways To Make Changes

3. Optional Question

SGID Feedback Form

FIGURE 13.1

students then discuss the questions on the SGID Feedback Form and the recorder writes down the points on which they reach consensus. After this process, which takes only 8–12 minutes, the facilitator or a student records the comments of each group on a central chalkboard. Another recorder copies everything from the board for later analysis. The facilitator has several major tasks. She asks students for clarification or amplification on ambiguous points, and she seeks to determine whether there is whole-class consensus on the issues raised. This can be accomplished by asking for a show of hands indicating agreement or disagreement with particular comments.

In the next phase of SGID, the facilitator analyzes and organizes the material to make it meaningful to the instructor. The comments can be arranged, for example, in order of frequency under the central headings of "Things to Continue;" "Things to Consider Changing;" and "Additional Suggestions." The facilitator should try to "chunk" data under common themes to be shared with the instructor. Not all comments are included verbatim, particularly ones with potentially hurtful phrasing, but representative ones can give the flavor of the interview.

The facilitator and the instructor then meet for a "debriefing" at which the interview process is described, the data are analyzed and discussed, and further action is planned. White (1991) describes the process as follows:

> Soon after the SGID, the facilitator meets with the instructor and summarizes the information, answering questions, explaining comments, and offering alternative interpretations of apparent contradictions. During this phase, the tasks are to support the instructor and faithfully represent student perspectives. Though the intent is not to persuade faculty to agree with the students, it is to highlight themes and explanations that integrate student and instructor perceptions. In this manner, the facilitator promotes reflection on the issues, not just on the words. The conversation then moves to strategies for change, and what the instructor can say when he or she returns to the class (pp. 7; 10).

Diamond (1988) suggests that these questions should guide the action plan:

1. "What happened in class to elicit this comment?"
2. "Is *x* within my control?"
3. "Is *x* important to student learning?"
4. "Can I reasonably change in regard to *x*?"
5. "Do I want to change in regard to *x*?"
6. "What action can I take now?—the next time I teach the class?" (p. 92).

During the final phase, the instructor must discuss the results of the SGID with the students. As with the Classroom Assessment Techniques advocated by Angelo and Cross (1993), the SGID process raises student expectations for positive changes. Thus, it is essential that faculty members offer explanations of those course elements they are willing and able to address. Wulff (1994) describes, for example, an instructor who reduced the workload in response to student input from an SGID. During the followup with students, she thanked them for their honest comments and announced that she would try to make the workload more manageable by previewing assigned readings, eliminating overlapping readings, and assigning only two major readings in the week that papers were due.

STUDENT FOCUS GROUPS

Focus groups are a common marketing tool used by corporations to gather data about new products, customer satisfaction, and a host of other issues. They tend to be action-oriented in that clients will typically base decisions on data they feel are representative and reliable. Such groups involve highly paid professional moderators; expensive facilities, including one-way mirrors and sophisticated recording equipment; and complex data analysis that takes into account both nonverbal and verbal responses. Focus groups conducted by

professionals in higher education hav
fundraising issues (Topor, 1996).
However, focus groups can also be c
skilled nonprofessionals trained to gath
courses or programs of study. Like ot
outside facilitator, focus groups involve
interview conference, the interview itse
Course focus groups are similar to SC
conducting structured student intervie
differences. For one thing, the SGIDs a
regular classroom hours, so it is likely
Focus groups rely on volunteers, so th
attracting students who particularly lik
focus groups are often drawn from hetero
there is greater opportunity for candid i
relationships in ongoing classes. Wright a
discovered focus groups to be an invalua
tions within a university setting. They fou
particularly when the facilitators were abl
that could not have been obtained from t
It is useful to have two facilitators: on
and another to record the responses. Ta
the data capture, but, even if the session i
facilitator to record at least the openin
directives given by the processing facilit
clues provided by students (frowning, shru
transcription. Having students number of
their number helps the transcriber, but m
ing" assures students that their names will
To begin the session, it is important that
session establish the ground rules for the d
specific rights which should be clarified:
scope of the focus-group interview; they
videotaping; they must be assured of confi
that they will be treated with dignity, co
important that the group agree to some nc
task participation, and constructive disagre
can include guarantees for equal participat
Effective focus-group facilitators will stru
discussing the objectives and possible que
design activities to maximize the data collec
will use a variety of techniques. For example

