Memoirs from a Small-Scale Course on Industrial Math

David G. Schaeffer

Introduction

Overview

In this article I describe my experiences from a graduate course on industrial math that I taught at Duke in the spring semester 1995. When I volunteered to teach this course, several months before the beginning of the semester, it seemed like a very manageable task; but as the time to start teaching drew closer, I grew apprehensive that I would be unable to develop a successful course in this rather unfamiliar area. My worries were all the more severe because (i) I was operating solo, without any colleagues at Duke working on industrial problems; (ii) I hadn't had significant contact with industrial problems for several years; and (iii) my skills on the computer are rather limited. In the final analysis, I believe the course was very successful, as I will explain below. In writing this article, I hope that others who may be considering such a course may profit from what I learned. Moreover, I want to add a personal note: I thoroughly enjoyed the course, and I plan to teach courses in this format regularly in the future.

Because of my own background, I restricted the focus of the course to problems with a PDE/analysis slant. With some digging I found that industrial math has many challenging problems of this type which, because of both their intellectual content and their obvious relevance to a larger social context, provide an excellent vehicle for learning techniques of applied math. Moreover, understanding the current job market, students appreciated the broader training such a course offers, and they also liked the opportunity for active participation included in the structure of the course.

For classes meeting twice per week (lasting 75 minutes), Duke's semester includes 29 meetings. These were divided among the following four activities:

i) lectures by me (16 classes)

ii) lectures by students (5 classes)

iii) student research (6 classes)

iv) guest lectures (2 classes)

(Regarding the difference between (ii) and (iii): in (ii) each student gave an expository account, without any attempt at a solution, of an industrial problem, while in (iii) student teams worked over a period of several weeks to solve a problem, eventually presenting their results in a lecture.) This article is divided into four sections which discuss these activities in sequence.

One important conclusion may be stated even before further discussion: in student research (iii) direct contact with the proposer of an industrial problem is a crucial aspect of students' learning experience.

Prerequisites and Enrollment

Despite my not advertising, a course on industrial math generated much interest around the university. Before the start of the semester a half dozen students from engineering and computer science came to me to inquire about the course, and ten people came to the first class,

David G. Schaeffer is professor in the Department of Mathematics, Duke University, Durham, NC. His e-mail address is `dgs@math.duke.edu`.

five of them from outside mathematics. However, the prerequisite that I set for the course—the undergraduate PDE course covering separation of variables—served to squelch outside interest. Although I set the prerequisite at what I regarded as a low level in order to include people from outside the department, no one from outside the department stayed beyond the first class, and several of the people making inquiries told me explicitly that the prerequisite excluded them. In the end the enrollment consisted of five graduate students, all in mathematics.

Lectures by Me

The sixteen lectures I gave were allocated as follows:

	TOPIC	NO. CLASSES
(a)	Overview of the course	1
(b)	Survey of an industrial problem	1
(c)	Techniques of applied math	8.5
(d)	Presentation of problems for student research	1.5
(e)	Optimal control	4
		16

The lecture on topic (b) was intended to help students give their lectures, as described in the section "Lectures by Students" below. Essentially for one class I pretended I was a student, and I gave a lecture of the kind I hoped they would give.

In the fairly standard lectures on topic (c), based on Lin and Segel [3], I covered nondimensionalization, scaling, asymptotic series, singular perturbations (both boundary layers and rapid oscillations), and similarity solutions. I also included a case study from my own research to illustrate the use of singular perturbation theory.

The lectures on topics (d) and (e) both related to student research, which is discussed more extensively in the "Student Research" section. In (d) I surveyed the list of industrial problems from which I asked students to choose one for research. In (e) I covered a mathematical topic that was needed in attacking one of the problems that the students chose for their research. (The choice of this topic was one of several aspects of the course which became definite only as the semester evolved.)

The various topics were covered in the order listed above, except that the survey of industrial problems, topic (d), was sandwiched in the middle of topic (c), following the fifth of the eight-and-a-half lectures. (See Table 1 in "Student Research" section for more information about scheduling.)

Lectures by Students

Each student gave one 75-minute expository lecture on an industrial problem of his/her own choosing. I asked the students in their lectures to address the following points: (i) the physical background of the problem, (ii) formulation of a mathematical model, (iii) difficulties of the mathematical model, (iv) partial results, and (v) open problems. No solutions were expected in this part of the course.

