

Course transformation case study

Carl Wieman – January 26, 2007

The following is a case study example of a course transformation as an illustration of many of the ideas listed in the course transformation template. This particular effort was carried out in a course for which there were obvious problems and where there were two faculty available to work on fixing it. There was also a postdoc who was doing research on student learning of quantum mechanics, the topic of this course. So this example was chosen because it illustrated a large variety of changes-- almost everything was examined and redone. Under other circumstances/courses, only some fraction of these different aspects might be addressed, and the solutions might be different.

The example course was Physics 2130 "Introduction to modern physics" (third semester physics course taken by engineering majors) at the University of Colorado. This transformation example was carried out by Carl Wieman and Kathy Perkins, who team-taught the course, and Sandra McKagan, a new postdoc who completed a Ph. D. in theoretical physics but had a strong interest in education. McKagan worked on this as a research project in physics education about how students learn quantum mechanics, and is in the process of writing up several articles for publication on different aspects of what her research connected with this course have revealed.

Step 1. Setting learning goals.

A. Examining the customers and their needs.

We first reviewed who takes this course and why. We learned that it is a required course for mechanical engineering majors and an elective for a few other engineering departments, with most students who elect to take it coming from electrical engineering and aerospace engineering. It is the terminal physics course for all of these students. The mechanical engineering students are roughly split between sophomores and juniors, and the students from other majors tend to be largely juniors and seniors. The course is taught via two 75 minute lectures per week, with no recitation or lab sections. (Changing these scheduling constraints was considered to be impractical to try and do without clear evidence it was necessary. Subsequent evidence shows that reasonable learning is possible within these constraints.) Enrollments vary between semesters but there were typically 125 students in the fall semester and 60-70 students in the spring semester. [These numbers jumped to nearly 200 per semester after our transformation efforts were put in place, and included several students planning to major in physics. This enrollment change somewhat complicated the course transformation process.]

We met with several relevant faculty from the three engineering departments whose students take the course (mech., electrical, and aerospace eng.) and gathered their ideas as to what students should get from this course. Condensing a number of emails and one hour-long meeting, a few consistent themes were clear. First, the course as previously taught was badly aligned with what was considered useful for the students to know. The material which they thought was not useful was special relativity and spending lots of time on how to solve various the differential equations that occur in the few model systems that are analytically solvable. What they said was most important was for their students to understand the quantum mechanical origins of basic chemical and material properties of matter. Also it was important for their students to understand the part that quantum mechanics plays in a number of modern technologies, particularly nanotechnology and electronics. These two emerged as clearly the most important educational goals of the course.

A discussion with physics faculty who had taught this class previously revealed that they also found special relativity difficult to fit into the course and of little interest or value to the students. So, we concluded that the course should drop special relativity and use the added time to concentrate on providing students with enough understanding of quantum mechanics that they could understand the basics of technological applications. As discussed below, in formulating more specific learning goals, we also decided that it made sense to reduce the large amount of time that had been traditionally spent on having students learn ways to solve the Schrodinger equation in a variety of special model cases.

1.B. Establishing learning goals for the transformed course.

In light of previous input, we (actually the postdoc) next examined the current way the course was being taught and what students actually were learning from it. The first part was done by examining the syllabus, lecture notes, homework and exams from previous semesters, talking to faculty who had taught the class, and sitting in on some lectures. The postdoc gathered data on student thinking and learning by setting up a "homework session," a time and place each week where students could come to work together on homework in the class and, as an incentive to students to come, get limited help from her as needed. These sessions resulted in providing large amounts of input on student thinking, revealing many basic misconceptions and confusions, as well as the profound difficulties most students were having in understanding the purpose, motivation, and structure of course. The postdoc also read research literature on teaching QM to collect what was known about student difficulties with the material. Finally, she collected syllabi from comparable courses at other institutions, looked at lists of topics covered in standard textbooks, and discussed with faculty who had taught the course in previous years as to what they saw as the most important things students should learn in this class.