Appendix C. Before and After Assignment Articles

Parallel Computing II – Discussion 1

For Monday, March 15, 2004

Read the article “Parallel Programming and Interoperable MPI” in Dr. Dobb’s Journal, February 2004. Bring in one discussion question, i.e., what you would like the class to discuss about the topic of this article.

Before Assignment and Article

Parallel Programming with Interoperable MPI

Building portable applications for diverse systems

William L. George, John G. Hagedorn, and Judith E. Devaney

Modern computing centers typically provide users with a variety of computing resources, ranging from single-processor workstations to high-performance parallel computers. Increasingly, this mix also includes Beowulf class machines—clusters of commodity PCs configured to operate as parallel computers. The Message Passing Interface (MPI) library, which provides C and Fortran interfaces to routines for sending data (messages) between processors, was designed to implement portable applications for diverse systems such as these.

Although parallel computations are normally run on single parallel computers, there is often a need to harness the resources of multiple clusters and parallel computers, forming what we call "multiclusters" to perform a single computation. (For information on related technologies, see the accompanying text box entitled "Multicluster Environments.") This might be required, for example, for simulations that are too large to be performed on any available individual parallel machine. Interoperable MPI (IMPI) provides a means of accomplishing this with minimal effort on the part of application programmers. IMPI is a set of protocols—implemented within an MPI library—that let multiple MPI libraries cooperate, acting like a single MPI library for programs running on a multicluster. The IMPI protocol specification is available at http://impi.nist.gov/impi-report/index.html. In this article, we'll examine IMPI and provide examples of how it can be used. For background on parallel architectures and programming, see the accompanying text box entitled "Basic Parallel Architectures and Programming."

William and John are computer scientists and Judith is group leader in the Scientific Applications and Visualization Group at the National Institute of Standards and Technology. They can be contacted at william.george@nist.gov, john.hagedorn@nist.gov, and judith.devaney@nist.gov, respectively.

A Crash Course in MPI

For readers unfamiliar with message passing, we'll briefly describe some basics of this programming style using C and MPI. Assuming you are running a program using *P* processes, each process will be identified in calls to MPI by an integer rank from 0 to (*P*–1). Listing One, for instance, sends an integer from the lowest rank process to the highest rank process.

Once this program is compiled and linked to the MPI library (*-lmpi*), it can be executed by a command-line utility program provided with the MPI library. Often this utility is named "mpirun." Assuming our executable is named "program1" and *-np* is the command-line switch for specifying the number of MPI processes (this syntax varies between MPI implementations), the command line to run our program with eight MPI processes could look like: *mpirun -np 8 program1*.

In Listing One, the *MPI_Init* and *MPI_Finalize* calls are required in all MPI programs. No calls to MPI routines can be made before the call to *MPI_Init* or after the call to *MPI_Finalize*. To get the rank of the local process, you call *MPI_Comm_rank*. To get the total number of processes, call *MPI_Comm_size*.

In most MPI routines, an MPI communicator is a required parameter. A communicator describes a set of processes (including the assignment of ranks to those processes) and defines a separate communications context. A message sent using one communicator can only be received by a call using the same communicator. The predefined communicator MPI_COMM_WORLD simply includes all of the processes; however, subsets of MPI_COMM_WORLD are possible.

