Report on MSRI 1994 Summer Workshop on Hyperbolic Geometry and Dynamical Systems

David Epstein, Jane Gilman, and Bill Thurston

During the summer of 1994 the three of us gave a two-week workshop on hyperbolic geometry and dynamical systems at the Mathematical Sciences Research Institute (MSRI, pronounced "emissary") in Berkeley. The participants were about fifty graduate students chosen by their universities, most with no particular background in the subject. The workshop used nonstandard teaching methods and treated nonstandard topics. We describe here the thinking that led up to the workshop, our expectations and methods, and how things turned out in practice. Since there has been little discussion in the mathematical community about teaching and learning at the graduate level, we think that an account of our experience (mistakes and all) may be informative and useful in stimulating dialogue about graduate-level education.

We are grateful to Sheila Newbery and Silvio Levy for reading, editing, and significantly improving this article and to the editor and anonymous readers for valuable comments.

David Epstein is professor of mathematics at the University of Warwick, Coventry, U.K. His e-mail address is dbae@maths.warwick.cs.uk.

Jane Gilman is professor of mathematics at Rutgers-Newark. Her e-mail address is gilman@andromeda.rutgers.edu.

Bill Thurston is director of MSRI and professor at the University of California at Berkeley. His e-mail address is wpt@msri.org.

Hopes, Goals, and Philosophy

The workshop was an attempt to combine for a graduate audience our individual efforts in recent years to develop new approaches to teaching at the undergraduate level. The impetus for these efforts has been our dissatisfaction with the methods of pedagogy familiar to most of us. These methods tend to be of two types: collective, formal lectures on the one hand and self-study on the other. Many students profit from both, and we use both in our teaching, but neither stimulates interaction among students or between students and teachers. The result of this is that much of the energy latent in these relationships is never harnessed to further pedagogical aims.

Teaching mathematics need not be a choice between giving formal lectures and leaving the students to discover everything for themselves. One can consciously vary the style of one's interaction with a class and see what works best. Our workshop was an opportunity to try different approaches. We had some common ideas about what these approaches should be, and Bill and Jane had even tried out some of them in a team-teaching context. (We will use names to identify subsets of the authors.)

First, we wanted to avoid the over-emphasis prevalent in math education on formal definition and symbolic reasoning. Such reasoning is of course essential, but it is often difficult to communicate, and when presented too quickly or magisterially it can dampen a student's enthusiasm. In our workshop we led the students into the subject matter through hands-on experi-

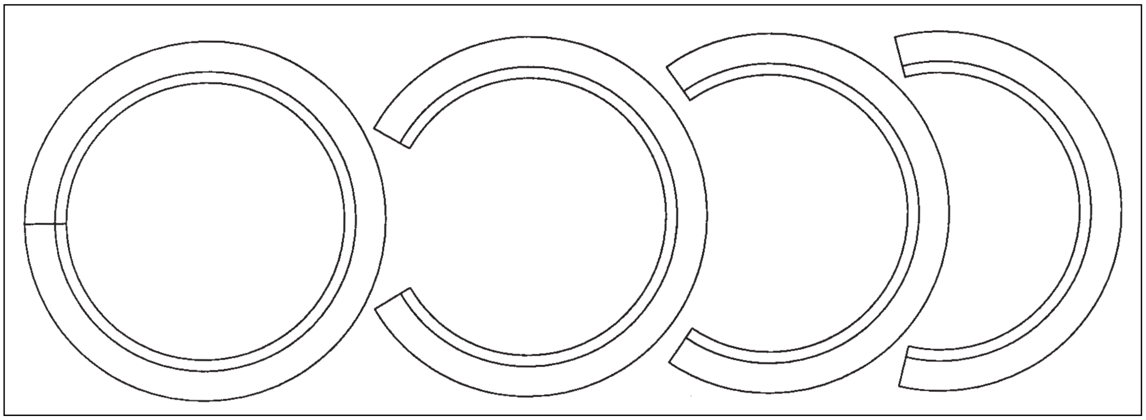

Figure 1. Gluing the outer edge of each strip to the inner edge of the one to its right causes the strips to warp. The result is paper of (approximately) constant negative curvature. The more strips are added, the more total curvature the surface acquires. Enlarge the figure and try it yourself.

ments that, we hoped, would encourage them to make connections between their own sensory observations and the symbolic representation of traditional teaching. This had been very successful at the undergraduate level: we've learned that theoretical knowledge takes on a more immediate and inviting importance for the students when it concerns an object that is "really there".