Students chose problems to lecture on from Avner Friedman's wonderfully stimulating series [2], each article of which summarizes a talk given at the Industrial Math Seminar at the IMA. To assist students in choosing, I selected approximately 70 DE/analysis problems from the more than 100 problems in the six volumes (now seven), and I grouped these into broad categories containing 5–15 problems. These included both physically based categories (e.g., fluid flow, semiconductors, image processing, electromagnetism, and optics) and mathematically based categories (e.g., diffusive phenomena, free boundary value problems). I expected that students would first choose a category of interest to them and then choose a specific problem from that category. In actuality I'm not sure the students' interests were sufficiently well developed for them to make the initial choice of category; thus they struggled with choosing from a very long list of problems.

Although there are other references on industrial math problems (e.g., [4] as well as less formal proceedings from RPI and the IMA), we used only Friedman as a source for student lectures. This worked out well because of the great variety of problems contained in his books and because of the brevity of his treatment. Brevity especially was desirable for two reasons: (i) it made the task of lecturing on a problem seem less formidable to the student, and (ii) it meant that to give a coherent lecture the student had to fill in missing details. Indeed, in some cases, the student and I working together could not fill in such a detail, and we had to contact the original proposer of a problem (each article mentions the proposer). Although I was nervous about contacting complete strangers, in fact all these people responded warmly and helpfully to interest in their work from the academy.[1]

[1] *It was not always easy to locate the proposer. In one case the proposer had left a major corporation to found a successful consulting company which exploits the techniques on which he spoke at IMA; with the help of a friend at his former employer, we did reach this man. In another case we never did locate the proposer: since the time of his lecture at IMA he had gone to work for a defense contractor which went out of business following the collapse of the former Soviet Union.*

I took two measures intended to help students profit from others' lectures as well as their own. (i) Each student actually was involved with two problems, one as lecturer and one as "support person" for the lecturer. Both persons associated with a problem were responsible for understanding it. The support person was available to the lecturer to discuss questions while he/she prepared and was present at the lecture to help clarify issues if he/she got stuck. (ii) I asked all the students for each lecture to hand in brief written answers to a short list of questions about the lecture. (This list is reproduced in the appendix.) Here is some amusing evidence that these questions had their intended effect: At one of the lectures a student in the audience asked a question, was unsatisfied with the answer, and repeated his question with more emphasis while adding, "I need to know because I have to answer questions about this lecture."

I liked the results of the student lectures very much. They were lively, enjoyable affairs, and in the end-of-the-term evaluations students said they learned much from them. One student even suggested that the earlier lectures on mathematical techniques (the ones I gave) be given by students, including the lecturer's assigning problems on the material covered. The experience of choosing a problem on which to lecture, although difficult, was also instructive for students. They had to read several articles and evaluate them on a very specific criterion of great importance to them: "Do I understand this article well enough that I can speak on it in front of my colleagues without embarrassing myself?" Some students thought choosing a problem would have been easier if I had given them *less* time to choose.

My one regret concerning this part of the course is that I should have given students more help in preparing for their lectures. I required lecturers to meet with me once, at least a week or more before the lecture, and I invited them to meet as many times as they desired. If at the required meeting no problems seemed apparent to me, I did not insist on a second meeting. I tended to trust the student's judgement as to whether his/her understanding was adequate. In retrospect, I think I needed to take a more proactive role, quizzing students almost as in a preliminary exam, since in some cases weaknesses in understanding surfaced at the lectures.

Student Research

In the first subsection below I describe how student research was brought into the course and evaluate the results in general terms. In the second I discuss the mathematics of the two specific industrial problems that students worked on.

The Process

I planned to divide the class into teams of three people, each team working on a separate problem. Of course in a class of five students, that plan meant there were two teams, one of which had only two members. I preferred students working as part of a team rather than individually because: (i) by dividing a task, a team could tackle a bigger problem; (ii) as part of a team, students could learn from discussions with colleagues; and (iii) team research is the norm in industry.

To find appropriate problems for student research, I sent by e-mail or regular mail approximately twenty letters requesting problems. I contacted colleagues at universities with a developed industrial program (Claremont, IMA, Oxford, RPI), academic colleagues who have worked individually on industrial problems, colleagues in industry, and grants officers. (I am grateful for the generous responses of the people listed in the acknowledgments.) During the semester another good source for research problems surfaced: four of the five student lectures (cf. "Lectures by Students" section) raised open questions that I considered appropriate for students to work on in the course. (The fifth lecture also raised interesting open questions, but since they were in computational fluid dynamics, they were much too ambitious for the course.)