In Listing One, the communication is performed with the most basic MPI communications routines *MPI_Send* and *MPI_Recv*. The parameters to these routines describe the message to be sent/received (*message*, *count*, and *MPI_INT*), the rank of the destination process (*dst* for *MPI_Send*) or source process (*src* for *MPI_Recv*), an arbitrary tag value, and an MPI communicator for message matching. The status parameter to the *MPI_Recv* routine holds details of the message once it has been received.

according to the ranks given to the IMPI clients.

Some IMPI Usage Patterns

Now that we have described the basics of parallel message-passing programming with MPI and how to start an IMPI program, we now turn to how IMPI can be used to expand the power of MPI programs. There are several types of applications that we anticipate will use IMPI to great advantage, but most likely there

Basic Parallel Architectures and Programming

At the highest level, most parallel scientific programs can be characterized as either task parallel or data parallel. These terms describe how the program obtains parallelism.

A task-parallel program consists of multiple independent tasks that can be completed with little or no communication between the tasks. Typically, one process is in charge of assigning the tasks to the available processors and collecting the results. The SETI@home project is one example of this model of parallel processing (http://setiathome.ssl.berkeley.edu/). The amount of parallelism available in a task-parallel program increases as the number of independent tasks increases.

A data-parallel program typically operates on large multidimensional arrays with a main loop that updates these arrays once per iteration while converging toward a solution. Each iteration of the main loop completes identical, or nearly identical, calculations to update each of the elements in one or more of the large arrays. The amount of parallelism available in a data-parallel program increases with the size of the arrays. Data-parallel programs often require communication between the processing nodes during the update calculations. Therefore, the distribution of the data among the processing nodes is an important consideration when designing these programs.

A classification scheme also exists for describing hardware that supports parallel programs. This classification scheme focuses on one of the most important considerations in parallel programming—access to main memory by the processors. At the top level, a parallel machine can be described as either a distributed-memory machine or a shared-memory machine.

In a distributed-memory machine each processor has its own local main memory that is not directly accessible by other processors. Sharing data between processors is accomplished via message passing; that is, explicitly sending data from one processing node to another. These messages are sent over a network that connects the processors.

In a shared-memory machine, all processors have equal access to all available memory. Like distributed-memory machines, shared-memory machines also need an interconnection network; however, in this case the network connects the processors to the main memory. For best performance, applications must avoid contention on this network by avoiding multiple simultaneous requests for data stored in the same area of memory. Even on a shared-memory machine, message-passing style programs, using MPI, perform well. In this case, each process only directly accesses portions of the memory that hold its data. Message passing is implemented (within MPI) using standard memory-to-memory moves. Unlike most multithreaded shared-memory applications, this results in very scalable parallel applications due to the low contention on the processor-to-memory interconnection network.

Of course in the real world, these classification schemes are blurred. Currently, many high-performance parallel machines available from manufacturers such as IBM, Sun, Hewlett-Packard, and SGI are distributed-memory machines with high-speed interconnection networks, in which each processing node in the network is a 2- to 16-processor shared-memory multiprocessor. Additional hardware and software is sometimes available for these machines that provide you with a shared-memory API on top of the basic distributed-memory architecture. Luckily, with vendor-tuned MPI libraries, all of these machines can still run your C/Fortran MPI programs without change and with good communications performance.

For an introduction to parallel programming with MPI, including books and online tutorials, see http://www.lam-mpi.org/tutorials/ and http://www.ERC.MsState.Edu/labs/hpcl/projects/mpi/.