Based on all this input, we came up with list of primary overarching course goals that we wanted students to be able to achieve from the course. These largely aligned with what engineering faculty said they wanted their students to get from class, but these course goals "operationalized" these ideas in terms of specific topics and what students should be able to do with these topics in terms of application, calculations, etc. They also included foundational topics that we decided were essential for students to master before they could effectively achieve the goals of applying ideas to technology etc. How these goals could be made consistent with an intellectually coherent structure and modestly consistent with the topics listed in the official university course description also played a substantial and small part respectively in the final decision. During this process we discovered that for most students, and even some faculty, the material covered in this course was perceived as a largely unconnected set of different topics. This lack of intellectual coherence made it difficult to learn. So we put a significant effort into figuring out how one could create a course that was intellectually much more coherent. (As an editorial aside, there are indications that we have achieved major improvements in this respect, but the students are still not completely connecting up all of the different ideas, so further work on this is needed.)

The subsequent list of goals was discussed further with faculty familiar with the course in order to converge on a list that was largely a consensus of viewpoints. It was largely but not perfectly a consensus as a few faculty had strong but idiosyncratic opinions about the importance of particular topics or skills.

2. Choice of presentation of material.

To address student beliefs and to be consistent with research on engagement and motivation, as well as the desire for engineering faculty to have the students relate the material to real life things, we found as many real world examples as possible of the ideas of the quantum mechanics that we wanted to present. We then presented the ideas in terms of explaining what was happening in these contexts, and how that could be generalized.

A major and quite successful effort was the creation of a number of interactive simulations that addressed particularly difficult ideas in effective novel ways, or could provide considerable illumination of fundamental ideas. Several of the simulations recreated virtual (and easier to understand) versions of several key experiments that established the fundamental basis for quantum mechanics. Others showed pedagogically valuable gedanken experiments. These were done as part of the PhET project and went through the standard cycle of development and testing to ensure student usability and pedagogical effectiveness. These simulations were then used extensively in lecture and on homework problems.

Some general features that were designed into the homework, lectures, and exams were:

- a. Make explicit connections of the subject to students' experiences in everyday life. This was based on research showing that connecting with students' experience is beneficial both in making the subject more interesting to them and in giving them background with which they can build their understanding, by providing them with a context. These features are important in any subject, but we found them to be particularly important in this class. It had a profound impact on student's views about the course compared to previous years.
- b. Emphasize active sense-making and reasoning about the material. Include in-class activities (e.g. clicker questions) that require students to engage in this process.
- c. Include a focus on developing a conceptual understanding of the science.
- d. Use analogies, illustrations, and visual aids (simulations) in class to help students develop visual models, build their understanding, and draw connections to everyday life. Students' inability to have visual models of what was being discussed or to see how it connected with real world phenomena, and the related difficulty of understanding the instructor's goals for the class, were frequently cited sources of frustration in previous years. There was a clear and dramatic change in students' views on all three of these points after the course was changed.

Step 3. Creating assessment tools.

We developed a variety of means to assess the results. First, we created a multiple choice test to measure students' mastery of a number of specific concepts of quantum mechanics covered by the course. The creation of the multiple choice quantum mechanics concept test was based on input from several sources: physics education research literature; interviews with students prior, during, and after instruction; formal observations by faculty, graduate, and undergraduate TAs of the student group homework problem solving sessions on student confusions and difficulties and problem solving; and examination of answers to open-ended questions on homework and exams. The resultant multiple choice test addressed central difficult concepts of quantum mechanics and provided choices for answers that reflected common student incorrect reasoning, often presented in the words of students.

We refined the test by having several faculty look at it to ensure that the test questions cover appropriate topics and see if important topics were missing. Based on this, we threw out a number of questions that covered material faculty felt were not central or where the answers

were not sufficiently clear. (Surprisingly, we found faculty to be somewhat fuzzy in their understanding of one or two topics, so we also threw out most questions on these topics, feeling that it was unreasonable to expect students to master them.) Next we interviewed students on the test questions to ensure that they interpreted the question and possible answers correctly.

This test was first given to the class at the end of the semester before changes were made. The test was also given to the students as they completed the simultaneously offered modern physics course for physics majors that covered many of the same topics. The test was given on the next to last day of class. Although it is ungraded, students were motivated to do their best because we explained that it helped us to see which topics we particularly needed to review on last day of class to help students prepare for final exam, and it would help guide their studying for the exam because these questions illustrated what we felt were essential concepts. We have repeated this practice of giving the test on the next to last day of class in subsequent semesters as we changed the class. We now also give the exam to students on the first day of class to evaluate their incoming knowledge. We explain that they should not expect to know many of the answers, but our purpose is to gauge their level of background knowledge to better tailor the class to meet their needs.