Another fundamental concern had been to lessen the ethos of competition among students. We prefer to emphasize cooperative learning, which reinforces the students' own capacity to discover ideas and tends to be more fun as well. Cooperative learning is hard in a lecture context, where members of an audience are expected to just absorb rather than find or share ideas. Because the finding and sharing of ideas was exactly what we wished to foster, our approach was to promote group learning and try to draw everyone into active participation. Our decision to team-teach was itself part of a conscious effort to provide a model of interaction for the students.

A third goal was to encourage students to explore in greater depth subjects of particular interest to them. Traditional courses, because of the pressure to cover material, often allow little opportunity for the type of spontaneous, in-depth exploration that nourishes a student's motivation to learn.

Finally, and perhaps most importantly, we wanted to share with students some of the beauty of mathematics, or rather some of what makes mathematics beautiful for us.

Scissors and Glue Sticks

Even with the strength of our shared teaching philosophy, it was a somewhat daunting prospect to lead a nonstandard workshop for a large, hand-picked graduate audience for the first time. Competition to get into the workshop had been stiff, and the students' expectations would be high. Another challenge for us was the heterogeneity of the class—some students had years of research experience already, some none.

We decided to assume no background in hyperbolic geometry, but to try to find material that would be new and interesting to as many students as possible. In the end, a large number of relevant mathematical topics were discussed in the workshop, including many not raised by us. We can't do justice to the latter, since they were too numerous and our knowledge of them came mostly from snatches of overheard conversations. Instead we mention the topics that we raised ourselves, and especially those that we approached in an unconventional way, at the same time that we try to give a feel for the dynamics of the workshop.

When students entered the MSRI lecture hall on the first day, they were each handed a glue stick and several pages of annular shapes to cut out (see figure above). Half of them were given scissors. They were told to begin cutting and pasting to make "hyperbolic paper" while other students arrived, office assignments were given out, and attendance and lunch orders were taken. They reacted well to the unusual demands we made of them and worked assiduously and with clear enjoyment on a task they might well have decided was more appropriate to a class of six-year-olds.

Students were then asked to discuss the following problem in small groups. Let S be the operation of adding one, and let D be the operation of doubling. How many integers does one obtain if one starts with 1 and applies all possible words in S and D of length at most n? Is there a connection between this problem and the shapes one gets by gluing together the annular cutout shapes?

Our goal in handing out paper, scissors, and glue sticks was to draw people in, equalize individuals with different backgrounds, and break the spell of symbolic representation, conventionally regarded as the be-all and end-all of mathematical thought. Hyperbolic paper gives a natural introduction to the hyperbolic Riemannian metric on the upper half-plane and a striking physical demonstration of the rotational symmetry at any point of the hyperbolic plane. We did not expect the students to discover the formula for the metric on the upper half-plane without guidance, but when we did introduce it, we explained it in a way that stems naturally from the paper models. In a traditional course, students usually don't even think of asking why the formula is what it is.

Bill offered a \$100 prize to the first person making a disk of hyperbolic paper three feet in diameter. (The target size had been conservatively chosen so Bill wouldn't have to pay up—it was several times the radius of curvature of the paper, which is controlled by the inner and outer radii of the strips.) Several students investigating this challenge were able to prove by themselves how hyperbolic area grows exponentially, and estimate the amount of paper it would take to win the prize—on the order of 10^6 sheets. Others gained an intuitive feeling of exponential growth simply by doing a lot of cutting and gluing.