Before the start of the semester, my greatest anxiety concerned not being able to find appropriate problems. As it happened, I found there was a wealth of such problems. Indeed, I chose to cut out some good problems in order to simplify the students' choice.

About one-fourth of the way through the semester, I handed out a list of nine problems for research, and I spent a lecture and a half surveying them. Of course I discussed the mathematical problem and its physical background; I also evaluated its level of difficulty, its importance, and the skills needed to make progress on it. I included some easy problems on the list. For example, one problem involved only a routine computation with Fourier series; it seemed virtually certain to me that the students could solve it in the few weeks available to them. Another problem involved a straightforward application of bifurcation theory; the principal challenge would have been to understand a poorly written applied paper and to abstract a precise mathematical model from it. Other problems were more challenging and potentially more rewarding. One problem had a first step appropriate for the course and a second step that would have been an excellent thesis topic related to my own research. (Despite my emphasizing this incentive, no one chose to work on it.)

I let the students choose their own problems, including the composition of the team working on it. I regarded this choice as an important learning experience for them. I was gratified that they avoided the easy problems, even expressed contempt for them.

In the introduction, I listed 6 classes as being devoted to student research. This included (i) 2 class periods during which I cancelled the ordinary class meeting and instead used the time to meet with one of the student teams to discuss their problem and their work on it and (ii) 4 class periods during which students reported on their work to the others (two lectures per problem). In fact, these numbers underestimate the time I devoted to this part of the course. I spent 7 or 8 additional hours meeting with students to discuss their research, plus significant private time thinking about the problems. I hasten to add that I enjoyed this part of the course most of all; among other things, I felt that real learning was taking place.

In general terms, I believe the student research part of the course was a great success as regards students' learning but was incomplete as regards delivering something of value to the proposer. The latter situation reflected the fact that there was not enough time in the semester to complete the projects. Indeed, I had originally planned for students both to lecture and to submit a short paper on their work; in practice, they were under tremendous time pressure simply to complete enough of their work in order to give a meaningful lecture on it. (I jokingly justified my error in not allowing more time for research by saying that I was trying to reproduce industrial conditions where somebody above you sets unrealistic deadlines.)

Here is anecdotal evidence that contributes to my feeling students learned greatly from the research component of the course. It concerns a student whose lecture on an industrial problem from Friedman (cf. "Lectures by Students" section) was very formal and greatly lacking in intuition, so much so that I wondered to what extent he himself knew when he understood a problem. At the end of the term to present the background for his team's project, he had to understand and explain a lot of new physics, and he did a much better job the second time around. Moreover, best of all, following the talk, he commented in private, "I think I was convincing when I understood the topic, but not when I didn't."

In the future, when teaching this course, I will seek to provide adequate time for research by (i) starting the projects earlier and (ii) creating a mechanism to continue the research beyond the end of the semester.

Regarding (i), as indicated in Table 1, I began the course by lecturing for 12 classes, and then

ACTIVITY	NO. CLASSES	
Lectures by me [a]	actual	proposed
Lectures by students	12	4
Deadline for choosing research problem	5	5
Lectures by me [b]		
Student research [c]	4	12
Guest lectures	2	2
Student research [d]	2	2
	4	4
	29	29

Table 1: Use of classroom time

[a] *In the actual schedule I covered topics a, b, c, and d listed in "Lectures by Me" section. In the proposed schedule I would cover topics a, b, and d.*
[b] *In the actual schedule I covered topic e listed in "Lectures by Me" section. In the proposed schedule I would cover topics c and e.*
[c] *Preparation and conferences—see "Student Research" section.*
[d] *Lectures on results—see "Student Research" section.*

I scheduled student lectures on problems in Friedman for the next 5 classes. I required students to choose a research problem following the last of these student lectures. I had originally planned an earlier deadline. However, once I saw that the problems from Friedman provided fertile source material for student research, I postponed the deadline to this late date so that, if they desired, students could explore open questions from one of the lectures as their research project. Thus there remained only four weeks between problem selection and the first lecture to present results, too short by a substantial margin.

With hindsight were I to teach the course again, I would revise the schedule as indicated in the last column of Table 1 in order to allow adequate time for student research. Briefly, I would postpone the eight lectures on techniques of applied mathematics until after students had selected their problems, so that students would be pursuing their research concurrently with these lectures. With this revision of the schedule, there would now be *eight* weeks for research. It might be possible to gain additional time for research by having students lecture on problems from the list for research. However, a good lecture problem needs to be readily accessible, while a good research problem needs to be up-to-date, and these may be conflicting desiderata.