This exam addressed understanding of concepts that were considered important by a consensus of faculty. As often happens with such exams, the student results revealed some shocking deficiencies in student understanding. In the first semester we gave the test, which was before we started changing the class, it confirmed one very widespread confusion that we had suspected from our interviews and observations, namely most students did not understand the idea of potential energy as used in the course. This basic confusion which was never realized or addressed by the instructors in previous years meant that it was impossible for students in this situation to understand the last 2/3 of the course material, and so they abandoned all attempts at that, and merely memorized algorithms. (We found similar though not quite a pervasive a confusion among the students taking the introductory quantum course for physics majors.) The results from this test factored heavily in our detailed design of the transformed course.

There are other learning goals that were not measured by this conceptual test, such as ability to do various types of quantitative calculations, justify the fundamental ideas of quantum mechanics in terms of empirical observations, and to use quantum mechanical concepts to explain and predict behavior of real world situations and devices. These are assessed through the exams and homework problems we created for the course. We designed homework and exam problems that were intended to do that, and now having seen student performance on those items, and having done sampling interviewing of students on their understanding, we have a pretty good idea as to which of these are and are not valid assessment questions.

Another important assessment tool is the weekly observational reports provided by the TAs on the problem solving sessions discussed below. We also collect reports of their observations of lectures and student behavior (interest, confusion, understanding,..) during lecture, gleaned from sitting in the midst of the regular students. Finally, we have a weekly online survey question that gauges student opinion and perceptions of mastery and confusion on specific topics.

After the first iteration of the course, we made a number of adjustments to our detailed learning goals and topics covered. As we progressed through the actual teaching of the revised course for the first time, we broke down the overarching course goals into detailed specific goals, which basically laid out in detail what we hoped students would be able to do after the completion of this course. In the process of teaching the course and examining the results, we

realized that there were a few topics that, in terms of their intellectual demands on students, did not fit in, in that they required an excessive amount of student time and effort to master relative to the importance of the material. For some of these topics we subsequently modified the course structure to better address the student difficulties we discovered, which made these topics easier to master. For other topics we dropped them because they did not seem important enough to justify the investment of time we now realized was required for useful mastery.

We had to make one major modification in our goals and subsequent course structure after the first semester of transformation. The first design of the course was based on the assumption, guided by the course descriptions, that students had covered classical electromagnetic radiation thoroughly in the previous semester and understood it. This turned out not to be the case and consultation with instructors of the previous course indicated it was unlikely that it ever would be the case. So we modified our course accordingly.

Step 4. Evaluation of resources available and resource needs.

- A 200 student classroom equipped for clickers (projector, screen, computer, receivers)
- 2—1.25 hour lectures per week, no smaller sections possible.
- Schedule and staff for the problem solving sessions: When it was clear that the enrollments for this course were much higher than in the past (~ x 2-3), we negotiated for an additional half TA to help with grading of essay questions on homework. We recognized the value of undergraduate TAs from our previous courses, and negotiated to have 3 provided. These were used to staff the homework sessions and facilitate in-class discussions, as well as provide regular feedback on both.
- In order to have weekly homework, we went to a computer graded system, that allowed a variety of quantitative and multiple choice answers, including checking "all that apply" out of a long list of possible items. This is the most useful format that we have found for having computer graded homework assignments that have a critical thinking component, particularly involving conceptual reasoning. Finally, we took advantage of the availability of the PhET project staff to develop interactive simulations on particular topics where such simulations were expected to be both suitable and effective.

Step 5. Structure of the class.

Having established the general course goals, broad intellectual structure of the course, and assessment plans, we next worked out the detailed structure. This included deciding on grading policies, choice of textbook, lecture and homework format, class calendar, how the TAs were to be used, etc.