Paper models also gave a concrete dimension to the discussion of developing maps. Pairs of pants (spheres minus three disks) are one type of topology that arises naturally when many people make hyperbolic paper and then glue their contributions together. A few pants were made in this way; to control their shape, hyperbolic trigonometry was needed. Interesting practical and theoretical problems arise if one wants symmetrical pants. To consider the set of all possible pants, we introduced Teichmüller space. Paper models of the pseudosphere also provided an introduction to horocycles, horospheres, and cusps.

Joe Christy, head of mathematical computing at MSRI, wrote *Mathematica* software to handle inversions. This was used for student investigations of limit sets of groups generated by inversions in the plane, though some students preferred ruler and compass to a computer. We then introduced the convex hull construction to give students a better hold on the structure; many had already produced their own proofs in the plane, so they appreciated the simplicity of the more powerful approach.

Maintaining the Level of Intensity

Each morning of the first week was filled with group discussions, more hands-on activity, and short lectures (twenty to thirty minutes) by us. We made a strong effort to ensure that everyone was following what was going on. A lot of mathematics was covered in this way, and we will talk in more detail about this experience in the next section.

But during the second week there was a tendency to revert to a standard lecture mode for much of the time. There were two main reasons for this. First, we felt we ought to offer more advanced topics to students who came with more substantial backgrounds in hyperbolic geometry and dynamical systems. Second, there is a lot of beautiful mathematics that we wanted to at least touch upon; since we could not cover everything in two weeks, we hoped to provide students with a sort of road map for continuing their studies.

We discussed many topics in a more conventional framework: the various models for hyperbolic geometry and the shape of geodesics, horocycles, and horospheres in them; orbifolds; traces in $SL(2,\mathbf{R})$ and $SL(2,\mathbf{C})$; topological groups and discrete subgroups; dynamics of group actions; minimal sets and invariant sets; proper discontinuity; Dehn twists; the Poincaré polyhedron theorem; circle packings on compact surfaces of constant curvature; Andreyev's Theorem; quasiconformal deformation theory and the measurable Riemann mapping theorem.

Not all of these topics were treated cursorily. For instance, we spent some time on ruler constructions in the projective model of the hyperbolic plane, which can be carried out by hand; and some students worked out particular constructions for themselves. We also used the viewing software *Geomview* [1] to study symmetry groups in two and three dimensions, and explore several aspects of hyperbolic geometry and three-manifold theory after a showing of the video "Not Knot" [2]. But, in general, the second week was more "traditional" than the first.

Our impression is that most of the students would have been happy to have continued with the style of the first week. Indeed, the change of style led to some students complaining that we had lost them—we had adopted the standard style and the outcome was standard, except that the students had fewer inhibitions about saying they were lost. But the more experienced students were stimulated by the advanced work.

Everyone Contributes

We tried to communicate clearly our goals and philosophy to the class. On the first day we said that each person was expected to contribute in some way to the workshop: by working on the notes (note-making, as opposed to note-taking), by contributing to the student-run afternoon session, by answering other students' questions, by teaching someone how to use emacs (a text editor) or *Geomview*, by writing software, by asking questions, and so on. We didn't expect everyone's contribution to be the same; it should have been something appropriate to the individual.

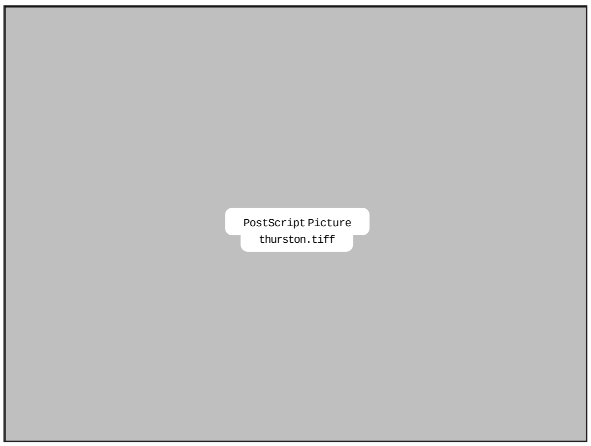

We also talked about the differences in background among members of the class and explained our view that mathematics is not a race but a process of gaining understanding. We revealed that it was no accident that we had only enough scissors on the first day for half of the students. The budget was big enough to buy scissors for everyone, but, by sharing, students would be more likely to talk to each other about what they were doing as they worked.

