Problem Solving: What I have learned from my students

James Sandefur, Math Department, Georgetown University

May 5, 2005

1 Background

Most of our students can quickly learn the techniques we teach and can apply them to problems that are similar to the ones we have worked in class. On the other hand, these same students often have difficulty generalizing techniques both to multi-step problems and to problems that are somewhat different from what they have seen. They often exclaim "You haven't shown us how to work this." In other words, most of our students are not adept at either problem solving or at constructing simple proofs.

I will first describe what I mean by problem solving and how that differs from constructing proofs. Consider a two-person game in which there is a pile of n beads, and on each turn a player can pick up one, two, or three beads. The winner is the person who picks up the last bead. If each person plays optimally, who wins the game, the person who goes first or second, and what is the optimal strategy? **Problem solving** consists of the students determining the winning strategy and developing insight into why the strategy works. They might play the game using different sized piles of beads. Eventually, they discover that the key is whether n is divisible by 4 or not. If $n = 4k$, then the person who goes second should win. If $n = 4k + 1$, $4k + 2$,or $4k + 3$, then the first person should win. The reason depends on the fact that no matter how many beads the first player takes, the second player can take a number of beads so that the total picked up by the two players is 4. This is problem solving. Once the problem has "been solved", the next task is to construct a **proof** that the solution is correct. In this case, we could use induction to prove that the second player wins when $n = 4k$. Then we could prove that the first person wins as a corollary of the first result, that is, if $n = 4k + i$, $i = 1, 2$, or 3, then the first player takes i beads, and we then apply the $n = 4k$ result to the remaining pile of beads. Our students have had very little practice in constructing proofs. They think that once they have found the solution, they are finished.

In watching students try to construct proofs, it is clear that they are missing many key ideas when they don't go any further than gaining insight. For example, I might ask the students to prove the implication that if the product of two integers, nm, is odd then each of the integers must be odd. They might say that if each of the integers is odd, then the product must be odd. They might even write the equation

$$nm = (2k + 1)(2i + 1) = 2(2ki + k + i) + 1$$

and claim this proves the result. Many do not to realize that they have proven the reverse implication, that the product of two odds is odd, which is not what they were asked to show. They may not even realize the statements are different. To gain insight they should consider the false implication that if nm is even, then n and m must be even.

We have tried numerous approaches to improving our students' ability to understand and construct proofs, such as presenting more proofs in calculus and requiring a proof based course, Foundation of Mathematics. These efforts have failed to achieve their objectives for many reasons. If we didn't test our calculus students on proofs, they would not pay attention when we presented a proof. If we did test the students, they would memorize the proofs with little understanding. In Foundations, the students saw the proofs as content specific and did not generalize the approaches to their other courses. Also, many students took Foundations in their senior year, so faculty teaching other courses could not build on what was learned in Foundations. But one of the most important reasons that we failed is that none of us really knew how students learn to solve problems.

2 The Research Study

In 2001, I began teaching the Foundations course. Simultaneously, I began a research program to try to understand how students learn to problem solve, and what I could do to enhance that learning. I was reminded of a story about some scholars during the Middle Ages who were arguing about the number of teeth in the mouth of a horse. Finally, one of the scholars suggested they actually go look at a horse. So it was with me: if I was going to understand how students learned to problem solve, I would have to actually watch them solve problems.

As a pure mathematician, I was not sure where to start. Fortunately, I was aided by a small grant from the University, and the support of staff from our Center for New Designs in Scholarship and Learning (CNDLS), particularly, Susannah McGowan. Over the past 4 years, with CNDLS help, I have been conducting video Think Alouds (TA), in which I videotape students doing their homework. The particular structure is to videotape each student working on the problem alone, then bring the students together and videotape them while they continued working on the problem in a group. Someone would be there to encourage the students to verbalize what they are thinking while they work on the problem, and to give them hints or prompts if they became stuck.