—W.L.G., J.G.H., J.E.D.

especially during initial development, can be of great help in debugging the code and in experimentally determining a set of reasonable simulation parameters. In this case, we can use IMPI to run two or three subprograms, all aware of each other and connected via MPI. These extra programs are used for monitoring and controlling the main simulation. This is akin to the model-view-controller (MVC) style of program, except that the coupling between the model/view/controller is much looser. With the size and computational complexity of the models (simulations), the time between view updates may be from minutes to hours or even longer. Figure 2 shows a configuration of IMPI clients for this type of parallel application. In Figure 2, MPI processes are colored to indicate the various values of the IMPI_CLIENT_COLOR attribute. Code similar to that shown for pipelined programs could be used to create MPI communicators for each of the distinct parts of the program (simulator, monitor, and controller), and then, assuming the simulator is a pipelined program, create the communicators for each of the stages of the pipeline. The outline for this type of IMPI application is as follows:

One client is assigned to the monitor, one to the controller, and three to the simulator. Each MPI process, represented in Figure 2 by a circle, has a value for the IMPI_CLIENT_COLOR attribute that is cached onto MPI_COMM_WORLD. These attribute values, which match the associated client numbers in Figure 2, are emphasized here by mapping each value to a separate color; see Listing Three.

As with the pipelined program case, communication between the viewer, the controller, and the simulator is enabled by creating special MPI communicators using the MPI routine *MPI_Intercomm_create*.

So, the model part of Listing Three contains the simulation that is to be run on one or more clusters. This part of the program can be a data-parallel or large-grain pipelined program, as previously described, or any other type of MPI program. It is also possible for this simulator to be a multithreaded program that runs on a large shared-memory machine that uses MPI only to communicate with the view and controller parts of the IMPI program.

The second program referenced in Listing Three is a monitor program (the view portion of MVC) that performs the following steps in a loop: accept image data from the simulation, possibly once every iteration of its main loop; render this data into a form suitable for the target display; and display the image, either on your workstation or other suitable device. If the simulation is not working as expected, you will know this as early as possible. To minimize the effect of this monitoring on the performance of the simulator, the communication between the simulation and the monitor can be reduced by decimating the image data or reducing the frequency of image updates.

There are several types of applications that will use IMPI to great advantage

The third program, if needed, allows for some amount of interactivity with the simulation, perhaps letting you modify the controlling parameters of the simulation or, more drastically, allowing you to kill or restart the simulation from within the main simulation. This control could also let you turn on/off the monitoring of the simulation as needed.

Conclusion

IMPI lets legacy MPI programs run unaltered on multiclusters consisting of two or more computing resources such as parallel machines, clusters, workstations, and PCs. Also, applications can be written specifically to run in such a multicluster, allowing greater control over various aspects of the application such as large-grain pipelining, load balancing, and file I/O. One major design advantage of IMPI over other available techniques to the problem of running on a multicluster is that IMPI uses the vendor-tuned MPI libraries for optimum communication within each parallel machine, while still allowing the unrestrained use of all of MPI.

The freely available MPI library LAM/MPI (http://www.lam-mpi.org/) supports IMPI. Full implementations of IMPI are available from Hewlett-Packard, MPI Software Technology, and Pallas GmbH (for Fujitsu). Other implementations of IMPI are anticipated in the future. Furthermore, the National Institute of Standards and Technology (NIST) IMPI test tool (http://impi.nist.gov/ImpiTT.html) lets you test IMPI implementations for conformance to the IMPI protocol standard. For a detailed background on MPI, see *Using MPI: Portable Parallel Programming with the Message Passing Interface*, by William Gropp, Ewing Lusk, and Anthony Skjellum (MIT Press, 1999).

Note: Certain commercial equipment, instruments, or materials are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

DDJ

Listing One

```
#include <mpi.h>
int main(int argc, char *argv[])
{
  int my_rank, src, dst, tag, message, nprocs, count;
  MPI_Status  status;
  count=1;
  tag=100;
  MPI_Init(&argc, &argv);
  MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
  MPI_Comm_rank(MPI_COMM_WORLD, &my_rank);

  src=0;

  dst=nprocs-1;
  if (my_rank == src) {
    message=42;
    MPI_Send(&message, count, MPI_INT, dst, tag, MPI_COMM_WORLD);
  }  else if (my_rank == dst) {
    MPI_Recv(&message, count, MPI_INT, src, tag, MPI_COMM_WORLD, &status);
  }
  MPI_Finalize();
  return  0;
}
```