One excellent mechanism for getting everyone involved was student-run sessions. We hit upon it serendipitously. Despite our attempts to plan in advance by e-mail, we did not have a complete plan at the start of the workshop. We had carefully thought out the first few days, and we counted on the energy and good sense of the students themselves to continue in the same tone. But we knew that the three of us would need to spend time together regularly to review the day's events and plan for the next day.

What we hadn't realized was that this requirement could dovetail with the students' own need for time to work independently of us. The way it worked out was this: we kept the afternoons free of our lecturing and arranged a student-run session every day, starting at around 4 pm, after students had had some time to think over what had been done in the morning session—time to get stuck, develop insights, and so on. We never gave any real thought to what would happen in these sessions; we figured it was up to the students to work it out.

This was a hit. We couldn't resist being present at the beginning of each student session, and it was fascinating to watch the students interact with each other. For example, at the beginning of the first session, one of the natural leaders in the group marched up to the front and started to chair the meeting. She made it clear that there would need to be different people chairing the session on other days, so this was not a power bid, but an (ultimately) successful attempt to keep the meeting focused and functioning effectively. Subsequent chairs of sessions emerged as volunteers, at times with some slight group pressure.

The students spent their first session gathering questions. This may have been prompted by our suggestion in the morning that one way to contribute to the workshop was by asking questions, and it probably seemed to them a safe way to start. Also, our morning activities had raised an enormous number of questions that would take some of the students the full two weeks and beyond to answer. But the students may have felt that they ought to have known some of the answers immediately. They were hoping to get some of the answers from their fellow students.

We three stayed at the back of the room, biting our tongues to avoid plunging in with our own comments—even if the person at the blackboard said something wrong. We left students to probe, question, and correct. They usually did this in a supportive and nondestructive way. We let the students know that we were there as a resource, ready to offer comments if asked, but that otherwise we would keep quiet. In fact we were virtually never called on throughout the two weeks.

After a half-hour or so the three of us would leave for our preparation of the next day's activities. On the first day, an hour and a half after the session's beginning, one of us peeked into the lecture room. About half of the students were still there; they were working in small groups gathered about various pieces of blackboard or around the paper models they had made, talking animatedly.