Through watching and rewatching the TAs, and discussing them with others, I have gained a better understanding of our students' difficulties. Using this understanding, I have experimented with the structure of Foundations, and continued use of TAs has helped me judge which approaches are more effective in supporting students as they become independent problem solvers. In the following, I will share some of what I have learned.

3 Initial Observations

From observing my students on videotape, it appears that students 1) are more than willing to do the work we require of them, 2) get stuck near the beginning of many problems, 3) cannot change directions when they are using an unproductive approach, and 4) do not use examples to help understand either the question or what approaches might help with the answer. Let me discuss each of these points in more detail.

Observation 1: Students are willing to work. I had the misconception that poor quality or missing homework was the result of students not making sufficient effort. Early on, I discover this is not always the case. For my first TA, I asked a group of students to find all values for a and b such that the function

$$f(x) = \begin{cases} \frac{1}{ax+b} & x \neq -\frac{b}{a} \\ -\frac{1}{b} - \frac{b}{a} & x = -\frac{b}{a} \end{cases}$$

has a 3-cycle, that is,

$$f(f(f(x))) = x$$

for some x. The solution is to algebraically simplify the equation $f(f(f(x))) = x$ to

$$(a + b^2)(ax^2 + bx - 1) = 0$$

For a 3-cycle to exist, we would need 3 x-values, but solving the quadratic gives only 2. Alternatively, if $a + b^2 = 0$, then we have a 3 cycle for (almost) any x. This answer makes sense, since the problem asks about a and b values, and not x values. This seemed to be a reasonably straight forward problem to me, requiring a little algebra and some thinking about the question being asked.

Each of four students was videotaped for 30 minutes working on this problem on their own. They then continued working on this problem as a group for another hour. The idea is to have each of them think about the problem separately, then let them come together to share ideas. This is the system I have used over the entire 4 years.

After working for 90 minutes, the group did not even understand the question. They would solve for x, then note that this wasn't the answer since it asked for a and b. They would then substitute some values in hoping to magically find a 3-cycle, then would go back and solve for x again. What was surprising was that the students wanted to continue working on this problem, even though they had accomplished little in the 90 minutes. Thus, I learned my first general principle that I have seen repeatedly in TA's: A lack of results does not mean there is a lack of effort. Sure, there are some students who put in a minimum amount of work or quickly stop working when they get stuck. But I have found that most of our math students are willing to put a great deal of effort into their work if they find the question interesting. The sad part is that this effort is often wasted time, with the students gaining little from the experience. From this, I developed my first goal in revising Foundations.

Goal 1: Help students spend their time more productively.

Observation 2: Students are often stuck at the very beginning of a problem. The students who were trying to determine the a and b values that resulted in 3-cycles never even understand the question being asked. In another situation, one of these students was videotaped working a graph theory problem. For 40 minutes, the student just kept repeating the question, but never made a first step.

Goal 2: Help students get started on their problems.

Observation 3: Students keep repeating the same steps, even when they clearly do not help. In watching my students work, I noticed they were like a wind-up toy car that is stuck in the corner, wheels spinning but not going anywhere. Once students decide how to approach a problem, they have difficulty trying a new approach.

One group of students was trying to show that the Power Set P_{n+1} on $n+1$ objects contains twice as many sets as the Power Set P_n on n objects. I suggested they try constructing P_1, P_2, P_3, and P_4, the Power Sets on 1, 2, 3, and 4 objects, respectively. The strategy the students used to construct each Power Set was to first list the null set, then all the sets of size 1, then all the sets of size 2, and so forth. They did not use the Power Set on 2 objects to help construct the Power Set on 3 objects.