Listing Two

```
int  *stage,  stat,  stage_rank;
MPI_Comm  stage_comm;

MPI_Attr_get(MPI_COMM_WORLD,  IMPI_CLIENT_COLOR,  &stage,  &stat);
MPI_Comm_split(MPI_COMM_WORLD,  *stage,  0,  &stage_comm);
MPI_Comm_rank(stage_comm,  &stage_rank);
```

Listing Three

```
int *color, stat, rank;
MPI_Comm comm;

MPI_Attr_get(MPI_COMM_WORLD, IMPI_CLIENT_COLOR, &color, &stat);
if (color > 1) color=2; /* Simulator gets all clients > 1 */
MPI_Comm_split(MPI_COMM_WORLD, *color, 0, &comm);
switch (color)  {
case 0: /* Call the Controller */ break;
case 1: /* Call the Monitor    */ break;
case 2: /* Call the Simulator  */ break;
}
```

DDJ

After Assignment and Article

Parallel Computing II – Discussion 2

For Monday, March 22, 2004

Read the article "Faster Image Processing and OpenMP" in Dr. Dobb's Journal, March 2004. Bring in one discussion question, i.e., what you would like the class to discuss about the topic of this article.

Faster Image Processing with OpenMP

Multithreading, multiprocessors, and higher performance

Henry A. Gabb and Bill Magro

Processing digital images typically involves several filtering steps, some of which are time-consuming. In this article, we use OpenMP-based tools to show how you can use multithreading to improve filter performance on multiprocessor systems and/or processors that support Hyper-Threading Technology. (Hyper-Threading is an Intel-developed technology that provides thread-level parallelism on each processor, allowing multiple threads of applications to run simultaneously on one processor; see http://www.intel.com/technology/hyperthread/.)

The first step involves selecting a suitable filter. What constitutes a suitable filter? The filter should be compute intensive. There's no point wasting time threading a computation with instant response time. The potential speedup should merit the cost of thread creation in terms of both system and programmer time.

An important and time-consuming operation is the "radial blur," in which each pixel is replaced by the average of neighboring pixels along an arc. In digital image processing, blurring is used to convey motion, soften edges, give the illusion of distance, or focus the viewer's attention on nonblurred image regions. Radial blur conveys a rocking motion or movement around a point in the image. In Figure 1, the focus is on the lighthouse, which is also the blur center. The blur algorithm achieves this effect by averaging the pixels on an arc intersecting the target pixel (Figure 2). Calculating the positions of the neighboring pixels involves a fair amount of trigonometric computations because the pixels lie on an arc. Consequently, this algorithm is computationally expensive, and the amount of work per pixel increases as the blur angle increases.

Selecting a Threading Method

Having selected a suitable filter for multithreading, which threading method should you use? There are basically two choices—explicit threading or OpenMP. The OpenMP standard (http://www.openmp.org/) is a specification for a portable implementation of shared memory parallelism in Fortran, C, and C++. The spec provides a set of compiler directives and runtime library routines that extend Fortran, C, and C++ to achieve shared memory parallelism. OpenMP language extensions include work-sharing constructs, data environment, and synchronization. The standard also includes a callable runtime library with accompanying environment variables. Explicit threading methods (POSIX and Win32 threads, for instance) are best suited to task-parallel problems, or those where each thread performs a different function. A computer game is a good example of a task-parallel system. The task of processing keyboard or joystick input could be assigned to one thread, audio processing to another, and graphics rendering to yet another. The threads could also be given different priorities. This generality is a key advantage of explicit threading methods but it comes at a price.