Grading: we operated on a philosophy that grading policies should reward students for doing things we thought were important for their learning. This means that we weight doing homework correctly relatively heavily compared to most instructors. The two hourly and final exams do still count a significant amount though. We also assign some small number of points to each of several other non-exam aspects of the course, such as reading the textbook before class, checking their past homework to see what they did right and wrong, attending class and responding to questions in class. Our surveys have shown that students are much more inclined to do something, such as reviewing their past homework and reading the textbook before class, if it is graded at a nominal level, even if it is an activity like these that they readily recognize as valuable to their learning. We also changed the grading to an absolute scale rather than a curve. Our past experience and that of others has indicated that grading on a

curve inhibits student-student collaboration and educationally valuable interaction. We believe that such social interaction is valuable to their learning, in part, because this is the only obvious way for students to practice and develop the important skills of critically analyzing scientific arguments of others and evaluating their own understanding and arguments ("metacognitive skills")

Collaboration on the homework was strongly encouraged, and our observations indicated that all but a few students avoided merely copying other student's homework or allowing theirs to be copied. We suspect that the explicit discussions about collaboration and expectations and connections between homework and success on exams, as well as the essay questions that had to be written individually, helped create this environment. There is little net impact on the actual grade distribution in switching from grading on a curve, because the absolute scale that we set is quite high, and we reserve the right to lower the scale later, and often do slightly. (So a B ends up being >77 or 78 instead of the 80% stated at the start of the course, for example.)

After deciding how the course would be organized and the material presented, we selected a textbook that we felt was the closest match. It had some recognized deficiencies including not being particularly readable and not matching the intellectual structure of the course particularly well, but it covered nearly all the topics and at roughly the same level of mathematical sophistication and in approximately the same order as we presented them in class. After one semester we realized our mistake. There were lots of complaints about the textbook from the students and indications from the problem solving sessions that it often confused rather than helped them. We then switched to a textbook that was considerably more readable and, not coincidentally, more closely matched the intellectual organization of the material we had decided upon, but did not cover a number of the topics we included, and it covered nearly all the topics at a less rigorous level of mathematics than we used. When surveyed, the students indicated that they found this a far better tradeoff. They routinely read and learned from this second textbook, and there were few if any complaints about what we had perceived as the deficiencies of the text.

Step 6. Preparing lectures.

Building on our past experience, we chose to use PowerPoint lectures generally built around a series of clicker questions for each lecture. These choices were based on our past data showing students retain very little from our lectures that do not involve peer-instruction enhanced by clicker use, and that the most memorable aspects of lectures are clicker questions, drawings, and animations, particularly those showing amusing but useful analogies. All of these are features that are facilitated by use of PowerPoint. Also, the PowerPoint slides are particularly easy to reuse and modify in subsequent semesters, greatly reducing preparation time for ourselves and subsequent instructors.

Based on successful results from past semesters in other courses, we had students all use individual clickers and they were randomly assigned seats with assigned consensus groups made up of 3-4 students in adjacent seats. The seat locations of the groups were switched front to back of the lecture hall, half way through the class. Students were assigned seats on the first day of class, with an explanation immediately given as to why we were using assigned seats (the established value of peer discussion, and advantage to having people that are always available and expected to talk with you and share and examine ideas; why this means assigned seats). Although there is some documented value to assigning groups based on a variety of more detailed criteria, we felt the benefit was not enough to justify the substantial added time and effort required.

We used clicker questions as discussed below.

a. Various formats for in-class questions:

- i. question posed → students discuss in groups → student vote as a group, i.e. consensus required. (Used most often later in the course.)
- ii. question posed → students discuss in groups → students vote individually
- iii. question posed → silent vote → students discuss in groups → students vote again (Requires more time, but can lead to charged discussion. Used more at beginning of course.)
- iv. question posed → silent vote (Used for very brief quizzes usually used to check that they had read the assigned material in advance of lecture, which greatly improved their ability to participate meaningfully in clicker question discussions. Occasionally we also used this format to test understanding of a specific point.) These types of questions were graded on correctness. The other types listed above were typically only graded on participation, with only an occasional question graded on correctness of answer. All indications were that students took the questions seriously nearly all of the time, whether or not they were graded on the correct answer in that well over 90% would get the correct answer on a fairly easy question, and obviously wrong answers were almost never selected.

b. Both the instructors would actively promote student-student discussions within the groups and listen in on student discussions to understand their confusions.

c. Undergraduate TAs attended class and promote student-student discussions during clicker questions by looking for groups that were either off topic or not talking during time consensus discussions, and encouraging them to do so by asking the members of the group directly about the question.