I had thought the students would see that they could use P_2 to construct P_3 by listing P_2, all of the sets on the first 2 objects,

$$\emptyset, \{1\}, \{2\}, \{1,2\}$$

plus all those sets with the third object added,

$$\{3\}, \{1,3\}, \{2,3\}, \{1,2,3\}$$

The students had not observed this connection between Power Sets. So I suggested that the students look for connections between the Power Set on 2 objects and the Power Set on 3 objects, again hoping they would see how one could be used to construct the next one. Instead, the students tried comparing sets according to how many elements were in the set, so P_2 and P_3 both contain the null-set, and there is one more subset of size 1 in P_3 than in P_2. This approach becomes more complicated when comparing the number of subsets with 2 or 3 elements, or when looking for the relationship between P_3 and P_4. This approach would clearly be difficult to generalize from P_n to P_{n+1}. Once these students developed a strategy for constructing Power Sets, they were having difficulty changing to another approach.

Goal 3: Help students learn how to try a variety of approaches.

Observation 4: Students avoid looking at examples and instead try to solve the general problem. One group of students was asked to show, if the integer n has an odd factor greater than 1, then 2^n+1 is not prime. During the individual portion of the TA, one student was juggling symbols, not getting anywhere. We then had the following conversation.

I said, "You seem to be wandering around hoping something will pop out." He replied, "Yeah, that's what I usually do." I said, "Did you think about checking a few examples to see if the statement is correct?" to which he replied, “Yeah, heh heh, ummh, that would also have been helpful for problem 3 (a problem assigned earlier).”

Once this student starting checking some values, he discovered that whenever n was odd, then 3 was a factor of $2^n + 1$. This helped him prove part of the result. Continued work with examples helped the student discover the general pattern, which his group was finally able to solve (after several more prompts to keep them going in a productive direction).

Students are not only reluctant to try examples, but when they do try examples, they tend not to reflect on how the examples can help them solve the problem, just as the group working on Power Sets did not reflect on how finding P_3 could be used to find P_4.

Goal 4 Help students learn to use examples to develop a better understanding of the problem and to reflect on examples to help them understand the general situation.

4 What helps?

Of the numerous approaches I have used to help accomplish my goals, the most successful strategies have been 1) to give prompts or hints that help the students structure their approach to a problem, 2) to have the students submit multiple drafts of their solutions, 3) to have students develop their own conjectures through the use of examples, and 4) to have students write each solution two ways, one which includes their thinking about the problem and a second that is a formal proof. Let me discuss each of these strategies in more detail.

Strategy 1: Give prompts. I have had more success with prompts than with hints. Let me explain the difference. A prompt helps the students organize their approach. They tend to be the questions that we ask ourselves when we are solving a problem, and they tend to be the same, regardless of the problem, such as

What is given?

How can I rewrite what is given in a form that is easily used, possibly using a definition?

What must I show?

How can I rewrite what I must show in a form that is usable?

What method of proof seems most likely to work, and why?

On the other hand, a hint is specific to the particular problem. It may suggest how to accomplish one of the steps of the problem. I gave the students a hint to look at the difference between the Power Set on 2 objects and the Power Set on 3 objects. As another hint, I might have suggested that the student determine common factors of $2^n + 1$ when n is odd.

What I discovered as I began teaching this course is that prompts are the questions I ask myself, almost subconsciously. To develop good prompts, I have learned to observe my own thinking as I work problems, somewhat of an internal Think Aloud, and focus on what I am asking myself as I work problems. Until I started observing myself, I didn't even realize I was asking myself questions. It had developed to the point that it was subconscious.

Let me give an example of how these questions work. Students generally do not know the difference between what is given and what must be shown, and they are also reluctant to use definitions. For example, students intuitively know what it means for an integer n to be odd, but it rarely occurs to them to use a definition of odd, such as there is an integer k such that $n = 2k + 1$. Suppose students are to show that the product of two odd integers is odd. In response to the questions, what are we given and how can we rewrite it in a usable form, the students would write, "We are given that n and m are odd. A usable form of this statement is that there are integers i and j such that $n = 2i + 1$ and $m = 2j + 1$." In response to the questions, what must we show and what is a usable form of that statement, the students would write, "We must show that nm is odd. A usable form of this statement is that there exists an integer k such that $nm = 2k + 1$." From here the students would compute

$$nm = (2i + 1)(2j + 1) = 2(2ij + i + j) + 1 = 2k + 1$$

where $k = 2ij + i + j$. What is amazing is that for many students, I need to keep repeating these questions throughout the course. It takes most students a long time to internalize these questions.