Explicitly threading an application is often time consuming and error prone. It is an invasive process that can require significant modifications to existing code. First, concurrent tasks must be encapsulated in functions that can be mapped to threads. Second, POSIX and Win32 thread tasks only accept one argument, typically requiring changes to function prototypes and new data structures. Third, applications that were not originally designed for concurrent execution can have serious thread-safety issues. Any global or static variable, for example, is a potential target of a race condition. Finally, porting explicitly threaded code between different

Henry and Bill work at the Intel Parallel Applications Center. They can be reached at henry.gabb@intel.com and bill.magro@intel.com, respectively.

from the origin instead of incrementing a placeholder; see Example 2(b). The offset is now a pure function of the loop indices. This makes the outer-loop iterations completely independent. They can be executed in any order without changing the final results. The revised pixel loop for the *RadialBlur* function is in Listing Two.

Time	Thread 0	Thread 1
T_0	x=0 y=0	
T_1		x=1 y=0
T_2	pixelOffset=f(1, 0)	
T_3		pixelOffset=f(1, 0)

***Table 2:** Storage conflicts can lead to incorrect results.*

Now it's time to begin threading the pixel calculations with OpenMP. Again, the outermost pixel loop contains the most work so this is where we'll put the OpenMP parallel region. Variables are shared among threads by default in OpenMP. This is satisfactory for many variables in the pixel loops. Those that are read but never written can be safely shared among threads (*filterRect, maskPixel,* for instance). It is also safe to share arrays that are indexed using the counter of the parallel loop. Provided *pixelY* is private, for example, threads will never write to the same pixel row (the index of a parallel loop is private by default in OpenMP). However, to guarantee correct results from parallel execution, threads must maintain private copies of some variables.

Analyzing and correctly classifying the variables in small parallel regions is not too difficult. Manually analyzing larger parallel regions, with nested function calls, however, can be time consuming and error prone. Fortunately the Intel Thread Checker (http://www.intel.com/software/products/threading/) automatically detects OpenMP errors. A parallel program must preserve data dependence constraints if it is to produce correct results. Violations of these constraints are called data races or storage conflicts. For example, consider this pseudocode:

```
for (y = 0; y < Height; y++)
{
  for (x = 0; x < Width; x++)
  {
    pixelOffset = f(x, y);
    ProcessPixel (pixelOffset);
  }
}
```

The offset to the current pixel is a function of the loop indices *x* and *y*. If multiple threads execute this code, it is easy to interleave statements that demonstrate how storage conflicts lead to incorrect results; see Table 2. The variables *x, y,* and *pixelOffset* are shared between Threads 0 and 1. At T_1, *x* and *y* are overwritten by Thread 1. As a result, both threads calculate the same offset and proceed to process the same pixel. Storage conflicts also occur on *pixelOffset*. Making these three variables private to each thread solves the problem.

To analyze the pixel loop for data dependencies, a single OpenMP parallel for pragma was placed before the outermost loop—no variables were explicitly classified as private. Next, the plug-in was compiled with the Intel C++ compiler using the *-Qopenmp* and *-Qtcheck* options.

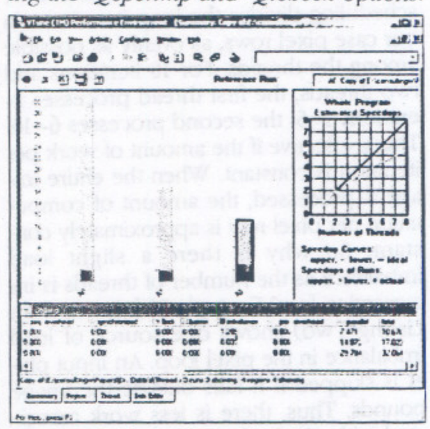

***Figure 3:** Thread Profiler showing performance improvement in the filter for one (A0), two (A1), and four (A2) threads. Green indicates time spent in the parallel region, light blue time spent in serial calculation, and red time lost to load imbalance.*

Consider Example 3(b), also taken from Listing Two. When editing a digital image, it is common to operate on sections of an image. These sections often have irregular shapes. There is no work associated with pixels outside the selected region (governed by the *maskPixel* array). Remember that static scheduling works best when the amount of work per loop iteration is approximately constant. The amount of work per pixel row is far from constant if the image section is triangular, for example. This creates a severe load imbalance when the image rows are divided statically among threads.