Regarding (ii) some work on these two problems has continued on an informal basis. In particular, the team for Problem (b) (cf. discussion below) is currently preparing a report of their findings[2], and one student on the team for Problem (a) is performing the checks needed to build confidence in the numerical solution found dur-

[2] *This represents an unusually high level of commitment in that two of the three students on this team chose to leave graduate school at the end of the semester.*

ing the semester[3]. In the future I will make these arrangements more formal: I will inform students from the start that I expect them to carry their projects through to completion, including a written report, the work continuing after the semester as independent study (for credit) as needed.

The Problems

I should mention that, in both cases discussed below, although the problems were of interest to the proposers, they were not central to their work. In particular, the proposers had never allocated the time for a concentrated attack on them.

(a) A Cooling Problem

This problem was proposed, in response to my solicitation, by an engineer working for a metal-producing company. The underlying issue concerns balancing two competing effects in quenching an alloy. On the one hand, one needs to cool quickly to prevent the two components of the alloy from separating. On the other hand, cooling quickly creates thermal gradients and hence stresses which may degrade the usefulness of the product. The mathematical problem, as I understood it following a brief conversation with the proposer and as I presented it to the students, consisted of choosing an optimal cooling rate $c(t)$ in a solution to the one-dimensional heat equation on $0 \le x \le L$, as follows. Given that the temperature $u(x,t)$ in the sample satisfies

$$\partial_t u = \kappa\, \partial_{xx} u \tag{1}$$

$$u(x,0) = T_0 \tag{2}$$

$$\begin{aligned} &(a)\partial_x u(0,t) = 0 \\ &(b)\partial_t u(L,t) = -c(t), \end{aligned} \tag{3}$$

choose $c(t) \ge 0$ to minimize

$$\sup_{0\le x\le L}\ \sup_{0\le t\le \tau} |\partial_x u(x,t)| \tag{4}$$

subject to the minimum-cooling-rate constraint that for all x, t with $0 \le x \le L$ and $0 \le t \le \tau$

$$\partial_t u(x,t) \le -r, \tag{5}$$

where T_0 and r are given constants and τ is chosen so that at time τ the specimen is essentially at room temperature. (Remark: The boundary condition (3a) comes from symmetry: (1–3) models half of a specimen that is being cooled at both ends.)

Although this problem seemed hard to me because of the L^∞ norms in both the objective function (4) and constraint (5), nevertheless one team chose it. Following their choice, in a series of phone conversations over a couple of weeks, the proposer and I developed a more satisfactory mathematical model. I described the original formulation of the problem rather completely in order to convey the spirit of the problem; however, I will summarize only briefly the reformulated problem, since its specific details are not important for our purposes here.

In one respect the new model was more complicated than (1–5): we found that more detail was required in modeling thermal stresses than was included in the simple formula (4). On the other hand, in three other respects the new model was vastly more tractable than (1–5): (i) A closer examination of the modelling, including the physical intuition that the greatest deterioration occurs at the centerplane $x = 0$, led to a simpler constraint analogous in form to

$$\int_0^\tau F(u(0,t))\,dt = C$$

for some function F and some constant C; i.e., an L^∞ constraint, both infinite-dimensional and nondifferentiable, was replaced by a one-dimensional smooth constraint. (ii) A similar combination of modelling and intuition simplified the objective function, so that like the constraint it depended differentiably on u. (iii) The heat equation was approximated by a finite-element model *with just two elements*; thus a PDE was replaced by a low-dimensional ODE. In addition, there were other changes which made the model more realistic but did not affect the mathematical difficulty of the problem. For example, the boundary condition (3b) was replaced by Newton's law of cooling

$$\partial_t u(L,t) = -c(t)\,[u(L,t) - T_r]$$

where T_r denotes room temperature and $c(t)$ is the control to be determined, restricted to the range

$$0 < \varepsilon \le c(t) \le M < \infty. \tag{6}$$

The improved model consisted of a standard control-theory problem for a 4×4 system of ODE, as discussed, for example, in Chapter II of [1]. The Pontryagin Maximum Principle, which I presented in four lectures following the students' choice of research problem, provides elaborate necessary conditions for an extremum. However, we did not use the full statement of