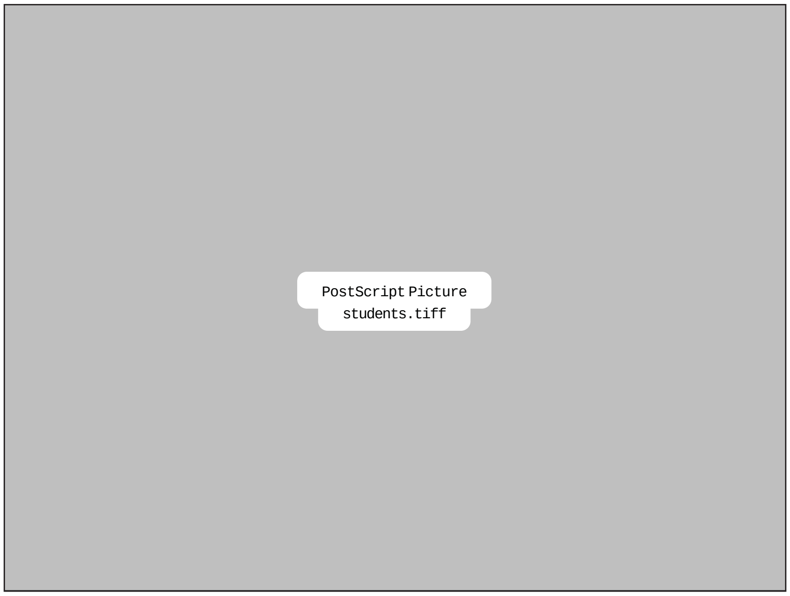

The second student session was even more fascinating. Students decided that unanswered questions were not so useful. Instead many had answers to problems they had been working on, explanations for matters we had left vague, answers to some of the easier questions asked by students the day before, and explanations of what they had done. They also decided (possibly after watching the three of us lose most people's attention when we went on too long at the blackboard) that there should be a time limit for each person. After a long debate, it was decided that each student wanting to make a contribution would choose how long to speak, and write his name on the blackboard during the day with the amount of time needed. Someone made the ingenious suggestion that the shortest presentations should always get priority.

The next problem was what if someone promised to talk for ten minutes and instead talked for thirty? To meet this problem a bell appeared. When people spoke beyond their time limit, the bell was rung by a time-keeper and they had to stop. The use of the bell was later modified by the introduction of a one-minute warning.

Team Teaching

In fact, the bell turned out to be useful in our morning lectures as well, particularly when two of us felt that the third, who happened to be talking at the moment, had lost the audience. All three of us were in the lecture hall together. Although in theory we took turns speaking, in practice we were all actively involved at all times, interrupting each other to challenge, correct, amplify, and ask questions. This aspect of team teaching is simultaneously demanding, nerve-wracking, stimulating, and inspiring.

In their subsequent comments, students showed that they were often uncomfortable with these interruptions:

> Having three teachers was a great idea…

> It was great! I love the different perspectives and the way you guys interact with each other. Except when you get up and interrupt each other ("no no, here, let me do it", etc.). Even if what the interrupter had to say was really useful, the interrupter was often impolite.

> They need to be more disciplined: interruptions and debates at the blackboard can be very tiring.

Sometimes we forgot to be polite, although the bell did provide a more civilized way of interrupting than we had managed in our previous team-taught courses. Often the interruptions involved debate about issues that were important, but the contexts of our disagreements were not properly explained to the students, and this may explain their discomfort.

In fact, we are even now in disagreement over how much disagreement we should exhibit.

> Jane says "Generally speaking, where our interruptions were intellectually motivated, they were part of an ongoing dialog between the three of us, and were therefore not as distressing to us as to the audience. This dialogue begins in the planning stage of the workshop, long before we meet the students, and continues during the workshop and after it ends. Our debates involve not only substantive issues but also our different perspectives on how we think about mathematics and what we each see as important. In fact I believe that sharing debates over such issues is one of our *gifts* to the students, even if they do not always perceive it that way."

> David says "The students were not invited to participate in the debate, and all they saw was the conflict. The interruptions were not only motivated

> by intellectual disagreement about the appropriate approach to a mathematical topic; there were plain errors by one or the other of us that required correcting. And some of them took place because the other two of us felt that the lecturer was losing the audience."

Student Reactions

Evaluating educational experiments is difficult: anecdotal evidence is partial and subjective, statistical evidence often concerns itself with spurious or misleading measures, and most student evaluation forms concentrate on the teacher instead of on what the student has been stimulated to achieve.

So what evidence do we have? The most significant body is the notes produced by the students. They go well beyond what we presented, and include many original approaches and insights. (These notes probably deserve publication, but this would require a significant editorial effort, and we don't know if it will get done.) Also, notes are often written by the outstanding student in a graduate course and form the basis of his or her future career as a mathematician; but here they were produced with nontrivial contributions from a big fraction—probably more than 75 percent—of the participants.

Another way of determining our success is by looking at the student evaluations, which were solicited spontaneously by one of the participants right at the end of the workshop. She circulated the following questions:

> 1. What parts of the workshop were the most useful for you (if it was at all)?
>
> 2. How well do you think it worked to have three teachers?
>
> 3. What do you think about the level of formality of the mathematics?
>
> 4. How do you think it would work in other settings?
>
> 5. If there was another workshop like this, what would you keep the same, and what changes you would make?
>
> 6. Would you come again?