For this reason, OpenMP provides alternative schedule methods. When the amount of work per iteration varies, dynamic scheduling is often best. In a dynamically scheduled loop, iterations are assigned to threads as they become available to do work. Unlike explicit threading, OpenMP scheduling changes are easily accomplished via a schedule clause (Listing Two). Dynamic scheduling incurs slightly higher system overhead than static scheduling, but in the case of our filter, its effect on parallel performance is beneficial.

Conclusion

While a simple image filter might only merit threading for very large images, radial blurring is sufficiently compute intensive that even small images benefit from parallel computing. In fact, the implementation shown here uses an approximation to simplify the calculations. Blurring can be made more accurate, and consequently more compute intensive, using barycentric coordinates (see the accompanying text box entitled "Barycentric Coordinates") but this exercise is left to the reader. Most image processing methods are data parallel, so OpenMP is a fast and easy route to parallelism. It normally requires little code modification and is easy to debug and tune with existing programming tools.

DDJ

Listing One

```
void RadialBlur(void* data, int32 dataRowBytes, void* mask,
                                 int32 maskRowBytes, Rect* tileRect)
{
   uint8* pixel = (uint8*)data;   // Points to top left pixel of the image
   uint8* maskPixel = (uint8*)mask;
   uint16 rectHeight = (uint16)(tileRect->bottom - tileRect->top);
   uint16 rectWidth = (uint16)(tileRect->right - tileRect->left);
   // Variables for radial blur
   uint32 arc;        // Accumulate values on blur arc
   int n, step;       // Number of pixels on blur arc and increment
   int count;         // Number of arc pixels within rectangle
   float offset;      // Angular offset from blur pixel
   float theta;       // Angular increment of arc
   float angle;       // Angular sweep from rectangle center to blur pixel
   int R;             // Distance from rectangle center to rectangle corner
   int xr, yr, r;     // Distance from rectangle center to blur pixel
   int16 w, h;        // Coordinates of pixel at rectangle center
   float *ct, *st;    // Cosine and sine tables
   float xx, yy;      // Real coordinates of arc pixel
   int i;
   uint32 pixelOffset;
   const double PI = 3.141592654;
   Rect* filterRect = &gFilterRecord->filterRect;
   PSImagePlane plane;
   // Used to determine the addresses of pixels on blur arc
   plane.data = &gFilterRecord->inData;
   plane.bounds.top = filterRect->top;
   plane.bounds.bottom = filterRect->bottom;
   plane.bounds.left = filterRect->left;
   plane.bounds.right = filterRect->right;
   plane.rowBytes = gFilterRecord->outRowBytes;
   plane.colBytes = gFilterRecord->outColumnBytes;
   // Initialize radial blur variables
   // Find rectangle center (w,h)
   w = gFilterRecord->imageSize.h / 2;
   h = gFilterRecord->imageSize.v / 2;
   R = sqrt (w * w + h * h);
   angle = (float)gParams->angle / 180.0 * PI;   // Convert to radians
   n = 4 * angle * sqrt(R) + 2;
   theta = angle / ((float)(n - 1));
   // Setup cosine and sine tables
   if (((ct = (float *) malloc (n * sizeof(float))) == NULL ) ||
      ((st = (float *) malloc (n * sizeof(float))) == NULL ))
   {
      *gResult = memFullErr;
      return;
   }
   offset = theta * (n - 1) / 2;
   for (i = 0; i < n; ++i)
   {
      // Build sine and cosine tables
      ct[i] = cos(theta * i - offset);
      st[i] = sin(theta * i - offset);
   }
   // Loop over pixels in the image
   for(uint16 pixelY = 0; pixelY < rectHeight; pixelY++)
   {
      for(uint16 pixelX = 0; pixelX < rectWidth; pixelX++)
      {
         bool leaveItAlone = false;
         if (maskPixel != NULL && !(*maskPixel) && !gParams->ignoreSelection)
            leaveItAlone = true;
         if (!leaveItAlone)
         {
            // Find (x,y) coordinates of pixels on arc
            xr = pixelX - w;
            yr = pixelY - h;
            r = sqrt (xr * xr + yr * yr);
            if (r == 0)
               step = 1;
            else if ((step = R / r) == 0)
               step = 1;
            else if (step > (n - 1))
               step = n - 1;
            for (i = 0, count = 0, arc = 0; i < n; i += step)
            {
               xx = w + (float)xr * ct[i] - (float)yr * st[i];
               yy = h + (float)xr * st[i] + (float)yr * ct[i];
               if ((yy >= filterRect->bottom) || (yy < filterRect->top) ||
                   (xx < filterRect->left) || (xx >= filterRect->right))
                      continue;   // (x,y) outside rectangle
               ++count;
               // Convert (x,y) coordinate to pixel address
               pixelOffset = ((int)xx * gFilterRecord->outColumnBytes) +
                             ((int)yy * gFilterRecord->outRowBytes);
               arc += *((uint8 *)gFilterRecord->inData + pixelOffset);
            }
            if (count != 0)
               *pixel = (uint8)(arc / count);
         }
         pixel++;
         if (maskPixel != NULL)
            maskPixel++;
      }
      pixel += (dataRowBytes - rectWidth);
      if (maskPixel != NULL)
         maskPixel += (maskRowBytes - rectWidth);
   }
   // Clean-up sine and cosine tables
   free(ct);
   free(st);
}
```