Lectures are designed around specific learning goals for the class, and usually are built around a set of clicker questions which the students normally discuss within their groups before answering. Frequently, in the followup to the clicker question we then call on students to give their reasons for their answer, before the answer is revealed. This allowed the reasoning, both correct and incorrect about the question to be addressed. We always follow all but the easiest clicker questions with a recap of what the correct answer is and why. We used to not do this when a large fraction of the class got the questions correct, but surveys indicated that, on most of our questions, even when students got it correct, they were not completely certain of their answer and wanted us to go over the correct answer. On more challenging questions where a substantial fraction of the class got it wrong, there were invariably follow-up questions from students. A single clicker question could generate as much as 10-15 minutes of student questions and follow up discussion.

We post our PowerPoint slides on the class website immediately after class is over, so students can use these for notes rather than writing down everything presented in class themselves. It is essential that they not have to concentrate entirely on taking notes, if they are to be able to concentrate on the ideas and the scientific discussions with their peers triggered by the clicker questions.

The well-known danger of using PowerPoint slides is that one can go through material far too fast for students to follow, but this is avoided by having frequent clicker questions. We find

that 14 not-very-dense slides, and 4-5 substantial clicker questions (possibly with some additional brief questions, such as the quizzes on assigned reading) is fairly typical for a one hour class. An unusual problem in this course was the number and depth of the student questions. It would often be necessary to cut off questions after a number of minutes in order to make any progress through the material. We became skillful at, after any obvious confusion on a topic had been dealt with, looking for questions that lead naturally into the next topic to be covered and moving the class along in that manner. We also found students asking much deeper questions than is typical, sufficiently deep in many cases that we were not able to provide immediate answers and had to defer responding until we researched the topic.

In the first semester the class was transformed, there would typically be 50-100% more material prepared than would end up being covered in that class. We carried out suitable adjustments and so in subsequent semesters classes usually were about the right length, although it was not atypical to leave the last slide or two as material left for students to review on their own.

We greatly reduced the amount of in-class time spent on mathematical derivations in this course, which freed up class time for student discussions. Students have resources (problem solving sessions, posted lecture notes, book) to help them with the homework, so derivations that were important and previously were done in class were shifted to homework problems as part of solving an actual physics problem rather than as an abstract mathematical derivation. This was found to also make the derivations much more meaningful for the students. Other derivations were simply deleted as being unnecessary and inconsistent with the learning goals.

7. Getting student buy-in

Starting with the first day, we explained to the students why the class was being organized and run in this fashion. This included a discussion of the research on why these novel methods have shown to lead to better learning. This discussion largely eliminated complaints about the assigned seats, and greatly reduced resistance to other novel aspects of the course. A couple of other times during the term we advertised the value of the voluntary homework problem solving sessions discussed below, repeating testimonials spontaneously offered by students who found them extremely useful. We would also briefly discussed how the purpose of the sessions was to enhance their individual learning and what sorts of collaboration would best facilitate this learning. Our observations of the student-student interactions within the sessions showed that this was quite successful--rather more so than we would have guessed. The strongest inducement for the students to come to these sessions though was clearly the nature of the homework assignments themselves, as the most difficult assignments, which included some early ones, brought a large fraction of the class to the homework sessions, and for many, coming to these sessions to do their weekly homework then became routine.

8. Homework

Homework was a large part of the class; thereby recognizing its importance in the learning process. Problem sets were assigned and due weekly, except for weeks in which there were exams. A computer based homework system was used that allowed a mixture of multiple choice and numeric answers that are machine graded. It also has online submission of essay answers that were graded by TAs as discussed below.

8a. Homework problem sets included the following features:

- Problems in the homework were connected to the real world (real world for students that is). Substantial effort is put into creating problems where it is evident that there is a reason

someone would want to know the answer. This requires finding a context in which the physics of interest is relevant to some decisions one would make on a meaningful question, rather than having problems be simple abstract exercises. While the problem sets required a lot of time and were quite challenging to complete, students recognized the value of the problems and rated the homework as a very important contribution to their learning.

- Problem difficulty was high, requiring extensive reasoning and analysis. As well as helping learning, the difficulty also encourages the students to work collaboratively and come to the problems solving sessions. Our experience is that students work much more effectively in this environment than when they work alone, with much less time spent on unproductive struggling without making progress, or wasting large amounts of time because of simple math errors that they fail to catch.
- Students were asked to explain their reasoning behind their method of solving a problem or the logic of the answer through the use of a few essay questions and long answer questions. Students were asked to provide scientific arguments.