Even though the prompts seem obvious, the students tend not to think in this way. They have usually approached mathematics as a series of equations that need to be rewritten, and they do not pay close attention to the specific words in a problem, especially the words "for all" and "there exists." They start learning that if a statement begins, "for all $x \cdots$", then the method of proof will probably begin "choose $x \cdots$." If the statement includes "there exists $\cdots$", then at some point, we will probably have to find it. Some good early exercises for the students are to give them statements such as

$$p = \{\forall\ x,\ \exists\, y \text{ such that the equation is satisfied}\}$$

$$q = \{\exists\, y \text{ such that } \forall\, x, \text{ the equation is satisfied}\}$$

and ask them for which of the following equations each of these statements is true.

$$\begin{aligned} &2x + y = -2 \\ &x + y^2 = 3 \\ &(x-2)(y+3) = 0 \end{aligned}$$

These are later followed up with the use of prompts that focus students attention on the phrasing of the problem.

When students begin to work on a proof they usually wander around with no direction. They hope for luck. It is particularly important to give prompts that help students focus on what type of approach might work. Among the types of proofs they know are direct proofs (in which we may work forwards and/or backwards), contraposition, contradiction, and induction. When I give the prompt, "what type of proof seems most likely to work and why?", the students know that: 1) if the statement is to show the existence of something (for every x, there exists a y such that), then a direct proof may work, in which we actually "find" y; 2) if the conclusion of the statement includes the word "not", then contraposition or contradiction might work (there does not exist an x such that) because it is difficult to show something doesn't exist, so let's assume it does exist and find a problem; and 3) if the conclusion is that something is true for every positive integer n, then induction might work.

Strategy 2: Multiple drafts: For difficult problems, hints will clearly help the students. The problem is that a hint may help one student and be of no use to another. In one TA, students were asked to show that a graph G on n vertices can be isomorphic to its complement G^c only if

$$n = 0 \text{ or } n = 1 \bmod 4$$

The complement of a graph on n vertices is the graph G^c which has edge uv if and only if uv is not an edge in G. The idea behind the proof is that for G and G^c to be isomorphic, they must have the same number of edges, which means the complete graph on n vertices must have an even number of edges, which can only occur if there exists an integer k such that

$$n = 4k \text{ or } n = 4k + 1$$

I gave the students the following hint. What do you know about the total number of edges in G and G^c? Answer: They are the same. At first, this hint made no sense to the students. After a short period of time, the students actually constructed some specific examples, finding isomorphic pairs, G and G^c, with 4 vertices and again with 5 vertices. They noticed that the complete graph on 6 vertices has 15 edges, which cannot be divided equally. Then one student exclaims, "Oh, so when he asks what we know about the total number of edges in G and G^c, he means the total number of edges in G and G^c." The hint finally made sense, and the students progressed quickly to the solution. This hint worked well.

As described previously, when I suggested that students look at the difference between the Power Sets on 2 and 3 objects, these students looked for the wrong connections, and the hint was of no value.

Hints work better if they are somewhat student specific. Because of this, I have instituted a multiple draft system. I the students will have a problem with a given deadline. The problem will be structured in several parts, with each part giving the students a prompt, such as clearly state, in a usable form, what must be shown. The students are required to type all solutions (most use Word with Equation Editor) and email the solutions to me by the deadline. To ease in grading, I have each group of 3 or 4 students send me one solution. Because it is electronic, I can give each group individual hints, depending on the problems I see with their solution. My hints and comments are embedded within the document using the Track Changes feature of Word.