[3] *When the work on Problem (a) is completed, we will communicate it to the proposer orally; probably no written summary will ever appear. Interested readers may contact me by e-mail to see if a report on Problem (b) has become available.*

the Maximum Principle; rather, we invoked only a corollary of it which asserts that the optimal control for such problems has the so-called "bang-bang" form. Applied to the present example, this principle implied that for every t, the optimal control lies at one or the other endpoint of the interval $[\varepsilon, M]$ in (6). We then developed a plausibility argument that the optimal control had the form

$$c(t) = \begin{cases} \varepsilon & \text{if } 0 < t < t_1 \\ M & \text{if } t_1 < t < t_2 \\ \varepsilon & \text{if } t_2 < t < \tau; \end{cases} \tag{7}$$

in other words, the optimal control started on the low level and jumped only twice. Assuming this form, the optimal control could be found within the two-parameter family obtained by varying t_1 and t_2 in (7). The student team wrote a numerical routine to search this two-parameter family for the optimum[4]. In fact, the numerical process appeared to converge to a solution, but there was not enough time to run the checks that would be needed to have confidence in this solution.

In my view, from this work students acquired both specific mathematical skills and philosophical sophistication about applying math to "real" problems. Regarding the former, they learned about control theory, nondimensionalization and scaling, and numerical solution of ODE with a vividness that nonparticipatory courses cannot match. Although the gain in sophistication is harder to articulate, I regard it as of equal or greater importance. Among other things, students saw the need for understanding the physical origins of a problem before applying mathematical analysis.

Besides providing inadequate time, my main regret concerning this problem is that the proposer was unable to visit Duke. (He would have liked to, but time pressures at the company prevented this.) Since I did not feel the students would be able to abstract an appropriate mathematical model merely through phone conversations, I handled this part of the project. It would have been much better for the students to participate in the simplification, since this process is a crucial component of industrial math. Also, as I saw with the other team, meeting with a proposer offers invaluable experience in trying to communicate with a person with a different scientific background. Therefore, in the future I will propose a problem for student research *only if* the proposer is able to visit.

In my opinion, industry received some modest benefit from the students' work on this problem. In the first place, although industry had experimented with controls of the form (7), they were unaware of the Maximum Principle; now this work may be pursued with increased confidence. Also, having a simpler mathematical model may increase qualitative understanding of the issues. For example, it appears that if the plate being cooled is sufficiently thick, one needs to set the control $c(t) = M$ for all t, and even then it may be impossible to reduce separation of the components of the alloy to desirable levels. Finally, assuming the student does complete the computations, industry will receive quantitative information about the specific cooling rates which achieve the optimum.

(b) A Porous-Medium Problem

This project grew out of one of the student lectures on a problem from Friedman involving fluid flow in a porous medium. When seeking to answer one of the questions raised by the IMA lecture, we quickly encountered an issue about which we felt the need to contact the original proposer. As it happened, he had given his lecture at IMA six or seven years earlier, and since that time his research interests had changed greatly. However, the student team still wanted to work on flow in a porous medium. Therefore, I contacted a hydrologist here at Duke, and he proposed several related problems, one of which the students then worked on.

The problem concerns estimating properties of a rock aquifer by measuring the rate at which water pumped into a well diffuses into the aquifer. In mathematical terms consider the domain Ω in cylindrical coordinates

$$\Omega = \{(r, z) : a < r < \infty, -L < z < M\},$$

which parameterizes an infinite horizontal layer of porous rock, bounded from above and below by impermeable layers at $z = -L$ and $z = M$, respectively, and through which a well of radius a has been drilled. Water is pumped into the portion of the well with $-\ell/2 < z < \ell/2$, which has been sealed off from the rest of the well. The hydraulic head[5] in the rock, $u(r, z, t)$ where $(r, z) \in \Omega$, and the head in the sealed section of the well, $p(t)$, evolve according to the boundary value problem

$$\partial_t u = \kappa \, \Delta u \tag{8}$$

[4] *Although my computer skills are limited, I am conversant with numerical methods. Thus, given that most of the students had learned to program in a course on scientific computation the preceding semester, I believe I was able to provide adequate supervision for their numerical work.*