8b. Grading homework

To deal with the time required to grade essay/long answer questions, half the questions are graded simply on whether an answer was submitted while half are graded on correctness (the students do not know which one will be graded on which basis.) Even so, there is a substantial grading burden so the rubric below was developed to grade the essay questions more rapidly. This makes it possible for the marker to simply put in a number for each part in an excel spread sheet, and that number is both recorded and transmitted to the students automatically. To provide additional feedback, after each homework set is graded, examples of good, medium, and bad answers are posted, along with a short list of common student errors that were seen.

Rubric for grading essay questions.

1. Identifies physical principle or principles that are relevant to answering the question: (2 if correct, 0 if irrelevant principle, 1 if have both some relevant and some irrelevant principles.)
2. Explains how the principle(s) apply to the situation described in the problem: (1 if correct, 0 if not)
3. Employs proper reasoning to explain the logic in going from how principle applies to the situation to the answer to the question. (2, 1, 0 according to level of correctness. If #1 or #2 are 0, this should automatically be 0.)
4. Clarity of writing. (2 if good, 1 if difficult but can be figured out, 0 if incomprehensible. If 0 here, all the others will be left blank.)

So if the answer is scored as 5/7 and the question is worth 1 pt, the student receives 0.71 pts.

A simple but quite effective question that we have on every assignment is the following "Pick one question from the previous week's homework that you got wrong, and explain what your error was." The students answer this question after the solutions are posted, but their own solutions have not yet been marked. The answer to this question is marked only on whether or not it is submitted, and not on content (although students are never told that). Our surveys indicate that students recognize that this going back and reviewing their answers and figuring out what they did right or wrong is valuable to their learning, but they say in spite of that recognition, they would not do it if it wasn't graded.

8c. Problem solving sessions are offered that encourage students to work on homework problems together. The problem solving sessions are scheduled for a time/place such that the most students are available to attend. (This is determined at the beginning of the course using clickers to survey the class.) The problems sessions are run as scheduled times where there is a location for students to get together, and a TA or faculty member available. Students come and go as they wish and form groups of their own choosing. A graduate or undergraduate TA, or other qualified person familiar with the class is present. Their main job is to facilitate student-student discussions about the homework so that the students work through the problems together and come to a better understanding themselves. The staff are not there to lecture nor to tell the students how to do the problem, and this is explicitly announced in class. This is also emphasized to the TAs in a training session before the class begins, and strategies are discussed each week for how to deal with students "just wanting to get the answer," and how to help students progress towards figuring out the answers on their own through Socratic dialogue. For example, when asked a question by one student, turning around and asking other students what they think as to the answer to the question and why? Students are also given advice by the TAs concerning other ways to encourage group formation and effective problem solving. This has proven to be very successful with this course. In other courses with younger and/or less motivated students it was reasonably successful, but not the unqualified success we observe in this course, where there was a large fraction of the students regularly participating and developing highly effective collaborative and self-reliant problem solving approaches.

Other benefits to the homework sessions are that they provide a great deal of information about what topics and problems are the most difficult to students, and successful ways to overcome these difficulties. This information comes from observing the students and listening to their discussions, as well as seeing what questions they ask the TA. At the weekly TA meeting, these observations are compiled. The homework sessions also make it possible to get much more accurate measures of how long it takes the students to do the homework, because one can see when a student is nearing the end and just ask them how long they have been there working. This makes it possible to fairly accurately hit our target of having students spend 4-6 hours per week on average on the homework, which in turn allows us to ensure the overall course workload matches the specified university standard. (Although this is the standard, our surveys indicate it is far greater than the amount of time students spend weekly on homework in most courses.)

9. Exams

These were designed to be consistent with the learning goals, but were otherwise based on personal preferences. As a pragmatic decision, the exams were predominantly multiple choice that could be machine graded, with a variety of numerical and conceptual problems. To be consistent with the homework and focus of the course, approximately one quarter to one third of the points on each exam were based on essay answers requiring explanations and reasoning.

10. Use of undergraduate TAs.

In addition to the graduate TAs assigned to this course in the standard manner, we hire three undergraduates to assist with the class who either have previously been through the course or are advanced physics majors. These undergraduates cost much less than a regular TA. We advertise these positions and select the TAs from what is usually a substantial pool of very good applicants. They are selected after an interview based on their understanding of the material,

interest in teaching, and social skills. Ones who have been through the class and we have observed to be effective working in the collaborative problem solving sessions are particularly valuable. These TAs attend lecture and mingle throughout the class to facilitate effective student-student discussion on clicker questions. They also provide us feedback about the student behavior, focus, and learning in the class. One is assigned for each class to simply observe and write up a summary of the lecture and how students were reacting to each part.