Each problem is worth 100 points. For most problems, I do not expect the first draft to be correct. If the students have made a good effort and are going in a productive direction, I might not deduct any points. On the other hand, if the students are going in the wrong direction, and need a major hint, I will give a 5 or 10 point deduction, which will be deducted from the grade of the final version of the solution. If it seems that the students did not put much effort into their first draft, then I will deduct 10 points, and but will not give a good hint. I want to encourage a good initial effort.

The students then have several days to write a second draft of the problem. I repeat this process, with additional possible deductions, until the students essentially have the correct solution. This way, the initial problem includes appropriate prompts, and I can individualize hints for each group of students. At the end, each student must have a solution that is correct except for minor details, but their grade may be in the 80's or even 70's if it took several major hints for them to get this solution.

This system in which each problem is graded several times is time consuming for me. Because of this, I assign fewer problems, but they are generally harder and more involved. Having solutions typed also makes the grading easier. For each problem, there are often multiple places where students make mistakes. I generally write a document with a collection of hints that might be given, and then can just copy and paste the appropriate hints into each group's paper. This also saves time.

There are other advantages to having multiple drafts. Students do not spend too much time working unproductively, since they can get feedback. To insure that each student in a group is involved, I require each draft to be typed by a different student. Having the electronic version of the solutions allows me to use some papers for illustrations in the class in later semesters. And if a student's solution is too good, I can quickly and easily compare it to solutions I have received in previous semesters.

Strategy 3: Making conjectures: Instead of asking students to prove a statement, I now try to have more problems that cause students to explore a situation and construct their own conjectures. This helps combine problem solving with construction of proofs. Students are encouraged to try to understand why the examples work by reflecting on their construction. When I began, I would have asked students to show that if n has an odd factor greater than 1, then $2^n + 1$ is composite. I now ask students to look for patterns in the factors of $2^n + 1$. They will quickly conclude that if n is odd, then 3 divides $2^n + 1$, and will prove it using induction. After some additional work, many students will observe that if $n = 2(2k + 1)$, then 5 divides $2^n + 1$. Continued work will lead students to the result that if n is not a power of 2, then $2^n + 1$ is not prime. The proof is still not easy for the students, but they get the sense of discovery, and even the weaker students can get some partial results. They also develop some ownership for the problem.

One difficulty that often arises is that students have so much fun looking for patterns, they forget to actually stop and prove some of their results. A second problem is that the problem must be worded carefully so that the students actually see some patterns. Many students are so disorganized that they may work for a long time without making any discoveries.

I believe one of the greatest shortcomings of post-secondary mathematics education is the lack of development of our students' ability to look for patterns and make conjectures. And yet, this is the basis for mathematical research.

Strategy 4: Two different types of proof: Students begin my course with only a vague notion of what a proof is. I might give them an implication, $p \Rightarrow q$. They will work from both ends until they get an equation that is known to be true. At that point, they think they have a proof. Consider the statement, "if n and m are consecutive integers, then $n^2 + m^2 - 1$ is divisible by 4." The students confuse what is given with what needs to be shown. They might turn the following in as their proof.

$$\begin{aligned} n^2 + m^2 - 1 &= 4k \\ n^2 + (n+1)^2 - 1 &= 4k \\ n^2 + n^2 + 2n + 1 - 1 &= 4k \\ 2n(n+1) &= 4k \end{aligned}$$

so

$$k = \frac{n(n+1)}{2}$$

While a mathematician would realize this actually contains the proof since $n(n+1)$ must be even, it has become clear to me that students do not know why this is a proof. They work the problem as n, m and k were all given constants and that they only need to check that everything balances in the equation $n^2 + m^2 - 1 = 4k$. They seem not to realize that n and m can be treated as known integers and that we must show an integer k exists that balances the equation.