All who addressed the last question (which means almost all the respondents) said they would like to participate in a similar workshop again.

Perhaps the question where there was most divergence was about the level of the workshop: most students thought it was just right but a significant minority felt it was either too easy or that we were assuming background that they didn't have.

We include here some excerpts from the responses. Out of fifty-plus participants, 24 returned evaluations; many people (perhaps those who were enjoying the workshop least?) had already left before they were asked to write their opinions on the last day. We are conscious of the possibility of bias in our selection of representative comments. Anyone wanting to see the full collection can consult `ftp://ftp.msri.org/pub/Misc/SummerCourse94/Evaluations`.

> …This has been undoubtedly one of the best mathematical experiences that I have had. … [An] excellent thing was the huge amount of examples. … Sometimes a good example is better than a proof. … [This workshop] recharg[ed my] motivation and enthusiasm in mathematics.

> …Presenting a few questions and leaving it to us to work it out proved to be a good way to start group discussions and cooperation. I think the idea that there is a section that the students are completely in charge of is excellent. It reinforces the idea that the students are responsible for their learning and gives them some freedom. I plan to use this in my teaching.

> …The last few days, with the more advanced lectures, we accumulated a peanut gallery of (I suppose) advanced grad students or members of the institute. They were annoying, talking during lectures and diverting lecturers with questions no one else understood.

> …I've had some hyperbolic geometry before, but never in a setting that required me to grapple so hard with the concepts.

…The most important aspect of this workshop was getting to see hyperbolic geometry from so many points of view. It was also good that there was a mixture of algebra, topology, and complex analysis (etc.) so that one could approach this subject from various backgrounds. It was also good to see these areas of mathematics all working together in one area. (This is quite different from the classroom approach.)

…Unfortunately, so much was presented, I soon became frustrated with the task of trying to make even some of the basic connections. … I'm left with a general sense of having glimpsed at a vast and beautiful subject, but very little that I didn't already know is sticking with me.

…What I liked best about the whole workshop: The atmosphere. The feeling of community amongst the students, amongst the teachers, between teachers and students, the theme of cooperation not competition, the way that the students were empowered with the ability to make suggestions and changes and come up with their own things they want to do. … Least useful: Long lectures that I hadn't figured out the background for.

It would have helped to have a written supplement with a higher level of formality. I think the informality and familiarity with which the teachers treated the material and each other made it much easier for us to attack these hard problems ourselves. It made math feel more like a natural activity, and successful mathematicians seem more like people I could emulate. I felt comfortable (comparatively) exposing my ignorance to the light that could obliterate it. It made learning faster, more enjoyable, and more complete.

The workshop was useful in several different ways for me. The constructions (with paper and with models) gave me a much better sense of the subjects being covered. … I think it would have been nice to have more activities/constructions later in the workshop which tied in with earlier ones (which were wonderful). I also think that allowing and encouraging the students to organize their own activities (e.g., note taking, afternoon discussions, advanced topics lectures) was useful. … If I didn't understand a morning session there was always a group of people to talk to about it.

I think it was one of the best educational experiences I've had [and] will have a lasting effect on my approach to research. I had fallen into a trap without really realizing it. I was unwilling to formulate my own questions about mathematics and try to answer them. I would only try exercises that were assigned or suggested, thinking all other questions would be too hard to approach. For this reason, I particularly liked the open-minded questions we were asked. It allowed us to look at a situation, study it, formulate questions, and answer them (sometimes)… Formal proofs are available in the literature but communication of the thoughts behind the argument is more rare.

Rewards and Regrets

What did we most enjoy about the workshop? The enthusiasm of the students; explaining things to people who were really interested in the answers and who had already spent a lot of effort and knew exactly what they needed to find out; new insights on matters we thought we already understood fully (like inversions); seeing the students present their own work; learning from their fresh slant on many familiar problems; realizing for the first time that one can build truly complex models out of hyperbolic paper, that one isn't limited to the pseudosphere and other examples with cyclic infinite fundamental group; and the feeling that we were doing something that the students really valued.