There is a weekly meeting of all the graduate and undergraduate TAs, where we review issues that came up in lecture and homework sessions and generally discuss how the class is going and what is working well and not so well.

11. Evaluation of learning and general aspects of the class.

In addition to the quantum mechanics concept test and homework and exam scores, we have online surveys that measure student's beliefs and their assessment of the contributions to their learning from various components of the course. The beliefs survey is a standard well tested instrument (the CLASS), and the contributions to learning survey uses the professionally developed SALG survey. We also have the weekly summaries from the lecture and homework session observations. The postdoc working on this class compiles all of this information into a document that then is a guide for what is reused and what needs to be modified in the subsequent semester. There is a weekly online submission asking a couple of brief questions and allowing general input about the class. Students get a nominal number of points from this that in principle contributes to their grade. These online submissions provide an ongoing monitor of particular student difficulties or issues. The postdoc also carries out regular in depth interviews with the students about what they are learning, and what they think about the class.

Although the first semester of this course went fairly well, and by a number of measures was clearly more successful than its previous incarnation, there were obvious rough spots. The feedback compiled in this way made it quite easy to make adjustments and so the second semester went well, and required relatively little work.

12. Passing on to new instructor

All the materials for the course including the website materials, syllabus, course schedule, homework, PowerPoint files with lecture notes and clicker questions, and exams was collected on a CD. This was passed along in its entirety to a new instructor. One of the instructors spent a couple of hours explaining the course design and the materials available to him. The new instructor has used the lecture notes and clicker questions with almost no change, and has made minor changes in the homework and exams. The course has run extremely well with very high levels of student interest and participation and learning, and requires much less work by the new instructor than is normally the case when teaching a new course.

Future work.

Although we are no longer in charge of and hence in a position to continue to improve this course, there were some places where there was room for further improvement. The most prominent are:

- 1) We discovered to our surprise there is a significant minority of the students who see mathematics purely as a process, and do not understand the use of mathematical equations to describe concepts and physical relationships ("algebraic reasoning"). It was not anticipated that this would be the case in this population of students (engineering majors,

roughly equally divided between 2, 3rd and 4th year), and hence there is little in the instruction to try and address this lack of algebraic reasoning.

- 2) While most students are quite engaged in the in-class peer discussions about the clicker questions, there are a substantial number that will relatively quickly agree upon an answer, and then stop thinking about or discussing the question. This clearly varies with the question asked, and would likely be improved if more of the questions asked were more challenging and engaging. As a likely related issue, there are also regular indications from the observations of homework sessions that some students are not absorbing as much as we hoped from the lecture, although it is not clear what are realistic expectations, without further study.
- 3) Conversely, by the second semester we had become sufficiently good in choosing and using clicker questions that we were fairly routinely faced with deep substantive questions from a large fraction of the students in the class; so many that it was impossible to answer them all and make significant progress through the material. The challenge of dealing with a large class where so many different students are engaged and want to participate in discussions and questions that it results in substantial resentment from some students that material is not being covered is a problem that we had never encountered before.
- 4) While our students do much better than students in previous semesters or the physics majors who take a different intro quantum class, our students do not score as high as we hoped for on the concept test. There are indications from interviews that some of this is due to flaws in a few questions, and those are being fixed, but on others the flaw lies in the course and needs to be addressed. (Our subsequent research on student learning of quantum mechanics is suggesting that there are some fundamentally flawed very basic assumptions about how introductory quantum mechanics can be learned, that need to be addressed if students are to fully master introductory quantum in such a course. This is an ongoing research project.)
- 5) There continues to be a significant number of student complaints about the amount of work required for the course. Since the homework sessions allow us to monitor fairly carefully how much time a representative sample of the students are spending we know that they are spending typically 4-6 hours a week most weeks, and as little as 3 hours on a few light weeks. Since the stated standard for the university is 6 hours for such a 3 credit course, we consider these complaints about the excessive workload to be a problem with the requirements of teachers in other classes, and not a problem with our own.