Listing Two

```
#pragma omp parallel for schedule(dynamic) \
            private(pixelY, pixelX, pixelOffset, arcPixelOffset, \
                    i, count, arc, xx, yy, step, x, y, xr, yr, r)
   for (pixelY = 0; pixelY < rectHeight; pixelY++)
   {
      for (pixelX = 0; pixelX < rectWidth; pixelX++)
      {
         // Find (x,y) coordinates of current pixel
         x = gFilterRecord->inRect.left + pixelX;
         y = gFilterRecord->inRect.top + pixelY;

         pixelOffset = (x * gFilterRecord->outColumnBytes) +
                       (y * gFilterRecord->outRowBytes);
         bool leaveItAlone = false;
         if (maskPixel != NULL && !gParams->ignoreSelection &&
             !(*(maskPixel + pixelOffset)))
               leaveItAlone = true;
         if (!leaveItAlone)
         {
            xr = x - w;
            yr = y - h;
            r = sqrt (xr * xr + yr * yr);
            if (r == 0)
               step = 1;
            else if ((step = R / r) == 0)
               step = 1;
            else if (step > (n - 1))
               step = n - 1;
            for (i = 0, count = 0, arc = 0; i < n; i += step)
            {
               xx = w + (float)xr * ct[i] - (float)yr * st[i];
               yy = h + (float)xr * st[i] + (float)yr * ct[i];
               if ((yy >= filterRect->bottom) || (yy < filterRect->top) ||
                   (xx < filterRect->left) || (xx >= filterRect->right))
                      continue;   // (x,y) outside rectangle
               ++count;
               // Convert (x,y) coordinate to pixel address
               arcPixelOffset = ((int)xx * gFilterRecord->outColumnBytes) +
                                ((int)yy * gFilterRecord->outRowBytes);
               arc += *((uint8 *)gFilterRecord->inData + arcPixelOffset);
            }
            if (count != 0)
               *(pixel + pixelOffset) = (uint8)(arc / count);
         }
      }
   }
```

DDJ