[5] *I.e., actual water pressure minus hydrostatic pressure $\rho g(z + L)$, converted to a length by division by ρg, where ρ is the density of water and g the acceleration of gravity.*

$$(9) \qquad u(r, z, 0) = 0$$

$$(10) \qquad \begin{array}{l} \text{(a) } u(a, z, t) = p(t) \text{ if } -\ell/2 < z < \ell/2 \\ \text{(b) } \partial_N u = 0 \text{ elsewhere on } \partial\Omega \\ \text{(c) } \lim_{r \to \infty} u(r, z, t) = 0 \end{array}$$

and

$$(11) \qquad \partial_t p(t) = R - B\,(2\pi a) \int_{-\ell/2}^{\ell/2} \partial_r u(a, z, t)\, dz$$

$$(12) \qquad p(0) = 0$$

where κ, R, and B are constants. Specifically, κ is proportional to the permeability of the rock and R is proportional to the pumping rate. The second term on the right-hand side of (11) represents the decrease in pressure resulting from water flowing into the aquifer; hence the coefficient B involves both the permeability of the rock and the compressibility of water. (Remark: The unsealed portion of the well is isolated from the rock by a collar so that the no-flow boundary condition (10b) holds there as well as at the impermeable layers.)

If

$$(13) \qquad L = M = \ell/2$$

(i.e., if the boundaries of the sealed portion of the well coincide with the impermeable layers), then the solution of (8–12) is independent of z and an analytical solution is possible. In practice this analytical solution is often used to estimate aquifer properties in situations where its justifying assumption (13) is not satisfied. The proposer asked: *To what extent is the analytical solution perturbed if L and M are greater than $\ell/2$?* In this case some vertical flow is superimposed on the primary horizontal flow.

The student team wrote a numerical routine to gather data about this question. Their work suggests that in the physically relevant parameter range

$$a \ll \ell \ll L, M,$$

perturbation changes the solution by approximately 10–20%. With this project also there was not enough time during the semester for them to make a convincing case that they had fully debugged their code.

In carrying out this work, students had to surmount a number of obstacles, all of which I regarded as significant learning experiences for them. For example:

- The proposer originally described the boundary conditions (10a, 11) in physical terms; the students struggled to translate his physical description into precise mathematical formulas.
- There was a compelling need to identify essential dimensionless parameters, since otherwise, as the students and I readily saw, the computations would have been prohibitive.
- To check the code, the numerical solution had to be tested against the analytical solution, which they obtained by separation of variables. The students learned, with some frustration in the process, that even a "small" mistake in either the numerical or the analytical solution invalidates such a test.

Regarding payoff, both the proposer and I feel the students' results are worthy of publication. However, if the students do not actually complete the work, then I would have to say that the proposer gave us more than conversely, since he met with the students for four or five hour-long sessions.

Guest Lectures

I originally scheduled three guest lectures, two by people from industry and one by an academic mathematician, not from Duke, who has worked on industrial problems. As it happened, one of the former had to cancel at the last minute for personal reasons, so there were only two guest lectures. Neither of these was the proposer of a problem for student research: one proposer was unable to visit, and the other was already here on the faculty at Duke.

In the future I will make two changes in this part of the course: (i) I will be sure that a nonresident proposer visits Duke, and (ii) I will invite the guest lecturers to speak about their working environment in industry as well as their mathematics.

Acknowledgments

I am grateful to the following people who contributed in various ways to this course: J. Abbott, R. Becker, E. Bissett, O. Bruno, E. Cumberbatch, D. Drew, A. Friedman, R. Goor, P. Hagan, A. Kapila, A. Nachman, J. Ockendon, E. B. Pitman, L. Reyna, S. Rojstaczer, D. Ross, J. Salzman, D. Schwendeman, and J. Trangenstein. Without their generous help the course would not have succeeded. Thoughtful feedback from Michael Shearer, Hugo Rossi, and an anonymous referee on a preliminary version of this paper improved it substantially.

References

[1] W. Fleming and R. Rishel, *Deterministic and stochastic optimal control*, Springer, 1975.

[2] A. Friedman, *Mathematics in industrial problems, Vols. 1–6*, Springer, 1988–1994.

[3] C. C. Lin and L. Segel, *Mathematics applied to deterministic problems in the natural sciences*, SIAM, 1988.

[4] A. Tayler, *Mathematical models in applied mechanics*, Clarendon, 1986.

Appendix

For each student lecture, I asked the other students to hand in brief written answers to the following eight questions:

- Describe the physical background of the problem.
- Formulate the mathematical model occurring in the problem. In general terms, classify its nature (e.g., nonlinear elliptic PDE with a free boundary condition).
- What are the essential mathematical difficulties of the model?
- Summarize the partial results that have been obtained so far.
- What open questions are raised by this problem?
- Are there any general lessons that may be extracted from this problem?
- What things did you like about the lecturer's presentation?
- What, if anything, might have been done differently to make the lecture more understandable or enjoyable for you?