Above all, because the students seemed to take our ideas seriously, a wonderful atmosphere evolved. There was tremendous energy going into discussions during the morning coffee break, during lunch and the long afternoon break, and during afternoon tea. Students seemed to be really talking to each other, listening to each other, helping each other.

What we would have done differently:

We lectured too much during the second week. At all times, we should have stuck to our plan of keeping periods of lecture to twenty or thirty minutes, interspersed with student activity. We managed this at first, but we didn't hold to it, in part because there was so much we wanted to talk about. This would probably have been less of a problem if the workshop had been three or four weeks long. On the other hand, we all were tired by the end of the second week; even the percentage of students that took active part in the student-run sessions had dropped. The level of intensity of the first one and a half weeks simply could not be maintained indefinitely.

One of our goals was to present the same material from different points of view. In some cases we did so, and it worked. But in other cases we ended up presenting a sequence of arguments where the first part was from one point of view, the second from another, and so on. This was too cryptic; the students aren't necessarily able to fill in the complementary points of view for all the parts of the material.

We should also have stuck with the idea that we were making minimal assumptions about students' backgrounds. In mathematics, we more often err by introducing sophisticated concepts too soon than by dwelling on the basic definitions and concepts too long. We followed tradition in making the former mistake. Even people who have already learned more sophisticated material usually profit from thinking about basics.

We wish we had had more discussion about the processes of teaching, learning, communicating, and understanding mathematics. It's hard to take the time to discuss "process", but it would have been worthwhile and it might have reassured those who were uncomfortable with the process. Also, since the participants in the workshop were mostly teachers and future teachers at the university or college level, some discussion would have been valuable for their professional lives—especially as many of them reported later in their comments that they planned to introduce some of these methods into their own teaching.

We wish also that we had had a more formal discussion about how we think and understand mathematics, particularly about different kinds and different levels of understanding.

Conclusions

The content of a mathematics course is traditionally defined as what the lecturer writes on the blackboard and what she says in front of the class. (Other definitions are the official departmental syllabus and the collection of examination questions asked over the past five years. They tend in practice to be respectively much wider and much narrower than the first definition. Possible explanations of this phenomenon are disturbing.)

The traditional point of view has some value, particularly at the graduate level: if the lecturer can cut a path through the thicket of elaborate technical know-how, the student can more quickly meet the real challenges of the subject matter. Nevertheless, we deceive ourselves if we do not concentrate on what the students do, rather than on what the lecturer does.

Our aim was to involve each student deeply in some of the material. We believe we amply achieved this objective, in spite of all the reservations expressed in the preceding section. In any course there are too many ideas and too many loose ends for any one person to pursue. To us, it was most important that people be engaged deeply in *something* that is interesting to them, not that they absorb *everything*. This contrasts with the usual classroom expectation that students are supposed to acquire a meager grasp of everything presented.

In fact, this experience has led us to wonder whether it might be better to reorganize some learning of mathematics into periods of total immersion, a technique known to be very effective in the learning of languages. Could successive periods of total immersion in different areas of mathematics be more effective than simultaneously following several different courses, each lasting a semester? Our experience is that it is ultimately more valuable for students to learn fewer things more deeply than to gain a superficial knowledge of many things. We feel that something like an immersion approach would be worth trying on an experimental basis.

Would we do it again if circumstances allowed? A resounding "Yes".

References

[1] M. Phillips, S. Levy, and T. Munzner, *Geomview: An interactive geometry viewer,* Notices Amer. Math. Soc. (1993), 985-988.

[2] Charlie Gunn and Delle Maxwell, *Not Knot* (video), A K Peters, Wellesley, MA, 1991. Produced by the Geometry Center, University of Minnesota, and available from A K Peters, phone (617) 235-2210.