I give problems that have several parts, corresponding to my prompts. The first part would be to restate what is given in an usable form (as discussed previously), and the second part would be to state what is to be shown in usable form. The third part has them show their "thinking", that is, write a solution such as the equations above. This helps them develop insight into why the statement is true. In the last part, they must construct a formal proof in which they begin with what is given and work forward until they reach what was to be shown. This last part requires the students to tie the other three parts together, and is crucial in helping students develop a better understanding into what a proof actually is.

There may be another part to the problem asking the students what proof method might work. This part usually follows the second part in which they write what needs to be shown. If they decide on contraposition or contradiction, then they would have to rework the first two parts, so that they have the contrapositive of what is assumed and what must be shown.

Strategy 5: Presentations. I have had good success with having students give group presentations. Early in the semester, groups will present simple problems, taking around 10 minutes. At the end of the semester, each group gives a larger presentation, usually consisting of the proof of a major theorem from the text and the solution to a problem. Students must learn these proofs at a deeper level than usual so they can help other students understand the problem. These presentations are similar in structure to the assigned problems discussed previously; that is, the students will describe what is given and what needs to be shown, and will present an informal proof, describing how they thought about the problem. Then they will give the formal proof. Most students now give Power Point presentations. The use of animation makes their presentation come alive. Students used to complain that they didn't get as much from student presentations as they got from my lecture. That is no longer a complaint.

Each group must meet with me twice before giving their presentation. The first meeting is so I know they are on the right track. In the second meeting, they actually give the presentation to me. This helps avoids embarrassing errors in the actual presentation. Several pointers that I give to the students are 1) make sure slides or transparencies are in large enough fonts, 2) each slide has one idea, 3) a computation is not continued from one slide to the next, 4) each slide should have very few words, and 5) they should engage the class, possibly asking questions about what the next slide should be. The fourth pointer is critical. Cluttered slides are difficult to follow. The goal is to only have the essentials on the slide and then fill in the details with the verbal presentation.

For the presentations to work, the students know that they will be responsible for the material on a final exam.

5 What Problems Remain?

Learning to teach problem solving has been one of the most difficult challenges of my academic career. After teaching this class 7 times in the past 4 years, and watching hours of videotape, I still feel I am only beginning to understand how students learn problem solving. And there are still many obstacles preventing more success.

One obstacle is finding a significant number of good problems that are at an appropriate level for my students. More often than not, a problem that seems easy to me is too difficult for my students. Equally often, I assign problems that are too easy, and the students gain little from working them. Giving appropriate hints is also quite difficult to do. In an effort to not give away the solution, my hints still tend to be too obtuse, and consequently, not much help to the students. Giving good hints is an art that I have not mastered.

Teaching problem solving is quite time consuming. Each semester, I try some variation in an effort to reduce the time I spend either meeting with students or grading papers. The only short cut that has had some success is the use of multiple drafts. This reduces the number of students who visit my office and the grading of later drafts is usually not that time consuming since they have already been given some direction. Most short cuts I have tried have not reduced my time commitment, but have resulted in less success in my students.

Each class presentation must be carefully prepared. For the students to gain from my presentation, I must carefully construct the steps they need to take through a problem. I must make sure that my presentations clearly model what they should be doing. I must also carefully choose similar problems for them to try in groups during class. The students need this supervised practice.

When most students enter college, they are not nearly as competent problem solvers as they should be. Some of this results from the lack of time spent solving problems and reasoning in the elementary, middle, and secondary mathematics classrooms. I do get the occasional student who has had some experience in proof and problem solving, but most of our students are just beginning to learn to reason when they take Foundations. I have been somewhat disappointed in how little they can accomplish in one semester. To be successful, we must reinforce what they have learned in our other courses. This means faculty must talk among themselves and share ideas, both those that work and those that do not work. And finally, we should continue to carefully observe our students to make sure we improve in our ability to help them.